StarTools Manual

Based upon StarTools V1.8

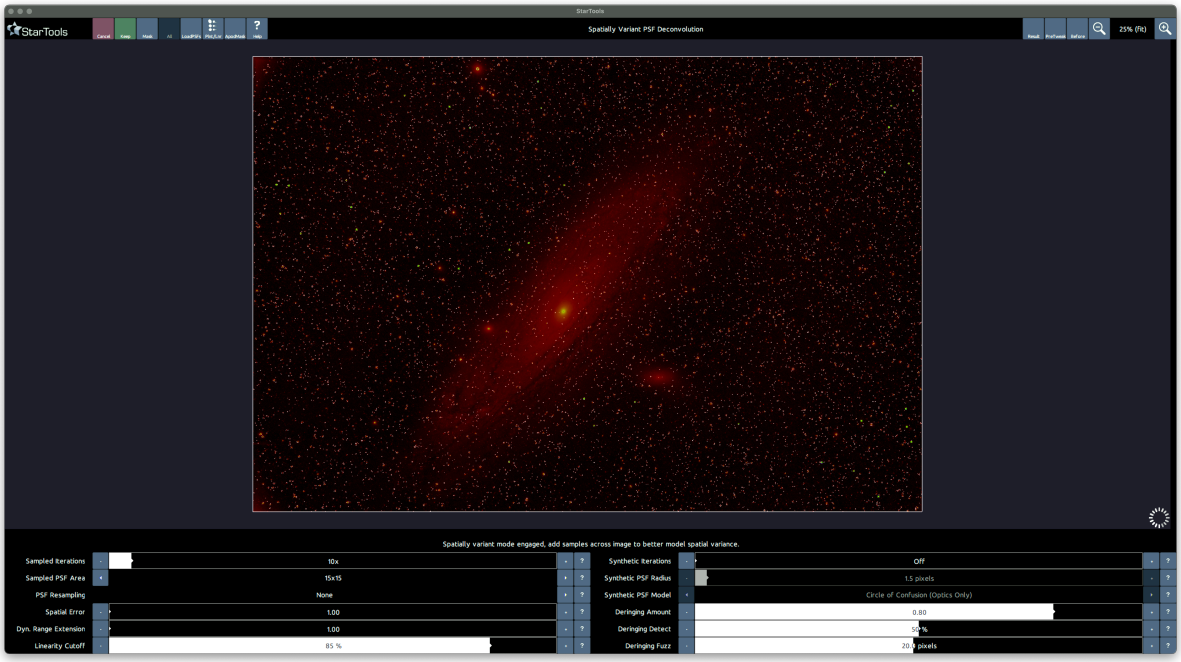

Issue 1.8.28 September, 30th, 2022

About this Manual

250+ pages? Don't panic! You will find what you need in a blink with the new structuring. Guy's excellent User Notes content has now been woven into Ivo's online material. Beginner or advanced user, you may use the tutorial section in progressing order:

'Core Workflow Step by Step' (page 22).

A new 'Tutorials' section was created, including a step by step guide based upon the workflow drawings and User Notes' methods, running you through all the modules of the basic workflow and more. This makes for a great StarTools crash course for beginners.

'Special Techniques' (page 33)

Also included in 'Tutorials' is an extra chapter providing training content for advanced learners. These 'how-to's' collected by Guy cover a lot of frequently reappearing questions on the forum. Be sure to check the content (next page) if you have any inquiry - it might well be answered already. You are welcome to send in any further methods you discovered/developed so they may be added in here.

'Modules & Features' (page 69)

Ivo's in-depth descriptions discuss tools and parameters in detail and are now complemented by a 'Wrap Up' section of Guy's User Notes to provide a module profile.

'Background Information' (page 209)

A newly created section provides further knowledge and in-depth explanations about topics frequently coming up on the forum.

'Starting with a good dataset' (page 47)

If you are new to StarTools, please check this first in order to skip frustrating trial&error attempts!

'Tracking' (page 60)

To better understand StarTools's approach, please check Ivo's 8 pages on Tracking. These are instrumental to understand why StarTools 'feels' so fundamentally different from other software and grasp the benefits of getting used to a different kind of workflow.

This manual includes:

- All sections of the 'About' tab on StarTools Homepage (Author: Ivo Jager)
- 'Quick Start Guide' of StarTools Homepage (Author: Ivo Jager)
- All sections of the 'Modules' tab on StarTools Homepage (Author: Ivo Jager)
- 'User Notes' on Forum (Author: Guy) - included in 'Tutorials', Wrap Up' and 'Background Information' sections
- Selected forum articles (Authors: Ivo Jager, Guy, RKonrad, elpajare, Mike in Rancho)
- Workflow Drawings (Jochen Scharmann)

This manual refers to V 1.8, including Guy's Notes content. It is subject to update so be sure to check for new versions.

Note: a german version of this manual is also available.

Clear Skies!

Jochen Scharmann

New in 1.8

Release 1.8.527

- The 'Spatially Variant Decon' module is the highlight of Star Tools 8. Its unrivaled functionality allows optimising the Point Spread Function **locally**, based upon multiple star samples.
- The brand new 'Narrowband Accent' module is providing an overeasy way to add narrowband accents to visual spectrum datasets. A seperate narrowband processing chain has been added to Luminance and Color paths in all modules, allowing indepenend but simultaneous processing.
- 'Repair' module has been made tracking aware, it may now be included in the workflow between 'Color' and 'Denoise'
- 'HDR' module is also tracking aware now, featuring a number of presets for common use-cases.
- A 'Correlation Noise Filter' in 'Wipe' module was added to remove any correlated effects such as debayer artifacts, so any downstream module may better focus on detail due to higher SNR.
- The new 'Protection' parameter plus enhanced deringing in 'Sharp' module will make it less prone to singularities, hence reducing artifacts.
- Further improvements include improved Tracking, improved Denoise and various code optimisations
- Native support for Apple M1 ARM CPUs.
- 'Layer' module now offers 2 new Filter mode options: 'Gaussian Highpass filter' and 'Laplacian over Gaussian 0 crossing' plus a new Cap mode 'Add 1/2 unity + clip'
- New 'AltStars' preset in AutoMask generator uses a star detection algorithm that is much better at detecting stars with full stellar profile included.

What is StarTools?

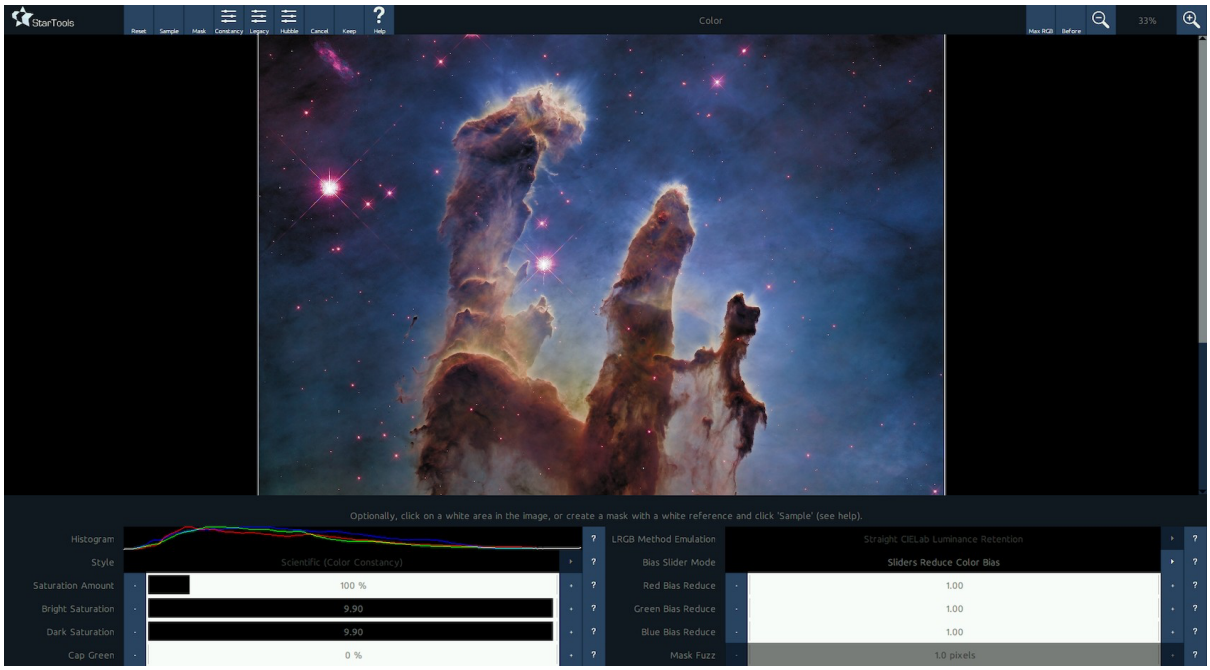

StarTools is powerful, yet very easy to use; no windows, no clutter, no process containers.

A new way of astronomical image and signal processing

StarTools is a new type of image processing application for astrophotography that tracks noise propagation as you process.

StarTools extensive knowledge of the past, present and - sometimes - future of your signal, allows you to do things users of other software can only dream of. These things include mathematically correct deconvolution of heavily processed data, mathematically correct color calibration of stretched data, and objectively the best noise reduction routine in the market that seems to 'just know' exactly where noise grain in your final image is located.

'StarTools is the best-kept secret amongst signal processing purists; those who fundamentally understand how StarTools achieves such superior signal fidelity.'

As opposed to other software, StarTools uses new brute force and data mining techniques, so your precious signal is preserved as much as possible till the very end. StarTools makes use of the advances in CPU power, RAM and storage space, replacing old algorithms with new, more powerful ones.

The best-kept secret amongst signal processing purists

StarTools is not just popular with beginners. StarTools is the best-kept secret amongst signal processing purists; those who fundamentally understand how StarTools achieves such superior signal fidelity. Yet, you don't need a mathematics or physics degree to understand the underlying theory; see the Tracking section to learn more.

We're incredibly pleased StarTools superior processing capabilities haven't gone unnoticed, now being the new tool of choice for a rapidly growing group of beginners, enthusiasts and institutions that numbers in the many thousands.

User friendly by mathematical nature

The software is 'user friendly by mathematical nature'. To be able to function, the engine needs to be able to make mathematical sense of your signal flow from start to finish. That's why it's simply unable to perform 'nonsensical' operations. This is great if you're a beginner and saves you from bad habits or sub-optimal decisions. It's not so much because we put 'guard rails' in, it's just that the application would break otherwise.

Interface

Don't be fooled by StarTools' simple interface - you are forgiven if, at first glance, you get the impression StarTools offers only the basics. Nothing could be further from the truth!

StarTools goes deep. Very deep. It's just not 'in your face' about it and you can still get great results without delving into the depths of its capabilities. It's up to you.

If you're a seasoned photographer looking to get more out of your data, StarTools will allow you to visibly gain the edge with novel, brute-force techniques and data mining routines that have only just become viable on modern 64-bit multi-core CPUs and increases in RAM and storage space.

If you're a beginner, StarTools will assist you by making it easy to achieve great results out-of-the box, while you get to know the exciting field of astrophotography better.

Whatever your situation, skills, equipment and prior experience, you'll find that working with StarTools is quite a bit different than most software you've worked with. And in astrophotography, that tends to be a good thing!

There is also a large UI version available

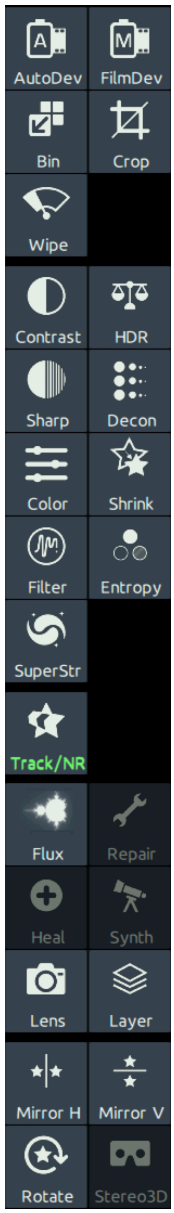

Example of the Main Interface

Navigation

Navigation within StarTools generally takes place between the main screen and the different modules. StarTools' navigation was written to provide a fast, predictable and consistent work flow.

There are no windows that overlap, obscure or clutter the screen. Where possible, feedback and responsiveness will be immediate. Many modules in StarTools offer on-the-spot background processing, yielding quick final results for evaluation and further tweaking.

In some modules a preview area can be specified in order to get a better idea of how settings would modify the image in a particular area, saving the user from waiting for the whole image to be re-calculated.

'Navigation within StarTools generally takes place between the main screen and the different modules.'

In both the main screen and the different modules, a toolbar is found at the very top, with buttons that perform functionality that is specific to the active module. In case of the main screen, this toolbar contains buttons for opening an image, saving an image, undoing/redoing the last operation, invoking the mask editor, switching Tracking mode on/off, restoring the image to a particular state, and opening an 'about' dialog.

In both the main screen and the different modules, a toolbar is found at the very top, with buttons that perform functionality that is specific to the active module. In case of the main screen, this toolbar contains buttons for opening an image, saving an image, undoing/redoing the last operation, invoking the mask editor, switching Tracking mode on/off, restoring the image to a particular state, and opening an 'about' dialog.

Exclusive to the main screen, the buttons that activate the different modules, reside on the left hand side of the main screen. Note that the modules will only successfully activate once an image has been loaded, with the exception of the 'Compose' module. Note also that some module may remain unavailable, depending on whether Tracking mode is engaged.

Helpfully, the buttons are roughly arranged in a recommended workflow. Obviously not all modules need to be visited and workflow deviations may be needed, recommended or suit your personal taste better.

Consistent throughout StarTools, a set of zoom control buttons are found in the top right corner, along with a zoom percentage indicator.

The icons in the top two panels roughly follow a recommended workflow.

Panning controls ('scrollbar style') are found below and to the right of the image, as appropriate, depending on whether the image at its current zoom level fits in the application window. Common to most modules is a 'Before/After' button, situated next to the zoom controls, which toggles between the original and processed version of an image for easy comparison. A 'PreTweak/PostTweak' button may also be available, which toggles between the current and previous result, allowing you to quickly spot the difference between two different settings.

All modules come with a 'Help' button in the toolbar, which explains, in brief, the purpose of the module. Furthermore, all settings and parameters come with their own individual 'Help' buttons, situated to the right of the parameter control. These help buttons explain, again in brief, the nature of the parameter or setting.

Zooming, panning and scaling

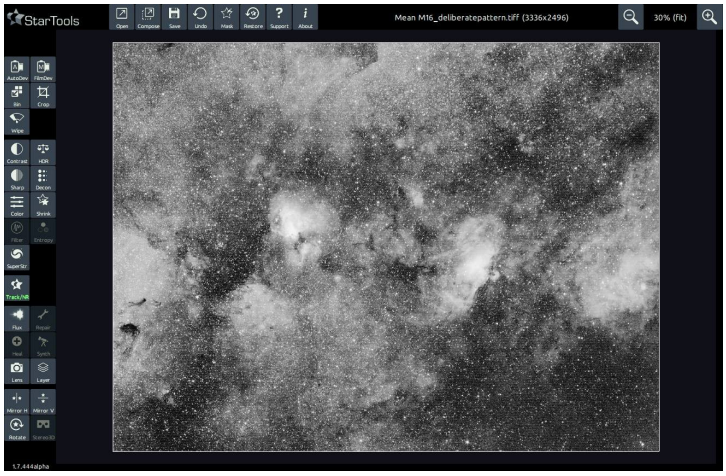

StarTools' astrophotography-optimised scaling algorithm can highlight latent pattern issues. It also was designed to show constant noise levels regardless of zoom level.

'StarTools implements a custom scaling algorithm in its user interface, which makes sure that perceived noise levels stay constant, no matter the zoom level.'

Even the way StarTools displays and scales images, has been created specifically for astrophotography.

StarTools implements a custom scaling algorithm in its user interface, which makes sure that perceived noise levels stay constant, no matter the zoom level. This way, nasty noise surprises when viewing the image at 100% are avoided.

Even more clever, StarTools scaling algorithm can highlight latent and faint patterns (often indicating stacking problems or acquisition errors) by intentionally causing an aliasing pattern at different zoom levels in the presence of such patterns.

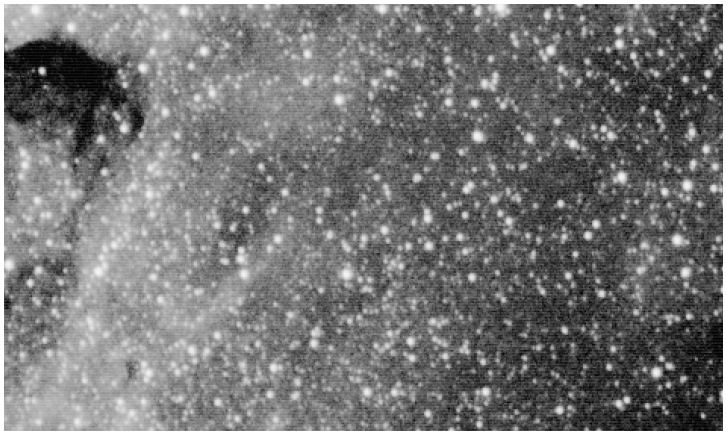

At 200% zoom level a barely distinguishable horizontal pattern can indeed be seen.

Changing parameters in StarTools

An example of a level setter control in StarTools

The parameters in the different modules are typically controlled by one of two types of controls:

1. A level setter, which allows the user to quickly set the value of a parameter within a certain range
2. An item selector, which allows the user to switch between different modes.

'The parameters in the different modules are typically controlled by one of two types of controls; A level setter, which allows the user to quickly set the value of a parameter within a certain range.'

Setting the value represented in a level setter control is accomplished by clicking on the '+' and '-' buttons to increment or decrement the value respectively. Alternatively you can click anywhere in the area between the '-' and '+' button to set a value quickly.

An example of a selector control in StarTools

Switching items in the item selector is accomplished by clicking the arrows at either end of the item description. Note that the arrows may disappear as the first or last item in a set of items is reached. Alternatively the user may click on the label area of the item selector to see the full range of items which may then be selected from a pop-over menu.

Hotkeys

As of version 1.5, StarTools implements some hotkeys for common functions;

Zoom out	**-** key	OK	**ESC** or **ENTER** key
Zoom in	**+** or **=** key	Blink before / after	**B** key
Zoom fit-to-screen	**0** key	Undo / redo	**B** key
Back	**ESC** key	Mask editor	**M** key
Cancel	**ESC** key	Open	**O** key
Done	**D** or ENTER key	Save	**S** key
Keep	**K** key	Screenshot	

Presets

Most modules come with presets that quickly dial in useful parameter settings.
These presets give you good starting points for specific situations, and for basing your own tweaks on.
Preset buttons can be distinguished by their icons; they bear the icon of the module you launched.

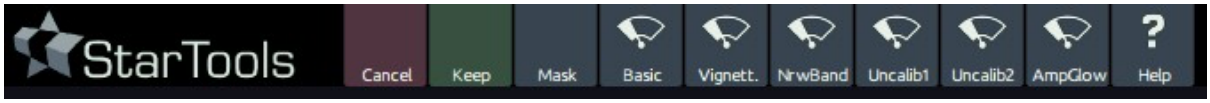

Mouse controls

As of 1.7, enhanced mouse controls are implemented;

Zoom in	Scroll wheel down
Zoom out	Scroll wheel up
Pan	Middle button + drag
Blink before/after	Right click

Touchscreen

StarTools can also be entirely operated by touchscreen with all controls appropriately sized for finger-touch operation.

Wrap Up

Main Window Purpose

To load the initial image(s), select modules to apply, change the mask, undo changes, and save the desired result.

When to use

- At the start of processing to load the files(s).
- To select the next module to apply (see below for example workflow).
- Between modules to set a mask or undo some changes.
- To turn Tracking off towards the end of the workflow.

Method

1. Open a file with the Open button, or combine multiple files with the LRGB button.
2. In turn, select the modules you want to apply (see above for example workflow).
3. Turn Tracking off. 'Denoise' at this stage.
4. Select any remaining modules required.
5. Save the desired result.

Description of Controls

Open

- Used to open single monochrome or color images:
- StarTools reads FITS, PNG and uncompressed TIFF files.
- StarTools works best with images that are linear and without preprocessing apart from registration and stacking
- The image should not have had any stretching, noise reduction, color balancing or deconvolution.
- For more information see Background Notes.
- After loading you are prompted about the image source. This will affect how you open the image in StarTools:
    - If it is a (stacked) RAW image - either monochrome or white-balanced color - select 'Linear, was not Bayered or is white balanced'.
    - If it is a (stacked) RAW color image that is not white balanced - select 'Linear, was Bayered, is not white balanced'. As of StarTools v1.5.366 standard OSC/ DSLR data that is not white balanced uses the Compose mode by default. See the Compose module notes for more details of the Compose mode.
    - If it is a JPEG or video source use 'Modified and not linear'.

Compose (v1.5 and later)

- Used to merge multiple files such as those representing Luminance and Red, Green, and Blue channels, or those representing the Hubble palette. Also allows creation of synthetic luminance, and different combinations of this data, e.g. LLRGB.
- See separate Compose module description for details on how to use it.

Save
Saves the current image:

- To save as a 16-bit TIFF file use the extension .tiff or .tif
- To save as a JPEG file use the extension .jpg
- To save as a 8-bit PNG file use the extension .png (StarTools v1.5+ only)

Undo
Undoes the effects of the previous module.

Mask
For general instructions on using the mask see Mask Module Use.

Track/NR
Turns off tracking and provides and opportunity to do the final 'Denoise' - See Tracking.

Restore
Allows us to return to the image at a previous stage in the processing - but with the option, for example, of reverting to Linear data but keeping Deconvolution intact - even though that was done after the global stretch.

- Original - revert to the original Linear, 'Binned and 'Cropped image.
- Linear, Wiped - revert to the Linear image after being Binned, Cropped and Wiped.
- Linear, Wiped, Deconvolved - revert to the Linear image after being Binned, Cropped, Wiped, and Deconvolved.
- Globally Stretched, Wiped, Deconvolved - revert to the image after being Binned, Cropped, Wiped, Globally stretched, and Deconvolved.

Modules
The modules are all accessed by clicking the corresponding button to the left of the image area.

- 'Lens' - Correct for lens defects.
- 'Bin' - Reduce resolution - improve signal to noise ratio.
- 'Crop' - 'Crop' image edges - remove stacking artifacts.
- 'Wipe' - Remove Gradient, light pollution and vignetting.
- 'AutoDev' - Automated global stretching.
- 'FilmDev' - Global stretching with more manual control.
- 'Decon' - Recover detail from seeing-limited data.
- 'Contrast' - Optimise Local contrast.
- 'HDR' - Optimise Localised dynamic range.
- 'Sharp' - 'Sharpen' detail.
- 'Flux' - Automatic feature recognition and manipulation.
- 'SuperStructure' - Light diffraction modelling.
- 'Color' - Correct and Adjust Color.
- 'Entropy' - Enhance local detail based on color (V1.5+).
- 'Filter' - Manipulate features based on color.
- 'Denoise' - Reduce noise.
- 'Layer' - Versatile pixel level manipulation.
- 'Heal' - Remove unwanted features.
- 'Shrink' ('Shrink') - Control star appearance.
- 'Repair' - 'Repair' star defects.
- 'Synth' - Augment or replace stars.
- 'Rotate' - 'Rotate' image.
- Mirror H - Make mirror image horizontally.
- Mirror V - Make mirror image vertically.

Tutorials

Quick Start Tutorial

A video is also available that shows a simple, short processing workflow of a real-world, imperfect dataset.

A Quick Generic Work Flow

Getting to grips with new software can be daunting, but StarTools was designed to make this as painless as possible. This quick, generic work flow will get you started.

While processing your first images with StarTools, it may help knowing that the icons in the top two panels roughly follow a recommended workflow when read top to bottom, left to right.

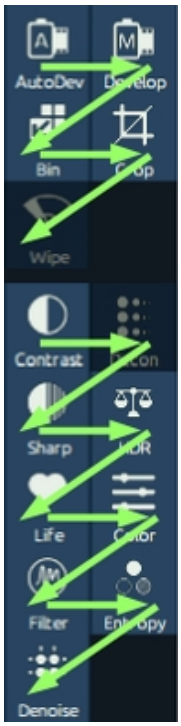

The icons in the top two panels roughly follow a recommended workflow when read left to right, top to bottom.

Step 1: Import, Start Tracking

Open an image stack ('dataset'), fresh from a stacker. Processing in StarTools is easiest and will yield vastly better results if the data is as 'virgin' as possible, meaning unstretched, not color balanced, not noise reduced and not deconvolved. Best results are achieved with data that is as close to what the camera recorded as possible.

Do not use any software that may meddled with your data prior to passing it to your stacking program. Avoid any pre-conversion tools or software that came with your camera. Make sure that any stacking software that you use is configured to perform as little processing to the data as possible. For example, if you use Deep Sky Stacker make sure that Per Channel 'Color' Calibration and RGB Channels Calibration are set to 'no'. Also make sure that, in Deep Sky Stacker, the final file is saved with settings 'embedded', rather than applied. 32-bit integer FITS files are preferable.

Counter-intuitively, a good stacker output will have a distinct, heavy color bias with little or no apparent detail. Worry not; subsequent processing in StarTools will remove the color bias, while restoring and bringing out detail. If, looking at the initial image, you are wondering how on earth this will be turned into a nice picture, you are often on the right track.

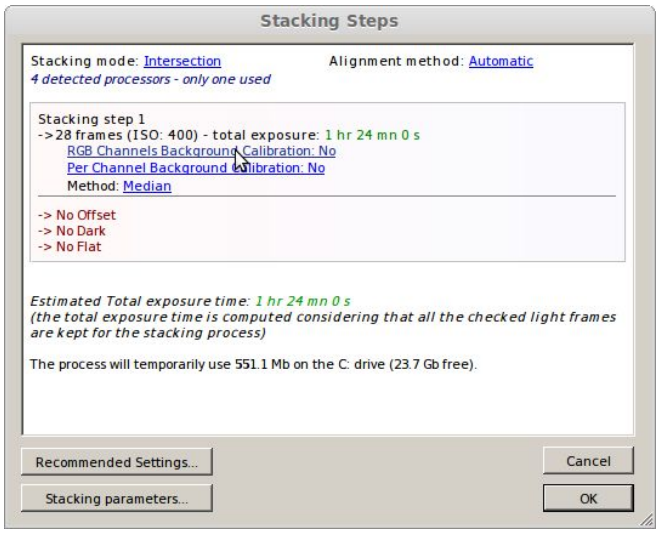

Upon opening an image, the Tracking dialog will open, asking you about the characteristics of the data. Choose the option that best matches the data being imported. If your dataset comes straight from a stacker, the first option is safe. Tracking is now engaged (the Track button is lit up green).

Step 2: Inspect Your Dataset

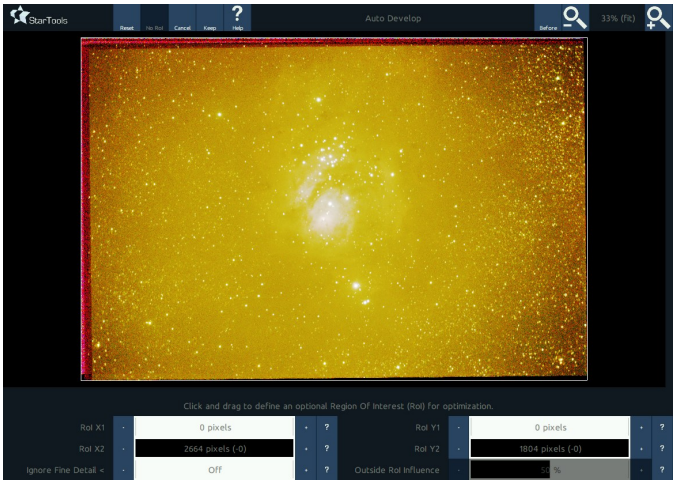

Launch 'AutoDev' to help inspect the data. Chances are that the image looks terrible, which is - believe it or not - the point. In the presence of problems in the data, 'AutoDev' will show these problems until they are dealt with. Because StarTools constantly tries to make sense of your data, StarTools is very sensitive to artifacts, meaning anything that is not real celestial detail (such as stacking artifacts, dust donuts, gradients, terrestrial scenery, etc.). Just 'Keep' the result. StarTools, thanks to Tracking, will allow us to redo the stretch later on.

At this point, things to look out for are:

- Stacking artifacts close to the borders of the image. These are dealt with in the 'Crop' or 'Lens' modules
- Bias or gradients (such as light pollution or skyglow). These are dealt with in the 'Wipe' module.
- Oversampling (meaning the finest detail, such as small stars, being 'smeared out' over multiple pixels). This is dealt with in the 'Bin' module.
- Coma or elongated stars towards one or more corners of the image. These can be ameliorated using the 'Lens' module.

Step 3: Prep

Fix the issues that 'AutoDev' has brought to your attention:

- Ameliorate coma using the 'Lens' module.
- 'Crop' any remaining stacking artifacts.
- 'Bin' the image up until each pixel describes one unit of real detail.
- 'Wipe' gradients and bias away. Be very mindful of any dark anomalies - bump up the Dark Anomaly filter if dealing with small ones (such as dark pixels) or mask big ones out using the Mask editor. You may also wish to use a mask to mask out nebulosity if using high values for the two Aggressiveness parameters.

Step 4: Final Global Stretch

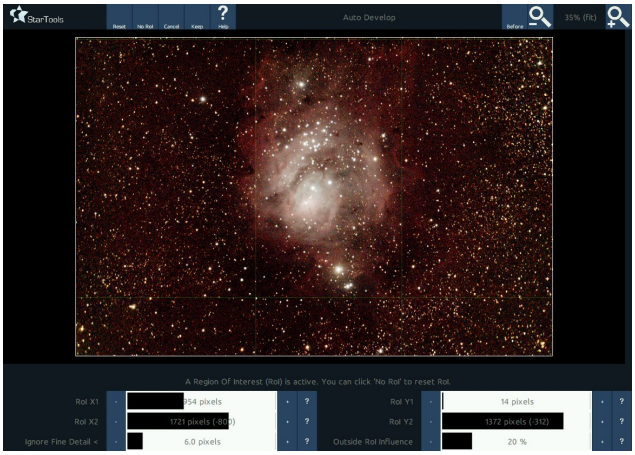

Once all issues are fixed, launch 'AutoDev' again and tell it to 'redo' the stretch. If all is well, 'AutoDev' will now create a histogram stretch that is optimized for the 'real' object(s) in your clean data. If your data is very noisy, it is possible 'AutoDev' will optimize for the noise, mistaking it for real detail. In this case you can tell it to Ignore Fine detail.

If your object(s) reside on an otherwise uninteresting or 'empty' background, you can tell 'AutoDev' where the interesting bits of your image are by clicking & dragging a Region Of Interest.

Don't worry about the coloring just yet - focus getting the detail out of your data first. If your image shows very bright highlights, know that you can 'rescue' them later on using, for example, the 'HDR' module.

Step 5: Detail Enhancement

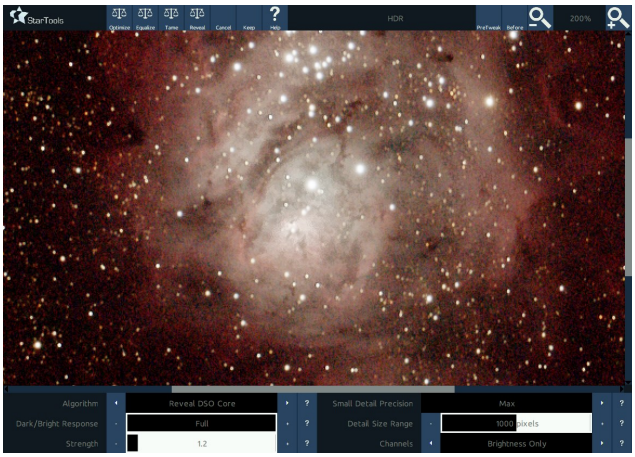

Season your image to taste. Apply deconvolution with the 'Decon' module, dig out detail with the Wavelet 'Sharpen' ('Sharp') module, enhance 'Contrast' with the 'Contrast' module and fix any dynamic range issues with the 'HDR' module.

There are many ways to enhance detail to taste and much depends on what you feel is most important to bring out in your image. As opposed to other software, however, you don't need to be as concerned with noise grain propagation; StarTools will take care of noise grain when you finally switch Tracking off.

Step 6: 'Color' Calibration

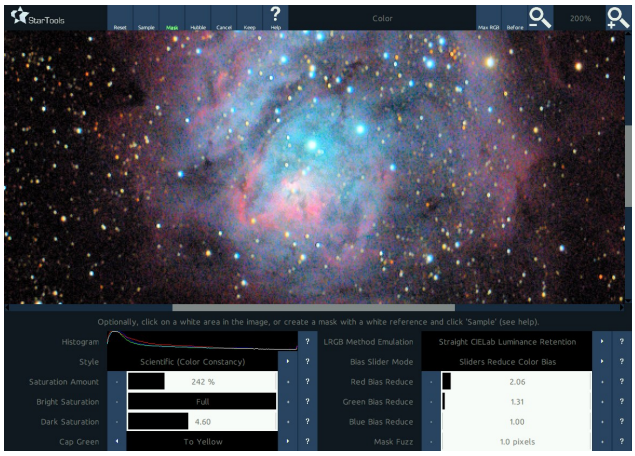

Launch the 'Color' module.

See if StarTools comes up with a good color balance all by itself. A good color balance shows a good range of all star temperatures, from red, orange and yellow through too white and blue. HII areas will tend to look purplish/pink, while galaxy cores tend to look yellow and their outer rims tend to look bluer.

Green is an uncommon color in outer space (though there are notable exceptions, such as areas that are strong in OIII such as the core of M42). If you see green dominance, you may

want to reduce the green bias. If you think you have a good color balance, but still see some dominant green in your image, you can remove the last bit of green using the 'Cap Green' function.

Step 7: Final Noise Reduction, Switching Tracking Off

Switch Tracking off and apply noise reduction. You will now see what all the 'signal evolution Tracking' fuss is about, as StarTools seems to know exactly where the noise exists in your image, snuffing it out. The most important parameters to tweak are Smoothness, in combination with Grain Dispersion.

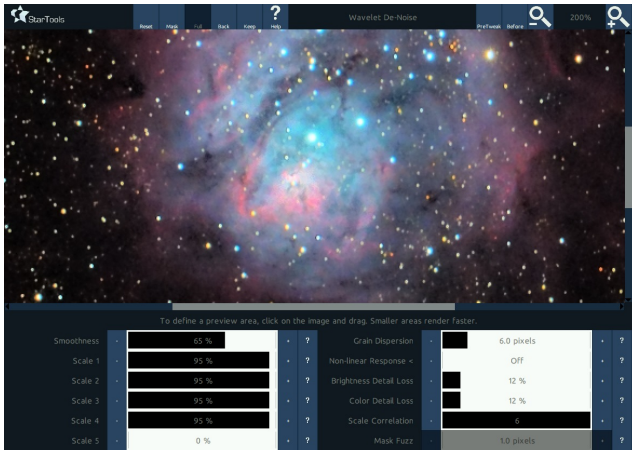

Creating the Stretch

High level overview on how the global stretch will be optimized using various modules (not to scale). 'AutoDev', 'FilmDev', 'Contrast' and 'HDR' modules, provide automated stretching at all levels of detail. These modules replace the traditional curve-based adjustment of stretching.

- Diagnostic Stretch ('AutoDev') will make gradients visible
- 'Wipe' will remove gradients, freeing up dynamic range down to dark anomalies limit
- Final Stretch ('AutoDev' or 'FilmDev') will increase dynamic range above of dark anomalies using lower mids freed up by 'Wipe'
- 'Contrast' will extend bottom range compressing dark anomalies and will optimize the stretch locally on large to medium scale detail
- Local stretch optimization on medium to small detail may follow using 'HDR'.

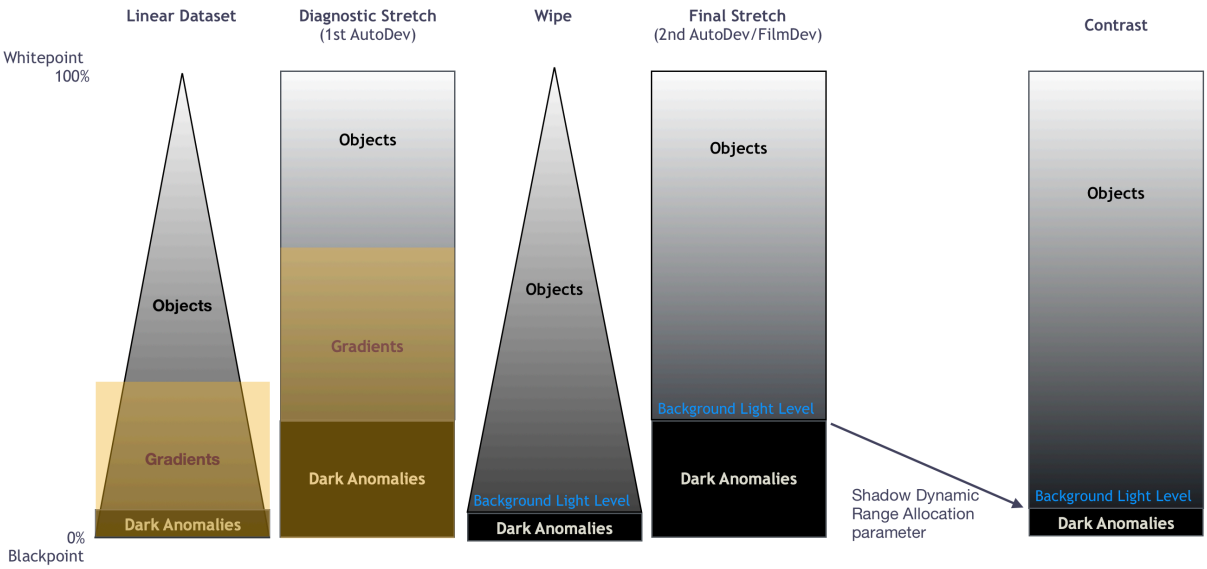

Core Workflow Step by Step

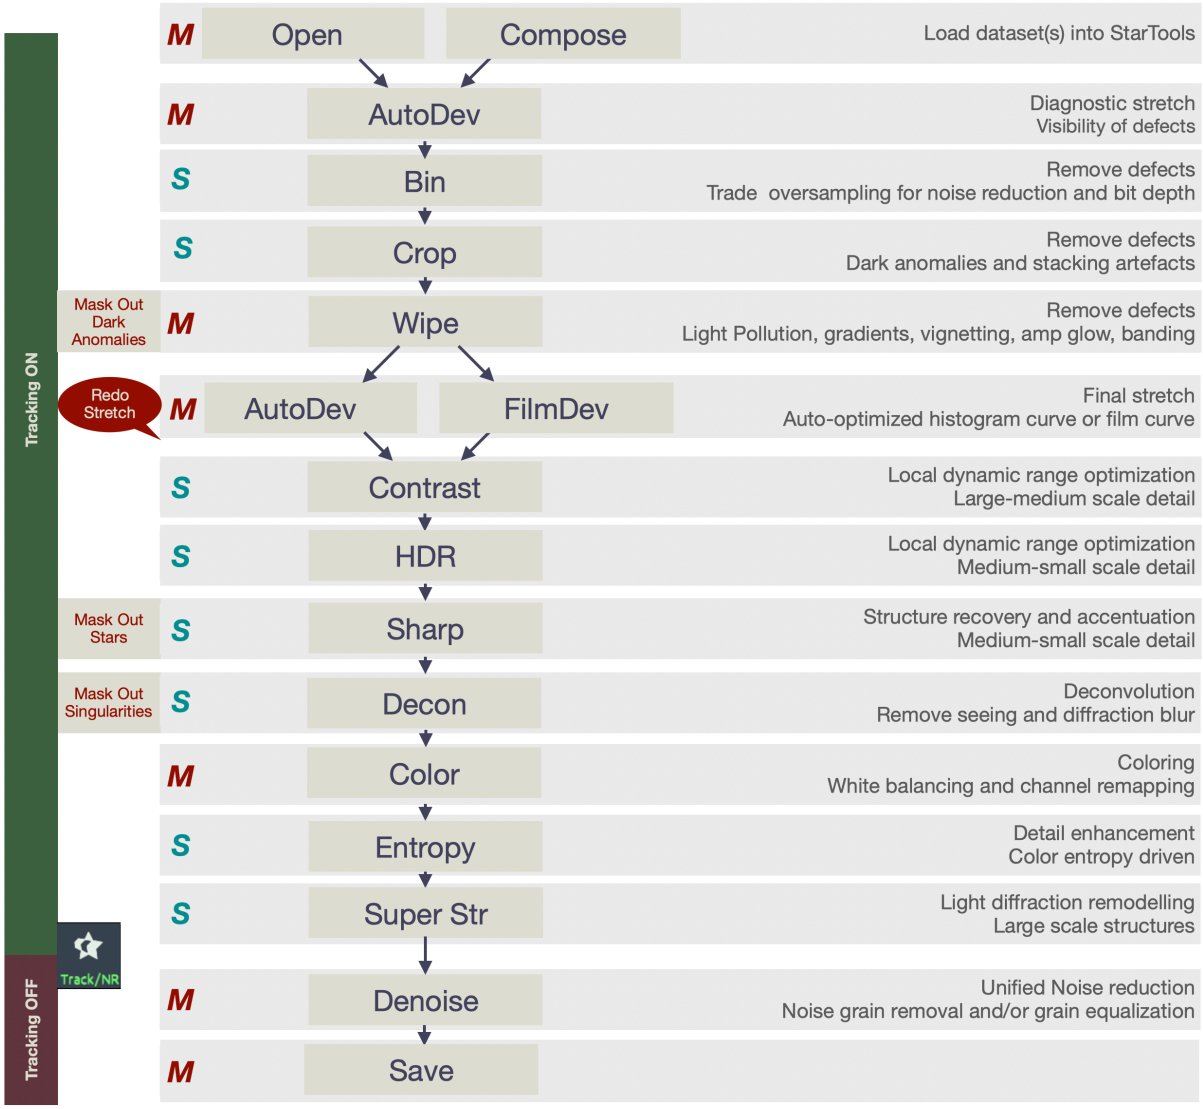

Basic workflow with mandatory ***(M)*** and suggested ***(S)*** modules. The latter may be used as desired / needed.

- StarTools workflow allows some flexibility. The above workflow will yield best results for most scenarios, doing many things right, but is not the only option
- Loading, stretching and noise reduction steps offer two options each
- Mandatory masks are noted. 'Sharp' will require an inverse star mask. 'Decon' module requires masking out singularities only
- The below step-by-step tutorial will run you through the main modules, including instructions and background information you need to know
- It is designed to provide a steep learning curve for operation, modules and most important parameters.
- More detailed Instructions, Explanations and Background information may be found in 'User Notes' and 'Modules'.

1) Loading Your Data into StarTools: Open or Compose

Choosing the Module

- 'Open' is used to open single monochrome or color images (e.g. DSLR, OSC)
- Used to merge multiple monochrome files such as those representing Luminance and Red, Green, and Blue channels, or those representing the Hubble palette.

Open

- After loading you are prompted about the image source. This will affect how you open the image in StarTools:
- A stacked RAW image of OSC or DSLR - select 'Linear'
- A stacked RAW color image that is not white balanced - select 'Linear, from OSC/DLSR with Bayer Matrix and not white balanced'. The data will be loaded using the Compose mode.
- A JPEG or video source use 'Non-Linear sRGB source'.

Compose

1. If you have any Luminance data, load that first. Else you must specify a synthetic luminance option in 'Luminance, Color' setting or 'Compose' will not be enabled.
2. Load any Red, Green and Blue data you have.
3. If there is a difference in exposure time between red, green and blue data set the ratio settings accordingly.

2) Diagnostic Stretch: 'AutoDev'

- The preliminary stretch will make object and all imaging issues visible so these may be tackled.

'AutoDev'

1. Load 'AutoDev' module. By default, it will make the whole dynamic range visible.
2. A 'bad' result will usually show up, because all imaging defects will be highlighted:

- Stacking Artifacts - remove later with the 'Crop' module.
- Color bias (remove this later with the 'Wipe' module)
    - Red or yellow/brown cast - skyglow that has been white balanced.
    - Teal, blue or green cast - skyglow that has not been white balanced.
    - Bright blue-green cast - skyglow filtered using a light pollution filter.
    - Missing yellow (e.g. no yellow stars) - indicates use of light pollution filter.
- Vignetting - darkening towards the corners, Amp Glow - remove later with the 'Wipe' module.
- Dust specks - remember to mask out when using the 'Wipe' module.
- Noise - will be addressed by 'Denoise' module.
- Banding - will be tackled using 'Wipe' module.
- Debayering Problems - checkerboard pattern. See the description here.
- Coma - fix later with the 'Lens' module.

3. Identify the sort of gradients present and note the 'Wipe' preset to be used later on:
    - If the brightness drops away at the corners choose 'Vignetting'.
    - If there is a general lightness in the background chose 'Basic'.
    - If there is an intensity in a particular area (e.g. edge) - you have 'Amp glow'.
    - If No Darks and/or Flats were used, try 'Uncalib1' or 'Uncalib2'.
    - For Narrowband datasets, use 'NrwBand'
4. 'Keep' the result.

3) Remove defects: 'Lens', 'Bin', 'Crop', 'Wipe'

- Any defects in the data will need to be addressed prior actual post processing.

'Lens'

- These modules are only required in special scenarios. See Full Workflow.

'Bin'

1. Load the module - this automatically 'Bin's the image by 50%.
2. Zoom in on the image to see if the image is still oversampled. Small stars would occupy more than 3x3 Pixel when oversampled.
3. Try different levels of binning until the smallest features occupy only a couple of pixels on their minor axis. Leave a bit of oversampling for the 'Decon' module later.
4. 'Keep' when finished.

'Crop'

1. Remove the stacking artifacts or crop the image to taste.
2. You may adjust the area to be kept using 'click and drag' or the sliders.
3. 'Keep' when done.

'Wipe'

1. After removing any stacking artifacts using 'Crop' module, load the 'Wipe' module.
2. Set the mask to exclude pixels you are sure are not background - e.g. the target, dust specks and other dark anomalies as described below:
    - Mask - Clear - Lasso around Dark Anomalies and Target as needed - Invert - Keep.
    - Or Mask - Clear - Flood Fill Lighter Pixels - click on an edge pixel on the subject to highlight it - Grow x 4 - Invert - Keep.
3. Choose a preset which best describes what you are trying to achieve (see above). If there are multiple issues choose the most aggressive.
4. Default values often work fine. Wait for the processing to complete.
5. Check the results - Use the 'Color'/Luminance' button to check both results if in 'Compose' mode.
6. Adjust the 'Dark Anomaly Filter' to mask the effect of small dark anomalies (see below). Increase until no further improvement.
7. Increase 'Correlation Filtering' to mitigate correlated noise and artefacts caused by dithering, debayering or fixed pattern sensor cross-talk issues. Values between 1/2x and 1x of the smallest detail (seeing limit or pixel resolution) will work fine.
8. If you see banding or other sensor issues, adjust the 'Synth' Dark/Bias using the appropriate setting for the issues seen. Wait for the processing to complete.
9. Reduce the 'Synth/Bias Edge Area' from 100% if required.
10. Adjust the 'Gradient Falloff' only when there is vignetting.
11. Use 'Before' button to compare with original.
12. Press 'Keep' when done.

4) Final Stretch: 'AutoDev' or 'FilmDev'

Choosing the Module

- After removing the gradients using 'Wipe', the stretch should be redone to use the freed up dynamic range.
- 'AutoDev' will optimize global stretch for the visual content to show the most detail in each dynamic range. Usually it is the best choice for most images.
- 'FilmDev' mimics a film tone curve to provide a classical yet non-optimized look.
- 'FilmDev' will be used if You may not get satisfying results using 'AutoDev' because:
    - There is no region of interest with sufficient detail for 'AutoDev' to use.

    - There is a lot of noise which disrupts 'AutoDev' operation.
- 'FilmDev' is also used for adding back an artificial skyglow at the end of processing.

'AutoDev'

1. Select Region of Interest (ROI):
    - Highlight the subject, or a part of it, to select the range of levels that 'AutoDev' should allocate dynamic range to.
    - Sometimes highlighting an area within the subject gives the optimum dynamic range. Try various selections.
3. Set 'Ignore Fine Detail' - to ensure 'AutoDev' ignores noise, dust specks and other dark anomalies. Increase the value until background and noise will not darken any further.
4. If you want to highlight detail in the shadows then increase 'Shadow Linearity' above 50%. Reducing this will darken shadows, pushing back noise.
5. If you make a mistake, the 'Reset' button discards all the changes since you started using the module.
6. 'Keep' the result when you are happy with it.

'FilmDev'

1. Increase Digital Development to set the black point of the global stretch - The 'Home In' button usually finds a good value.
2. Increase Dark Anomaly filter so the noise and other dark anomalies are ignored - usually when the background stops darkening.
3. Usually all other parameters work best with their default settings - but try changing them as described below if you want.
4. If you make a mistake, the 'Reset' button discards all the changes since you started using the module.
5. 'Keep' the result when you are happy with it.

5) Local Dynamic Optimization: 'Contrast', 'HDR'

'Contrast'

- To optimize medium-to-large local contrast (dynamic range) by doing local stretching.

1. Use one of the 3 presets depending on desired performance:
    - 'Basic' (default) will darken the image locally reducing glare
    - 'Local' will aggressively optimize dynamic level locally
    - 'Equalize' will darken highlights and lighten shadows creating a tranquil look
2. 'Contrast' will deallocate dynamic range used by small dark areas in favor of brighter features. This is controlled by 'Shadow Detail Size' and 'Shadow Dynamic Range' parameters.
3. For larger anomalies use the 'Heal' or 'Crop' modules first. There is no equivalent to the mask in the 'Wipe' module. Using the 'Heal' module must be done with Tracking off which may cause complications and so limit this as an option.
4. 'Keep' the result when you are happy with it.

'HDR'

- To optimize an images medium-to-small local contrast.

1. Select the appropriate preset:
    - Reveal - reveals details in almost-overexposed cores. Default Setting.
    - Tame - brings out detail in bright areas such as galaxy cores, like Equalize but doesn't brighten faint detail.
    - Optimise - accentuates detail it can find in both shadows and highlights.
    - Equalise - brings out detail in both shadows and highlights of larger areas.

2. Adjust the Context Size to target the size of detail you want. Changing Context Size starts a lot of processing and can be slow so get this set first. Selecting an area to sample will speed up processing. Other settings are quicker. If the default setting takes too long - try reducing it. Even reducing it to 30 and see if you get acceptable results.
3. Adjust the 'Highlights Detail Boost' and 'Shadows Detail Boost' to accentuate detail.
4. Adjust the 'Gamma Highlights' and 'Gamma Shadows' to accentuate detail in larger areas.
5. Adjust the 'Gamma Smoothen' parameter to smooth the transition between locally stretched areas.
6. Adjust the 'Signal Flow' if required
7. Toggle top "Before/After" button to see effect of module if needed.
8. If a sample area was selected press 'All' to process the complete image.
9. Press 'Keep' when done.

6) Structure Recovery and Accentuation: 'Sharp'

- To enhance small to medium scale structures.

1. Select the Structure Size you want to sharpen - and press 'Next'.
2. To create inverse star mask, chose between 'auto generate mask' or 'auto generate conservative mask'.
3. Select a suitable preset according for your object:
    - DSO (deep sky object) - default
    - DSO Dark
    - DSO Light
    - Planetary
4. Adjust the 'Amount' control to bring out the structures as much as possible without increasing the noise at all scales.
5. Within the selected Structure Size are 5 different layers of different scale sizes. If necessary, reduce the individual Scale setting to reduce the effect at that scale.
6. Toggle top 'Pre Tweak/Post Tweak' button to see the effect of last adjustment if needed.
7. Use Before/After button to compare between original and processed version.
8. 'Keep' the result when you are happy with it.

7) Deconvolution: 'SVDecon'

- 'SV Decon' removes seeing and/or optical diffraction customized to multiple image sections .

1. Select preset as per image type as a baseline - Deep Space or Lunar/Planetary.

'Deep Space'

Method

1. For DSOs:
2. This is a way of using the module which should give good results in most cases:
3. Select 'Auto-generate mask' to create an apodization mask - a form of star mask - or use the ApodMask preset to create one - this masks out overexposed stars which cause ringing.
4. Select a preview area rectangle in the image to see the effect of changes - this speeds up analysis. Choose a bright, detailed and noise-free area.
5. Click on 'Sampling' button to show the possible sample stars.

6. Keeping Zoom at 100% - should allow the stars should be large enough to evaluate and for you to position the sliders at the edge of the image easily to cover the whole image in 6-9 segments.
7. Select samples - at least two from each segment
8. Samples should be green without red centres. See 'Choosing sample stars' below.
9. Add more samples if there is a lot of variation in the way the stars are deformed.
10. Put a sample in any area where the stars have been deformed in a slightly different way.
11. Adjust 'Sampled PSF Area' so that the blue area just surrounds the selected stars green pixels.
12. Increase 'Sampled Iterations' until you see no further improvement
13. Set 'Synthetic Iterations' to 'Off' unless you want to use an optical PSF model
14. Zoom in and out so you can see the effect in the detail and as a whole.
15. Toggle top "Pre Tweak/Post Tweak" button to see effect of last adjustment if needed.
16. Normally other default values work well - but you can experiment if you want.
17. When done, select 'All' to apply this to the whole image - this may take some time.
18. Press 'Keep'.

'Lunar/Planetary'

1. For lunar/solar/planetary targets:
2. Click the 'Plnt/Lnr' button
3. Don't try and select a sample star
4. Select a Synthetic PSF model to model the atmosphere.
5. Increase 'Dyn. Range Extension' to increase the dynamic range.
6. Increase the 'Synthetic Iterations' until no improvement is seen.
7. Adjust the 'Synthetic PSF Radius' until ringing occurs - then back off a little.
8. Press 'Keep'.

Selecting Sample Stars in Apod mask

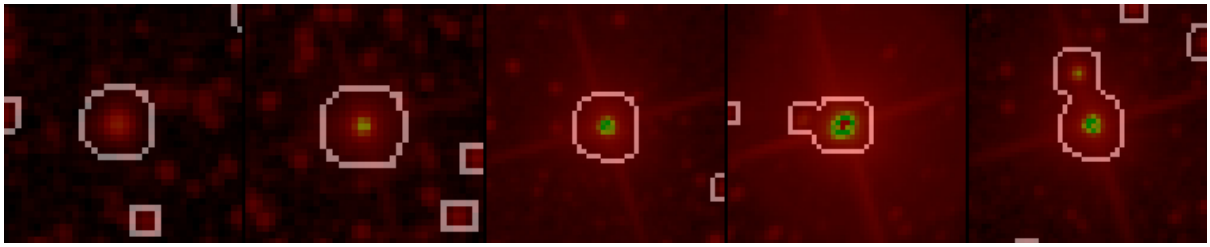

1. Panel 1 presents a basic red ball. These are not to be used for SVD, but if all you have is those, try picking the brightest ones.
    - A pure synthetic PSF might just be better, if good quality samples are not yielding expected results.
    - The Moffat models should quite effectively undo the effects of atmospheric turbulence and may provide better results than chasing samples in crowded wide fields.
    - Other imperfections/distortions in the PSF (whether local or across the entire image) may not even be visible at some scales.
    - Usually, atmospheric influences are the biggest contributors to out-of-the-ordinary PSFs, if the rest of your gear is dialed in reasonably well.

2. Panel 2 has some yellow in the core, no green.
    - Those should be OK as a secondary choice if You have nothing else in the area.
3. Panel 3 has some green.
    - This appears to be a good sample.
    - You may even want to increase the apodization mask so that more of the stellar profile is included.
4. Panel 4, (if one ignores the adjoining star), is green but red pixeled in the center.
    - Ideally such is not to be used
    - f you are quite sure that the star isn't over-exposing and your sensor's (and stack's!) response stays linear with stars this bright, then it may be used as a secondary choice .
5. Panel 5, (ignore the double star) is starting to show a green center but there's a yellow pixel (though not gone red) in the very center.

- CCD well saturation (e.g. photons don't get converted into electrons as much due to the wells becoming "full").
- Stacking data of variable quality. E. g.,
    - a star under poor seeing will not saturate as quick, because its light is "smeared out" over multiple pixels.
    - Under good seeing conditions, however, stars will over-expose quicker as the light gets concentrated into fewer photo sites.
    - If you Stacking a mixture of those, you will get stars with cores that have unreliable data (due to averaging over-exposing cores with non-over-exposing cores).

6. The apodization mask indicates which pixels may be included in a sample and which may not.

- The blue box indicates the location and area that will be sampled.
    - Any pixels inside the blue box AND inside the apodization mask get included in the final sample.
    - Anything outside the box is excluded from the sample.
    - Anything inside the box, but not set in the apodization mask is also excluded from the sample.
    - You can certainly touch up the mask to include any part of the stellar profile you wish (or do not wish) to include.

8) Coloring: 'Color'

1. At startup, if there is a full mask set, the module auto-calibrates the white point in the image.
2. Select Style 'Scientific, Detail Aware' for accurate color reproduction, or 'Artistic, Not Detail Aware' for a more classical look.
3. Leave default 'LRGB Method Emulation' option 'Straight CIELab Luminance Retention'.
4. If there are problems:
    - Use MaxRGB to look for issues (like unexpected green-dominant areas) - adjust the Red, Green and Blue bias to get a better balance.
    - Use Cap Green control to eliminate any unwanted remaining green.
    - If there are color problems around stars or other highlights - use Highlight 'Repair' (v1.6)
5. If you make a mistake, the 'Reset' button discards all the changes since you started using the module.
6. Select the preferred Style: 'Scientific, Detail Aware' for accurate color reproduction, or 'Artistic, Not Detail Aware' for a more classical look.
7. 'Keep' the result when you are finished.

9) Detail Enhancement: 'Entropy'

- To enhance local detail based on color information

1. Leave Default or choose preset of color channel(s) to be targeted
2. Press 'Do' to generate the entropy map and get initial results.
3. Try different settings of Strength, Dark/Light Enhance.
4. Press 'Do' to update to see the results of the changed settings.
5. Press 'Keep' when you have the results you prefer.

10) Light Diffraction Remodeling: 'SuperStructure'

- 'SuperStructure' will process large scale structures in the image (e.g. galaxies & large scale nebulosity) independent of the rest of the image to allow effects such as de-emphasizing busy star fields, emphasizing large structures or enhancing their color. It will bring back 'life' into an image by remodeling uniform light diffraction.

Isolate or DimSmall presets - to push back noise/star field

1. Select 'Isolate' preset - this sets most of the controls to their optimum.
2. Set 'Strength' between 30-100% - to adjust the impact of the 'SuperStructure' module.
3. Set 'Saturation' - Change this to adjust the amount of ‚glow‘

Brighten or Saturate - to add 'life' - increased glow and a more 3-dimensional effect.

1. Select preset.
    - 'Brighten' brightens detected superstructures
    - 'Saturate' saturates detected superstructures
2. Set 'Strength' between 50-100% - to adjust the impact of the 'SuperStructure' module.
3. Set Saturation - Change this to adjust the amount of 'glow'.
4. When you are finished 'Keep' the result.

11) Unified Noise Reduction: Noise Grain Removal and Equalization

Unified 'Denoise'

1. Click the 'Track/NR' module button to select the 'Denoise' Module.
2. Set Walking Noise Size and Angle parameters as necessary so as to eliminate any Walking noise in the image. Then set the Grain Size.
3. Select Grain Size so the noise grain and clumps can no longer be seen - do not take note of the actual image content in this view. Structures larger than the Grain Size are considered detail, not noise.
4. Click 'Next' - StarTools will do its initial attempt using that grain size with other settings at their default values. When complete screen 2 is shown.
5. Select a sample area to speed up the processing while you adjust these controls:
    - Adjust Brightness Detail Loss and 'Color' Detail Loss - to balance detail loss and noise reduction.
    - To control the balance between detail retention and noise reduction within the subject adjust the Scales (e.g. consider using Scale 5) and Scale Correlation. Scale Correlation identifies how much of the smaller structure is considered detail.
6. If you prefer a more film-like result where there is a uniform yet invisible grain in the background, increase the 'Grain Equalization' parameter.
7. Toggle top "Pre Tweak/Post Tweak" button to see effect of last adjustment if needed.
8. Press 'Full' to apply the effect to the full Image.

9. If you make a mistake, the 'Reset' button discards all the changes since you started using the module.
10. Press Keep to exit, keeping the results.

11) Save Your Result: Save

1. Save the current view using the 'Save' Button
2. Choose one of the supported formats:

- To save as a 16-bit TIFF file use the extension .tiff or .tif
- To save as a JPEG file use the extension .jpg
- To save as a 8-bit PNG file use the extension .png (StarTools v1.5+ only)

Full Workflow

Advanced workflow with mandatory ***(M)*** and suggested ***(S)*** modules to the left and special purpose modules to the right.

- StarTools workflow allows some flexibility. The diagram provides a suggested flow.
    - Modules between 'Color' and 'Denoise' might well be used in a different sequence, depending on purpose.
    - Final touch-up modules (tracking off block) might be used in a different sequence, depending on purpose.
    - 'Layer' and 'Flux' might be used at many position depending on purpose. Hence these are not shown in the diagram
- Check out 'Special Techniques' to explore the versatile use of the modules
- Loading, stretching and noise reduction steps offer two options each
- Mandatory masks are shown. 'Sharp' and 'Flux' will require an inverse star mask. 'Decon' module requires masking out singularities only. 'Shrink', 'Repair' and 'Synth' work on a mask including stars only

Further Modules

StarTools provides further modules for special purpose. Please see Full Workflow Diagram for where to fit in those. These modules will enable You to:

- Mitigate aberrations
- Decrease star size and improve star colors
- Synthesize stars from original pixels
- Remove unwanted objects or defects
- Layer tools and math
- Add narrowband accents
- Create depth information for 3D presentation

'Lens'

To correct coma:

1. Increase the 'Curvature Linked' setting until the stars in the corners are circular.
2. Adjust Center X and Center Y to correct if the centre of the lens is slightly off-centre.

To correct chromatic aberrations:

1. Adjust the Red Shift X and Red Shift Y until the red fringe of the star has moved into the stars centre.
2. Adjust the Blue Shift X and Blue Shift Y until the blue fringe of the star has moved into the stars centre.

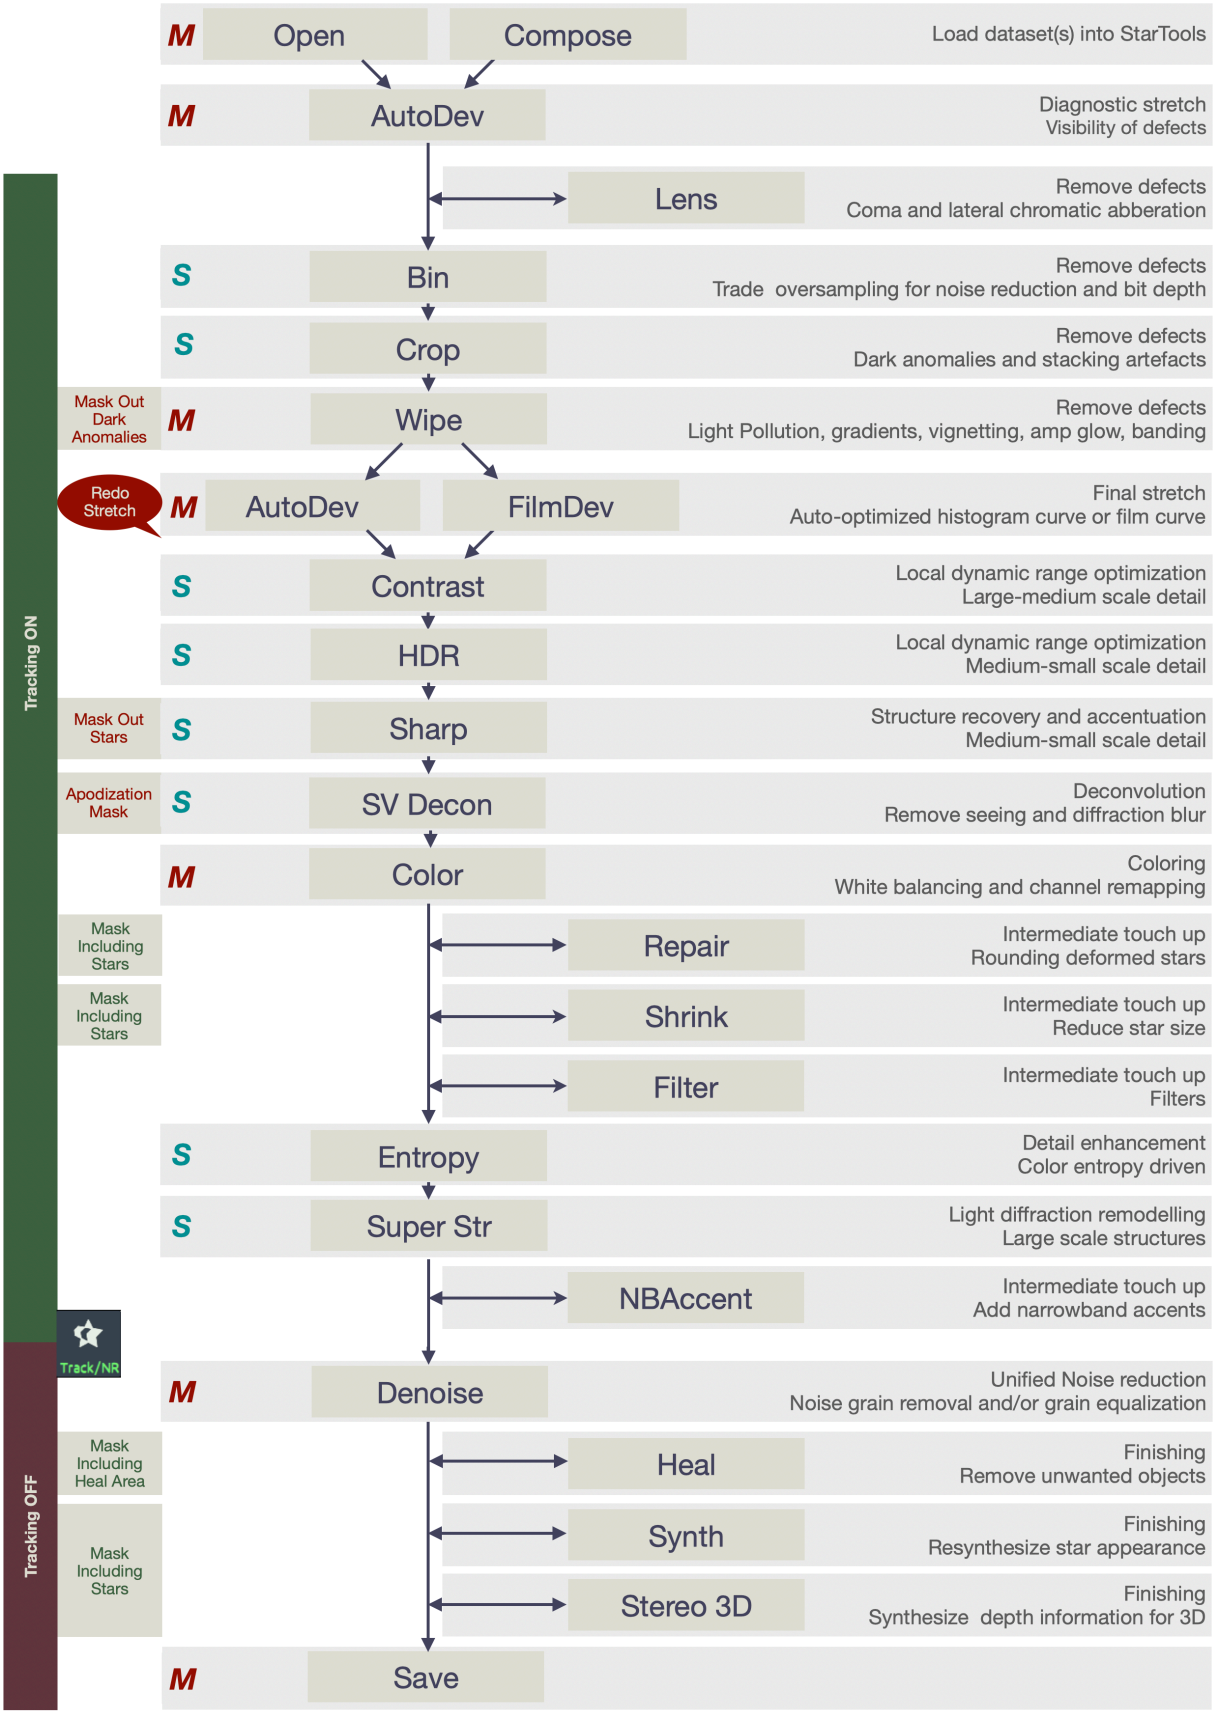
M
Open
Compose
Load dataset(s) into StarTools
M
AutoDev
Diagnostic stretch
Visibility of defects
Lens
Remove defects
Coma and lateral chromatic abberation
S
Bin
Remove defects
Trade oversampling for noise reduction and bit depth
S
Crop
Remove defects
Dark anomalies and stacking artefacts
Mask Out Dark Anomalies
M
Wipe
Remove defects
Light Pollution, gradients, vignetting, amp glow, banding
Redo Stretch
M
AutoDev
FilmDev
Final stretch
Auto-optimized histogram curve or film curve
S
Contrast
Local dynamic range optimization
Large-medium scale detail
S
HDR
Local dynamic range optimization
Medium-small scale detail
Mask Out Stars
S
Sharp
Structure recovery and accentuation
Medium-small scale detail
Apodization Mask
S
SV Decon
Deconvolution
Remove seeing and diffraction blur
M
Color
Coloring
White balancing and channel remapping
Mask Including Stars
Repair
Intermediate touch up
Rounding deformed stars
Mask Including Stars
Shrink
Intermediate touch up
Reduce star size
Filter
Intermediate touch up
Filters
S
Entropy
Detail enhancement
Color entropy driven
S
Super Str
Light diffraction remodelling
Large scale structures
NBAccent
Intermediate touch up
Add narrowband accents
Track/NR
M
Denoise
Unified Noise reduction
Noise grain removal and/or grain equalization
Mask Including Heal Area
Heal
Finishing
Remove unwanted objects
Mask Including Stars
Synth
Finishing
Resynthesize star appearance
Stereo 3D
Finishing
Synthesize depth information for 3D
M
Save
Tracking ON
Tracking OFF

'Shrink'

1. Launch 'Shrink' module
2. Create Star Mask. Mask-Auto-Stars-Do-'Shrink'-Grow-Keep
3. Pick preset.
   - 'Tighten' will tighten the stars around their core.
   - 'Dim' will reduce the luminosity of stars.
   - 'Un-glow' will remove the glow around stars.
4. Set 'Iterations' to whatever you like (more iterations draws color in further).
5. You might also want to increase the 'grow mask' parameter 1 or 2 pixels.
6. 'Keep' the result.

'Filter'

1. Specify the action you want to perform when you choose the color range (spectral line) by setting the 'Filter' Mode.
2. Optionally - create a mask to select only elements you want to act on. If you don't the selected action will apply to the selected color range anywhere in the image.
3. Select a particular color on which to do the specified action by clicking on that color in the image.
4. Repeat the selection to enhance the chosen effect.
5. If you make a mistake, the 'Undo' button undoes the last operation, the 'Reset' button discards all the changes since you started using the module.
6. 'Keep' when the result is right.

'NB Accent'

1. The narrowband data must have been loaded previously in the Compose module:
   - Add your luminance and RGB data as normal.
   - Press 'NBAccent' button and navigate to and load the narrowband data file.
   - Select the 'NBAccent Type' which best describes the data just added.

Page 1: Signal Stretch and accent scope definition.

2. The RoI will be inherited from the AutoDev module - change this if needed.
3. Adjust 'Ignore Fine Detail <' to remove any noise and other small detail you don't want to be accentuated.
4. Other parameters are normally left at their defaults but may be adjusted as described below.
5. Press 'Next'.

   Page 2: Image Accentuation colouring and brightness.
6. Depending on the subject - press the associated preset - 'nebula' or 'galaxy'.
7. Set 'Response Simulation' to the colouring appropriate to the narrowband channel(s) in the data.
8. Set the 'Strength' parameter to give the right emphasis to the accents.
9. Set the 'Luminance Modify' to define how much you want the luminance of the image to be modified by the accent.
10. Set the 'Color Modify' to define how much you want the colour of the image to be modified by the accent.
11. Set the 'Detail Size' to limit the size of the elements that appear as accents.
12. Other parameters are normally left at their defaults but may be adjusted as described below.
13. Press 'Keep' when done.

'Heal'

1. Open the 'Heal' module.
2. Create a mask highlighting the elements to be removed.
3. Set 'New Must be Darker Than' as needed
4. Set 'Grow Mask' as needed to replace pixels around the elements to be removed.
5. Set 'New Darker Than Old' to Yes.
6. 'Keep' the result when finished.

'Repair'

1. Create a star mask Mask-Auto-Stars-Keep making sure that the stars in the mask are separated by at least one 'off' pixel.
2. Load the 'Repair' module.
3. Select the Algorithm - Warp for oval stars, One of the Redistribute modes for more heavily distorted stars.
4. You can usually leave all other settings at the default values.
5. Experiment with the settings if the results aren't right.
6. 'Keep' the result.

'Synth'

1. Select the 'Newtonian' or 'Refractor' Preset according to the telescope used to take the image.
2. Set the Aperture value to that of the telescope used to take the image.
3. Set the Focal Length to that of the telescope used to take the image.
4. The default values may be sufficient for the other settings - but adjust any that you know are wrong.
5. Press 'Next'.
6. Define a star mask - Mask-Auto-Stars-Do-Keep
7. Set the Blur parameter so the diffraction pattern of the stars are consistent with the blur in the image caused by the atmosphere.
8. Set the Image Diameter value to the angular size of the image.
9. Set the Overlay Mode to 'Hybrid, Keep Stars'
10. The default values may be sufficient for the other settings - experiment if needed.
11. 'Keep' the desired result.

'Flux'

1. Select the Wavelet Library size - after a while you will be presented with Screen 2.
2. Select the desired preset - normally choose 'Sharpen' or 'Noise' - avoid using 'Detail' preset except in special circumstances.
3. Use Before/After button to compare between original and processed version.
4. Keep the result when done.

'Layer'

There are so many ways of using the 'Layer' module with so many different aims and no single method stands out - See the Special Techniques section for some of the more common uses.

The general process is:

1. If multiple images are to be combined make sure they are the same size and aligned the same as described in the Background Notes.
2. Start the 'Layer' module - the current image gets loaded into both foreground (centre) and background (left) fields.
3. Load another image - this will replace the image in the foreground.
4. Swap the images if necessary to put the right image in the foreground.
5. Create a mask if needed to protect certain parts from being altered.

6. Select the desired 'Layer' Mode.
7. Select 'Filter' Type or Brightness Mask Mode as required.
8. Adjust other settings as required.
9. 'Keep' the desired result.

'Stereo 3D'

1. Choose the mode that is most comfortable for you to perceive depth
2. For Side-By-Side 'Cross' or 'Parallel' modes - zoom out so you see both L and R images
3. You can process using the easiest mode to view (e.g. Anaglyph with glasses) and then convert later
4. Choose the way(s) of getting depth information best for this subject:
    - Dark detail in front of light background - e.g. dust clouds - Luma to Volume - Shadow Dominant
    - Light detail in front of dark background - e.g. M20 - Luma to Volume - Highlight Dominant
    - If there are still bright elements you want to bring forward - Simple L to Depth - Increase from 0%
5. Control the embedding of elements:
    - Large bright areas - e.g. emission or reflection nebulosity - Highlight Embedding - use higher values to embed
    - Bright objects within reflection nebulosity - e.g. bright stars within reflection nebula - Highlight Embedding - use higher values to embed
    - Small structures in front of large - Structure Embedding - use lower values.
6. Use 'Before'/'After' and 'PreTweak'/'PostTweak' button to see the effect of changes
7. Set the Mode to 'Depth Map' to see the areas that have been brought forward (white)
8. Choose the final Mode required
9. Export to: Web 2.5, Virtual Reality or Facebook 3D Photo as needed
10. Press 'Keep' when done

Hands On example IC1396 - Elephant's Trunk

The challenge with this image is to achieve a balanced stretching over a huge dynamic range throughout the whole image field: All detail up from dark nebula regions, faint Ha-regions up to bright stars should be allocated enough dynamic range so detail contrast would not suffer. Also star bloating should be avoided. In ST, this may be achieved without tricks like selective 'masked stretching'.

Just setting a small ROI in AutoDev, around the elephant's trunk, results in stretching to appear very focused on dynamic range around IC1396, causing the extended nebulosity and dust to be largely invisible. This is, of course, a choice that is wholly up to you. However, confining the dynamic range to just that area will neglect allocating dynamic range to the rest of the image. As a result, stars (which tend to span the entire dynamic range), will stand out against a mostly invisible background, rather than being "embedded" in the rest of the detail that is present in the complex. Again, this is a choice that's as valid as any, but from your post, it sounds like this was not a choice.

In the final AutoDev, if you can you might consider including a bright star within any ROI, so that the global stretch has that star to work with in setting the overall dynamic range. That may help keep bright stars under control if they are getting away from you (classic example of this is 52 Cyg in the Veil).

E.g., a simple image-wide AutoDev on IC1396, yields this:

Star location and prevalence tends to correlate with larger scale structures; where there are emissions there are stars. Where there is dust, stars are obscured. This is incidentally, the premise behind the Super Structure module as well and why the things like the Isolate or DimSmall presets in that module work so well.

Keeping the two balanced (super structure and star prevalence) will result in a much more pleasant (and informative!) image to look at.

Use local dynamic range manipulation (e.g. Contrast and possibly HDR if the image requires it) if you feel too much of an object is lost in the highlights (or even shadows). E.g. a default application of the Contrast module brings out IC1396 by locally re-allocating the dynamic range:

Keeping the two balanced (super structure and star prevalence) will result in a much more pleasant (and informative!) image to look at.

Use local dynamic range manipulation (e.g. Contrast and possibly HDR if the image requires it) if you feel too much of an object is lost in the highlights (or even shadows).

E.g. a default application of the Contrast module brings out IC1396 by locally re-allocating the dynamic range:

Then, after SVDecon and Color module, you should get something like this:

Notice the now visible Ha emission to the left, and notice how the distribution of stars feel like they "make sense", given the larger structures and different emissions that are now visible.

For your particular case, the Shrink module has a special "Unglow" feature that actively suppresses halos around "smaller" stars.

As mentioned above, you can further exploit the relationship/correlation between star location/brightness/temperature and physical processes/dust/emissions by using the Super Structure module. 'Dimsmall' or 'Isolate' presets can be useful for maintaining the super structures while pushing back stars and background. Be sure to set your Airy disk parameter to suit! To the left for wide field, higher for narrow field. A rule of thumb is the default 50% works well for 600mm focal length and APS-C (without major cropping), adjust from there.

If you left your stretch weaker, you may be able to come back in with a second SS (brighten and/or saturate) to help enhance the target against the stars and background.

Now, stars don't feel "out of place", but appear like a natural part of their surroundings; because they are!

No special masking techniques required, no star removal, relayering or any other uncouth/ selective processing shenanigans.

Finally, if you really wanted to, you could use the Filter module's "Reject" Filter Mode to reduce the brightness of the blue halos:

Season everything to taste obviously.

Note: I should also say that it is important to compare your images to other images taken in the same spectrum; IC1396 is a very popular narrowband object, and as such star prevalence will be much different/less in such narrowband composites.

Special Techniques

Mask Module

To revert to a module processing all pixels

Return to the default mask for any module - all pixels selected.

1. Clear
2. Invert
3. Keep

Create a Star Mask

Create a star mask based on the Auto 'Star' button preset values.

1. Auto
2. Star
3. Do

Star Mask when there is a lot of noise in the image

The above procedure can show false 'stars' When the image is noisy. This procedure may eliminate this problem

1. Auto
2. Star
3. Set Selection Mode to 'Highlights>Threshold'
4. Set Threshold to say 60%
5. Press 'Do'
6. If there are too few stars detected, reduce the threshold and press 'Do'. If noise is being selected increase the threshold. Repeat until you get the result you want.
7. Keep the result.

Creating an Inverse Star mask for 'Decon' use - AutoMask equivalent

This creates an inverse star mask with stars and their halos masked out. This gets the same results as the AutoMask button in 'Decon'. This is only possible to do when Tracking is on.

1. Auto
2. FatStars
3. Grow x4
4. Invert

Selecting background only

To select the background only:

1. Press 'Auto'.
2. Select 'Stars' preset.
3. Reduce Threshold until only the background shows - usually to about 10-20% .
4. Press 'Do'.
5. Repeat until only the background is visible.
6. Press 'Invert' to select the background.
7. Keep

Selecting Faint Stars Only

There are a number of cases when you only want to select the fainter stars.

1. First select all stars including small ones:
    - Use Auto->Stars preset
    - Set Feature Size to 1 - to select the smaller features only.
    - Increase the 'Filter' Sensitivity value to show fewer stars if needed.
    - Press 'Do'

2. Then select the brightest stars and subtract them:
    - Auto
    - Set Selection Mode to 'Highlights>Threshold'
    - Set Threshold to say 75%
    - Set Old Mask to 'Subtract New From Old'
    - Press 'Do'
3. Keep the result.

Making Luminance Mask of a DSO

1. Set Brush Mode to 'Flood Fill Lighter Pixels'
2. Click on edge of DSO - any neighboring pixels which are the same as or lighter will be selected. It may take a few tries to get the right coverage.
3. Click on the edge of any other areas of the DSO that have not been selected.

Making Luminance Mask of DSO excluding stars

Sets the mask of a DSO based on a brightness threshold - but excluding stars.

1. First make a mask based on brightness:
2. Auto
3. Set Selection Mode to 'Highlights>Threshold'
4. Set Source to 'Stretched'
5. Set Threshold to, say 30.00% - to select as much of the DSO as you want.
6. All other settings should be OK at their default values. Adjust if you want.
7. Press 'Do'
8. Next select all stars including small ones:
9. Auto
10. Use 'Stars' preset
11. Set Old Mask to 'Subtract New from Old'
12. Press 'Do'
13. 'Keep' the result

Masking out long thin wisps

For masking long thin wisps (e.g. the veil nebula). This masking technique is shown in this video between 0m45s and 3m50s.

1. Create a trail of individual green pixels over the detail you want to mask (using the 'Single Pixel Toggle' Brush Mode)
2. Use the 'Grow' button until they merge.
3. 'Keep' the desired result.

Compose Module

Loading Narrow Band data using the Hubble Palette

1. Load the Luminance data - if you have it - into the Luminance channel.
2. Load the SII data into the Red channel.
3. Load the Ha data into the Green channel.
4. Load the OIII data into the Blue channel.
5. If you want to create a weighted synthetic luminance channel set the 'Luminance, 'Color" to 'L + Synthetic L from RGB, RGB'
6. Otherwise if you have Luminance data (and don't want to use LLRGB) set the 'Luminance, 'Color" to 'L, RGB'
7. Set the Total exposure times for each channel.
8. Press 'Keep' when done.

See also the description: LRGB & Hubble palette

Creating a Synthetic Luminance channel from Wide Band data

1. Load the Wide band image data into the Red, Green and Blue channels.
2. Set the 'Luminance, 'Color" to 'L + Synthetic L from RGB, RGB' [/url]
3. Set the Total exposure time for each of the Red, Green and Blue channels.
4. Press 'Keep' when done.

There is a useful video Painlessly re-creating the iconic Hubble 'Pillars of Creation' in StarTools.

Re-centering blue halos back onto stars

To recenter a smeared blue channel back to align with the red and green channels try the following as described in Fringe killer filter add blue back to central star

1. Optionally 'Bin' the image - then save it.
2. In the Compose module click Blue and navigate to the file just saved. This will load only the Blue channel of the saved image. The other channels will be automatically set due to 'Channel Interpolation'. You end up with a black & white representation of the blue channel
3. Process this image so that you reduce the bloating of the stars in the blue channel
    - Do 'AutoDev' or Develop - this allows us to see what we are doing.
    - You could then use the 'Decon' module - set to 'De-ring Mask Gaps, Show Result', set Iterations to 1, set Radius to 3.2 (as appropriate).
4. 'Keep' the result.
5. Use Restore selecting 'Linear, Wiped, Deconvolved' option - this reverts to an image which is Linear but which retains the deconvolution you did.
6. Save the image.
7. Use the Compose module to load the Red and Green channels (also 50% binned) and load the saved image into the Blue channel.
8. You should now have an image with a better focused blue channel. Process as normal.
9. When using the 'Color' module, the Scientific mode may have problems with the star cores as they have artifacts in the blue channel due to the deconvolution. Either:
    - Reduce Bright Saturation right down and set Saturation Amount to 100% or lower. This will avoid recovering color in the highlights, or
    - Use a star mask to mask out the stars - they then won't be involved in setting the color balance, or
    - Use Artistic mode

Simulating RGB image using narrowband data

1. Create a synthetic luminance frame from S+H+O (for your detail)
    - import S, H and O as R, G and B
    - Keep the result
    - Save as luminance file
2. Relaunch the Compose module.
3. Create a red channel from (S+H),
    - load S as red, green, or blue.
    - Load H as green.
    - Keep the result
    - Save as red channel
4. Relaunch the Compose module.
5. Load the SHO synthetic luminance (which you created earlier) as L,
6. Load the new red channel (S+H) you created earlier as red

7. Load O as green and blue.
- "Set Luminance, Color" to "L, RGB".
- Now process as normal.

Adding Ha to RGB datasets (method 1)

- A quick & easy method: Create a 'new red' dataset out of R & Ha

1. Create new red channel (only needed for OSC/DSLR datasets)
    - Load Compose Module
    - Channel Interpolation On
    - Load full-color data into red channel Only using the Red/S-II button
    - Save that monochrome dataset - this creates a mono dataset of the red channel
2. Create a blend of this with Ha
    - Load Compose Module
    - Load these two datasets - the red channel data using Red/S-II button, and the Ha data using Green/Ha button
    - Channel Interpolation On
    - Set 'Luminance, 'Color" to 'L + Synthetic L from RGB, Mono' to create a mono blend
    - Blend both datasets using 'Red Total Exposure' and 'Green Total Exposure' settings
    - Save this blend - this is the new red channel
3. Combine this new red channel with the other color data
    - Load Compose Module
    - Set 'Luminance, 'Color" to 'L + Synthetic L from RGB, RGB'
    - Load new red channel using the Red/S-II button
    - Load full-color data set using the Luminance, Green and Blue buttons. You now have R+Ha, G and B
4. Optional: Create synthetic luminance from RGB and 'new Red' datasets:
    - Set Compose mode to 'Luminance, 'Color' to L, RGB'
    - Load RGB dataset to Luminance, Green and Blue
    - Load new Red dataset to R
    - Save synthetic L
5. Process as normal

Adding Ha to RGB datasets (method 2)

- A more flexible way is to independently process RGB and Ha Datasets and combine them using the 'Layer' module

1. Process RGB dataset (or Compose load of individual R, G, B datasets) as normal.
2. Process Ha dataset as normal
3. Save processed Ha dataset
4. Combining both processed datasets
    - Launch 'Compose' module and set 'Luminance, 'Color" to 'RGB, RGB'
    - Set 'channel interpolation to 'Off'
    - Load process Ha into red channel. This should yield a completely pure red (processed) image
    - Open 'Layer' Module
    - Load RGB image in the background
    - Set 'Layer' Mode to "Lighten", "Add", "Screen" or even "'Color' of Foreground"
    - Try different Brightness Mask modes and Brightness Mask powers
    - Try different filters and kernel radii to make the Ha "glow"
    - Over-drive the effect with Blend Amount

'AutoDev' Module

Balancing faint details vs. noise

A way to lift as much faint details (e.g. nebulosity) above the noise floor without boosting noise too much (final stretch):

1. Set the ROI to include all details - do not pay attention to noise at this step
2. Increase 'Ignore Fine Details' Parameter. Noise will be pushed down. Increase further until no further improvement.
3. Reduce 'Shadow Linearity' until faint details start to get impacted. It is OK if some noise remains at this point. Usually 'Denoise' modules should be able to handle.
4. Further measures may help controlling noise:

- In 'Contrast' module 'Shadow Dynamic Range' may be reduced strongly. Set 'Shadow Detail Size' to 10-15 Pixel.
- In 'SuperStructure' module, 'Isolate' or 'DimSmall' presets will push down noise further.
- In 'Flux' module the 'Noise' preset might help as a last resort.

5. See also tipps in 'FilmDev' Module.

'FilmDev' Module

Final touch up

if there's still a problem existing after denoise and final touch-ups, and you don't want to go back and start over, FilmDev can act as sort of an "unstretcher.":

1. At the very end of your workflow, launch FilmDev module.
2. Increase 'Gamma' parameter slightly above 1.0 to push noise remnants towards black.
3. Use 'Skyglow' parameter to lift the black level to 1% or above to drown the noise peaks.

'Wipe' Module

Masking out flaws (e.g. dust or artifacts) that were hidden before running 'Wipe'
See the full description in the topic 'Wipe' module - editing mask.

- Use the 'Wipe' module as normal and 'Keep' the flawed result.
- Use the 'Mask' button on the main screen to edit the mask to exclude the newly highlighted flaws.
- 'Crop' any remaining stacking artifacts.
- Use the 'Restore' button on the main screen - select 'Original' - this reverts to the Original (Linear, 'Binned & 'Cropped) image.
- Use the 'Wipe' module as normal - which now uses the updated mask.

'HDR' Module

Masking out the effect of HDR on stars

This is a technique for selectively removing the effect of the last module. It can be used to remove the effect of the HDR Module on stars.
Launch the Layer Module
Open the Mask editor
Create Mask - Auto Stars - Do - Keep
In Layer Module - Copy the Undo buffer to the foreground layer - Undo > Fg
You can mask the effect of the HDR module as required - The mask active (green) area is the area where the effect of the HDR module is undone.

Adjust Mask Fuzz to get a smooth transition
Adjust the mask to include as many of the pixels surrounding the stars as you want - use Grow or Shrink as needed.
Keep the result when done.'Sharp' module can also be used as a scale decomposition tool - being able to remove or attenuate features of a particular size. This allows different versions of the image with different scales to be processed separately and blended together at a later stage.

Mimicking HDR results of V1.7

To mimic the old module (including its processing times) reduce the Context Size parameter to to 25 or something between 30 and 45, depending on processing time available. This is a little more small-detail oriented, but you can back off the appearance of that effect by adjusting the controls.

Mimicking the Reveal DSO Core preset:

reduce Context Size, and only set Shadow Detail Boost, while leaving all other parameters (e.g. Gamma Highlight, Gamma Shadow and Highlights Detail Boost) alone. One difference though compared to 1.7, is that even at max settings, the detail recovered in 1.8 is much more carefully placed in its surroundings, and does not look artificial, while also perfectly conforming to the noise levels of its surroundings (the latter extends also covers the interaction with the Gamma Highlight/Shadow parameter - consistency is maintained).

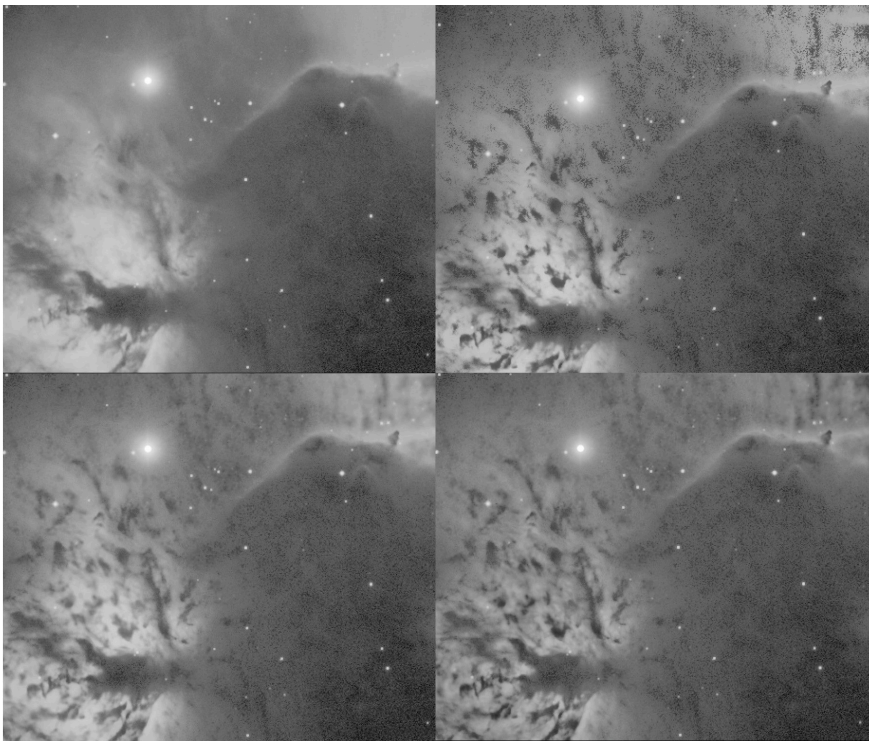

Top left original, top right 1.7 HDR, bottom left 1.8 HDR (Shadow Detail Boost 100%, nothing else), bottom right 1.8 HDR Shadow Detail Boost at 100%, with Gamma Shadow at 100%.
(obviously using the extremes of the parameters is not exactly advisable, but this is just for demonstration purposes)

'Sharp' Module

Scale aware processing

1. Make two (or more - depending on how many scales you want) copies of the image.
2. Blur one with, for example, a Gaussian blur.
3. Subtract the blurred image from the non-blurred image. The blur will have killed most of the fine detail (it's a blur after all), so subtracting the blurred image from the non-blurred image logically leaves you with just the fine detail.
4. You can now manipulate that fine detail - for example normalize it, use brightness/contrast operations on it, etc.
5. Add it (the modified fine detail) back to the blurred (coarse detail) image and voilà: you have an image that has had just a particular scale manipulated.

You could even do all this in something like The GIMP or Photoshop with the caveat that some additional trickery is required to manipulate and visualize pixels with a negative value. Now, nothing of course stops us from grabbing that fine detail we isolated, perform the same trick on that, so that you can start getting 'bands' of different detail sizes.
The radius of the kernel you use for the Gaussian blur in this case defines the image scale peak response.

Super Structure Module

Push Back noise using the Isolate preset with a mask

1. Open the 'SuperStructure' module.
2. Create a mask which is initially clear but which you put green in the areas of detail you want to exclude.
3. For long thin wisps (e.g. veil) - create a trail of individual green pixels (Using Single Pixel Toggle Brush Mode) then use Grow until they join up and cover the whole area of detail.
4. Press 'Isolate' preset.
5. Increase 'Mask Fuzz' to 3-4.
6. Press 'Do'.
7. Press 'Keep' when happy with the result.

Star Shrinking Under Tracking

(by Richard Konrad - notes by Ivo Jager)

1. Use Autodev
2. Perform 'Crop' as usual then wipe with a dark anomaly setting of about 3 px.
3. Redo Autodev with an emphasis on either the eastern or western nebula. I adjusted the ROI influence until I could see as much of the central nebula region as possible without enlarging the stars too much.
4. No deconvolution at this point as the image is too noisy (done after tracking and near the very end)
5. I did a linear mask of the stars (fat stars) with default setting then hit 'grow' one or two times - 'keep'
6. I masked the stars a second time under (stars) with default settings. Make sure you select 'add new to old'. At this point you will notice that the stars and noise are selected. Click 'shrink' (maybe twice) to remove the selected noise. Invert mask so stars are *not* selected
7. Under 'SuperStructure', select 'Isolate' and 'Do'. This will emphasize all but the stars. The following screen shots show the result just after '3. Autodev', then after '7. SuperStructure-isolate'

8. To further separate the nebula from stars, you will do a second iteration of isolate. First highlight all visible parts of the nebula using the default 'flood fill lighter pixels'. This is the fiddly part where it may take time to find faint central parts - be generous looking for all the 'red'. Click grow once.
9. So as not to emphasize the stars within the nebula, mask the stars with default settings but 'subtract new from old'. Once done, invert the mask to view the masked stars. If noise is also masked, click shrink then grow.
10. Invert then apply 'isolate'. Following is the result after 2nd isolate.
11. To sharpen details in the nebula, keep the same mask.
12. At this point, or even earlier, I prefer to desaturate the image so I can clearly see how dark the background is. In this case I might apply a small amount of equalization under the 'HDR' module.
13. Following, go to the color module and then disengage tracking for noise reduction.
14. In my case, I found further emphasizing the faint parts of the nebula were achieved by applying the default setting of 'life' over the whole image but changing the glow threshold to 0 and bumping up the 'airy disk radius' 2/3 towards the end. This gave a nice glow but kept more detail.

The Isolate and DimSmall presets in the 'SuperStructure' module are great for isolating larger scale nebulosity from a busy star field (or noise)! It really is extremely useful if you want to 're-focus' the viewer's attention on latent nebulosity, as it imparts coherence to the scene.

If you use it with a mask to selectively apply it (rather than globally), just be aware that you're getting into a 'grey area' with regard to documentary value, as you are now making guesses (though very educated!) at where nebulosity exists and where it doesn't exist - all without having recorded it through an instrument in order to substantiate your decisions. It's not exactly 'photoshopping' or 'painting on' nebulosity, but it may be frowned upon by some if presented as a scientific rendition of a dataset.

Of course, for a lot of us, astrophotography is more about making amazing pictures, rather than hard-core science. And where you stand on this is totally up to you!

'Color' Module

Color Balancing data filtered by a Light Pollution 'Filter'

This approach is summarized in the article Color balancing of data that was filtered by a light pollution filter.

1. Color balancing by sampling of filtered data will not give meaningful results.
2. Shoot luminance data with the light pollution filter in place.
3. Shoot color data without the filter in place.
4. Process both images separately.
5. Combine in 'Layer' module as described in the topic Mel 15 (in this example it combines luminance and color for Ha and RGB but the techniques are the same).

Color Balancing data collected using narrow band filters (e.g. Hubble Palette)

For cases where you have used the LRGB module to load the data collected using SII, Ha and OIII narrowband filters to the R,G and B channels respectively.

1. This is not intended to be true color so color balancing becomes a matter of taste.
2. Optionally create a weighted synthetic luminance frame - As described in the article: Ha,R,G,B -> Synthetic Luminance.
3. Adjust the relative proportions of SII, Ha and OIII using the R, G and B bias sliders until you get a balance that you like.
4. 'Keep' the result.

Adjusting Color of Stars and other features separately

Sometimes you want to adjust the color of the stars separately from the rest of the image:

1. Create a Star Mask.
2. In 'Color' module click 'Sample' - this sets the color balance assuming the average star color is white.
3. Invert the mask - this will apply the white balance to the other (non-star) features.
4. Adjust the saturation of the other features.
5. Make any other changes to the color balance you want.
6. Invert the mask - so it selects the stars again.
7. Adjust the saturation of the stars.
8. Make any other changes to the color balance you want.
9. 'Keep' the result.

Creating solar image in false Color

There is a false color palette, specifically for solar imaging.

1. Process your monochrome solar image as normal
2. When you get to the Color module:
3. Set the Matrix parameter to "False Color: Solar" - this will create a yellow false colour image as a starting point.
4. Adjust the bias sliders to change the hue. Reducing the green bias slider makes the yellow less prevalent and brings out the reds.
5. Make any other changes to the colour you want.
6. 'Keep' the result.

'Shrink' Module

Reduce bloom around bright stars

This is a way of using the module to reduce the halo or bloom around bright stars:

1. Create Star Mask.
    - Open Mask editor and click Auto - Stars
    - Set Selection Mode to 'Highlights>Threshold'
    - Set Threshold to about 90% - depending on the brightness of the stars you want to include.
    - Click 'Do' - The star cores will be selected (green).
    - Click Grow as needed to expand to cover the whole star.
    - Keep the result when done.
2. Set Mode to 'Tighten' - This mode will remove the halo that is around the star (as long as it is within the mask).
3. Set Iterations to whatever you like (more iterations draws color in further).
4. You might also want to increase the 'grow mask' parameter 1 or 2 pixels if not all the halo is selected.
5. 'Keep' the result.

'Filter' Module

Removing Purple Fringes around stars

For a full description of this technique see Fringe and halo killer.

1. Put the stars with fringes in a mask - making the stars green. Be sure to include the fringe.
2. Set the 'Filter' Mode to 'Fringe Killer'.
3. Set the 'Filter' Width to 1 - the lower value makes the filter more responsive to colors close to (but not exactly the same as) the selected color.
4. Click on the colored star halos in different places to eliminate the fringes completely.
5. 'Keep' the result when you are done.

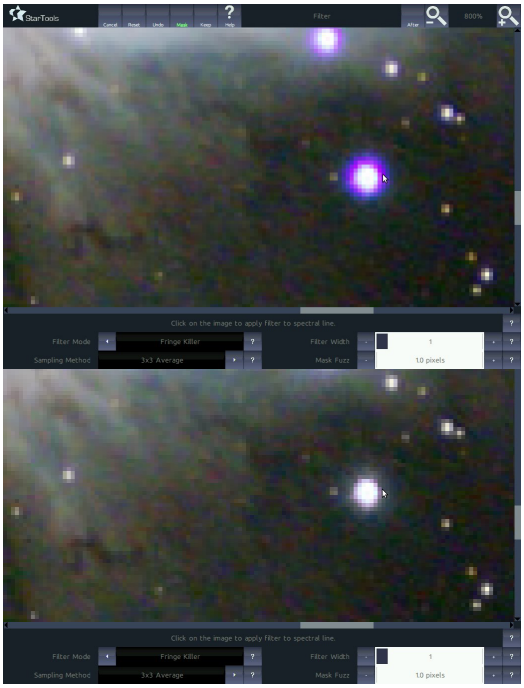

'Denoise' Module

Selecting optimum noise reduction settings

In cases where you are struggling to find the right noise settings this approach may help.

1. Select an area which includes background and large scale structures.
2. Increase Brightness Detail Loss to 30%, and reduce Scale Correlation to 2 - this allows us to see the effects of our changes.
3. Increase Grain Dispersion from 4.5 in increments until there is no further discernible smoothing of the background noise.
4. Increase Scale Correlation from 2 to 6 to see the increase in detail in the larger structures - stop when the level of detail is about right.
5. Reduce the Brightness Detail Loss - making sure the background noise is smoothed enough.
6. Keep the result.

Denoising only parts of the image

You can use the 'Layer' module to work on the "before" and "after" image of any module.

1. Switch off Tracking
2. Perform noise reduction with 'Denoise'
3. 'Keep' the result
4. Launch 'Layer' module
5. click 'Undo -> Background' to put the image as it was before denoise in the background layer
6. Use a mask to mix the background (not denoised) and foreground (denoised) as needed.

Denoise Separate Color Data with LRGB

The trick with color data is to apply heavy noise reduction. The human eye is much less sensitive to loss of color detail than it is to loss of luminance detail. Since, with LRGB, all the luminance detail will come from our luminance frame, we can be heavy handed with our noise reduction. Often this data set also has had a much shorter total exposure time allocated to RGB, so it is quite noisy.

In Unified Denoise Module:

1. Set all Scale settings 1-5 to 100%.
2. Set Brightness Detail Loss to 100%.
3. Set Color Detail Loss to 100%.
4. 'Keep' the result.
5. Merge this with the luminance frame.

Perform Denoise outside of Tracking

Note: This is not recommended as standard flow, but might help in special scenarios.

1. Switch on Tracking again
2. Indicate the image is non-linear to take the sRGB stretch into account
3. Turn off tracking and perform Noise reduction

'Flux' Module

Removing background noise using 'Flux' module

Should exposures be short and stacks have few subs, there is always a lot of background noise that should be eliminated as much as possible while respecting the details of the objects of interest. 'Fractal Flux / Noise' preset might help a lot when adjusted suitably.

1. Launch the 'Flux' Module
2. Select the 'Noise' preset
3. Increase or decrease the 'Positive Flux' parameter. This will process the bright area. Stars get a bit rounded by decreasing this.
4. Increase or decrease the 'Negative Flux' parameter. This will process the dark area. Stars get compressed and a bit rounded by increasing this.
5. 'Detail Filter' will have little effect
6. 'Brightness Mask Power' will have little effect
7. 'Filter Amount' will regulate the intensity of the filter, lower it until the brightness does not make the star core too fat.
8. 'Filter Radius' determines the radius up to which the filter will act. The larger setting, the larger the affected radius of the star will be.
9. 'Filter Fuzz' eliminates noise according to the size of the pixel that we put.
10. 'Keep' the result when you are happy with it.

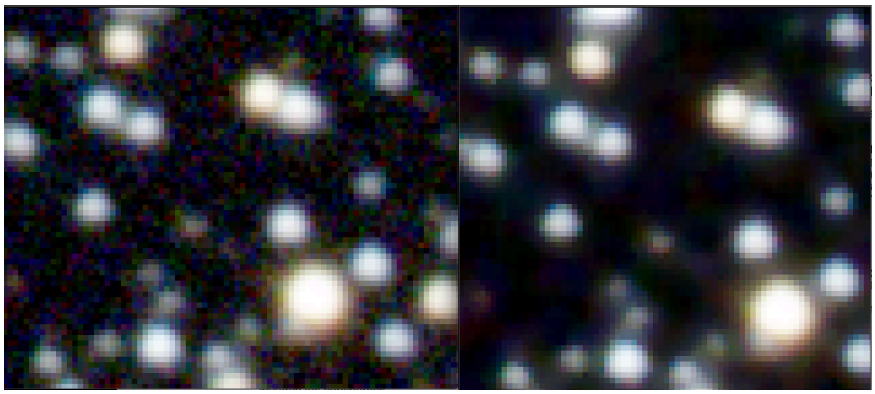

Before and after Fractal Flux / Noise

'Heal' Module

Using the 'Heal' module to remove satellite trails

See also the description in the topic 'Lens' flare and artifact removal.

1. Open the 'Heal' module.
2. Create a mask for the satellite trail by Clearing the mask and using the 'Line Toggle (Click & Drag)' Brush Mode and clicking the start and end of the trail.
3. Click 'Grow' a few times until the green line covers the whole trail.
4. 'Keep' the result
5. Set 'New Must be Darker Than' to 30%.
6. Set 'Grow Mask' to 3 or so as needed.
7. Set 'New Darker Than Old' to Yes.
8. Keep other settings at default Values.
9. The trail will be removed. 'Keep' the result when you are happy with it.

Creating a Starless Image

See also the description in the topic Creating a starless image.

1. Launch the 'Heal' module.
2. Create a mask for the stars - Mask-Auto-Stars-Do.
    - If the mask picks up details that are not stars - in 'Auto', adjust 'Filter Sensitivity' as needed (Note: lower values are more sensitive - i.e. more features detected).
    - If the mask picks up detail which is in a particular color channel you can ignore that channel. To do this - in 'Auto', adjust 'Exclude 'Color" to stop selection of elements of that color channel (channels are - Red, Blue, Yellow (Red + Green) and Purple (Red+Blue)).
    - Manually adjust the mask to remove any remaining masking of non-stars.
    - When you are happy with the mask 'Keep' it - we are then back in the 'Heal' module.
3. Set 'New Darker Than Old' to 'Yes'.
4. Set 'Grow Mask' as needed to ensure all pixels are masked.
5. The 'Heal' module should remove the masked stars with all other settings at their default values.
6. Go back and adjust the mask if there are outstanding stars or star halos.
7. 'Keep' the result when you are happy with it.

The module that removes the stars and substitutes them for a plausible background is the 'Heal' module. No need to change any of its default settings.

If you'd like to present your image without stars, then the quality of the result will depend on how well you can create a star mask that covers all stars.

To create a mask, start off with the default mask (Mask, Auto, Stars, Do) and see how well that works. If you find that the mask generator also picks up details that are not stars, you can increase (not decrease - slightly counterintuitive, we know!) the Sensitivity parameter. If you have detail in a particular color channel, you can tell the mask generator to not look in those color channels. For example, if your nebula is mostly red, blue and purple, you can use the 'Exclude Color' parameter in the Mask generator to disregard red, blue or purple details.

Additionally, you will probably want to do some manual touching up of the mask, either by manually switching off parts of the mask that contains non-star detai, or by fixing up very bright stars that have halos around them. Fortunately, the latter is not that hard - to select the halos, click next to the stars (on the boundary of the halo and the start of empty space), while 'Flood Fill Lighter Pixels' is selected for 'Brush Mode'. This will keep

looking for pixels (and filling them) that are brighter than the halo pixel that you clicked, effectively progressively filling every pixel going into the center of the star (since stars appear progressively brighter as you get closer to their cores) - very handy!

Doing all this for the image of IC1848, after a few minutes of manual labor, we end up with the following mask

Using the 'Heal' module to process stars independently

See also the description Using 'Heal' to process stars and background independently.

1. Create a starless image using a star mask in 'Heal'.
    - 'Keep' the starless background image when you are happy with it.
    - Save the starless image for use later.
2. Create a stars-only image:
    - Launch the 'Layer' module.
    - Click the preset 'Undo->Bg' - this sets the background (left panel) to what is in the undo buffer - which is the original image.
    - Set the 'Layer Mode' to 'Subtract' - and the result (right panel) is the extracted stars.
    - 'Keep' the resulting stars-only image.
    - Save the stars-only image - we will need it later.
3. Process the starless image:
    - Click 'Undo' - reverts to the image before the last modules' changes - to restore the starless image.
    - Process the starless image as you want (e.g. 'Sharp', 'SuperStructure', 'Denoise', 'HDR' etc.) until you are happy with it.
4. Merge the stars back in:
    - Launch the 'Layer' module - This loads the starless image in the foreground (centre panel).
    - Click 'Open' and select the saved stars-only image - This will load into the foreground (centre panel) and the starless image becomes the background (left panel).
    - Set the 'Layer Mode' to 'Add'.
    - The resulting merged image is in the right panel.
    - Click 'Keep' to keep the merged image.

Removing dust bunnies using the 'Heal' module

The 'Heal' module was partly designed for this task.

1. Put the bunny in a mask.
2. Run the 'Heal' module and set the 'New must be darker than' value to a percentage that approach the brightness of background level (it may be a little bit higher). You can do this by eye. I.e. if the mean background level brightness is about 10% of full intensity (white) then set it to 15% or so.
3. You'll see the 'Heal' module fill in a plausible background. The 'New must be darker than' value is simply there to make sure it is not adding 'plausible' stars as well. If you see it do that, just lower the value.
4. Once you're happy with the 'plausible' background, it's time to add back in any stars that were healed out. For this, launch the 'Layer' module, click Undo->Fg and set the 'Layer Mode' to 'Lighten'.

'Repair' Module

Improving the Stellar Profile

This technique is useful in cases where the stars don't look natural. It is described in the video M8 in 'Color' with modest data between 13m32s and 15m11s.

1. Load the 'Layer' module - keeping the star mask used by the 'Repair' module.
2. Grow the mask by about 2 pixels or so to make room for the stellar profile to taper off. Mask - Grow.
3. Set 'Layer' Mode to 'Lighten' and apply a small amount of Gaussian Blur - 'Filter' Type: 'Gaussian (Fg)' - 'Filter' Kernel Radius - say 3-4 pixels.
4. Optionally increase the Mask Fuzz to blur the edges further.
5. Experiment with the settings until it looks right. Keep the result.

'Layer' Module

Combining images of different exposure lengths

This is useful when you have a subject with a high dynamic range and you can't, in a single image, get the exposure right which gives you the detail you want in the fainter portions without saturating the brighter portions.

This Technique is shown in video StarTools: M42 H-alpha High Dynamic Range composite.

1. Make sure the images are the same size and aligned the same as described in the Background Notes.
2. Process each image separately - Use Develop instead of 'AutoDev' - save each result.
3. Restart StarTools.
4. Load both images into 'Layer' Module - with the longer exposure in the foreground (centre).
5. Set the 'Layer' Mode to 'Blend'.
6. Use the 'Filter Type' 'Max Contrast' or 'Max Distance to 1/2 Unity'.
7. You can combine the two by copying the result (Copy) of one filter, pasting to the Foreground (Paste>Fg), and then apply the other filter.
8. Adjust the Blend Amount (to 50% or 100% perhaps).
9. Increase the 'Filter' Kernel Radius until a smooth blend is achieved (can be near max. of 51 pixels).
10. You can repeat the process by copying the result and Paste to the Foreground (Paste>Fg) (or Background).
11. Keep the result.

Combining Luminance and RGB images using the 'Layer' module

This is a generic method to combine (synthetic) luminance and color data as described in Mel 15. It is an alternative to the method 'Processing Ha,R,G,B using a synthetic luminance frame' described in LRGB Module Use.

1. Make sure the images are the same size and aligned the same as described in the Background Notes.
2. Process luminance data as normal - trying to tease out the detail - and save it.
3. Process the R+G+B color stack, making sure the noise is minimal and there is a good continuum of colors - then save it.
4. Restart StarTools.
5. Load the luminance data.
6. Start the 'Layer' module - the current image gets loaded into both foreground and background.
7. Load the color data - this will replace the image in the foreground.

8. Select 'Layer Mode' 'Color Extract fg' to extract the luminance independent (normalized) colors from the color image.
9. Click Copy to store the composite output in the buffer.
10. Click Paste>Fg to paste the extracted colors into the foreground layer.
11. Select the 'Layer Mode' 'Color' Of fg' (or 'Multiply') to use the luminance information from the background and the color of the foreground.
12. If darker parts of the background are too bright - select Brightness Mask Mode 'Where fg is dark, use bg' or 'Where composite is dark, use bg'.
13. Use the Brightness Mask Power to subdue the coloring of the background.
14. Use the Blend Amount to increase/decrease saturation.
15. 'Keep' the desired result.

Creating a synthetic luminance frame

Creating a synthetic luminance frame, derived by combining the results from different narrow-band filters while they are linear and unprocessed, allows you to process the luminance and the color separately. The luminance is processed to enhance the detail, the other frames are processed to minimize noise - see M45 Advanced Processing In StarTools Part 2.

To create the synthetic luminance frame we combine all the data we have into a weighted average - weighting according to the exposure time collected in each channel.

Example: R:G:B:Ha 30m:30m:30m:110m T1:T2:T3:T4

See also the description: Ha,R,G,B ->synthetic luminance.

We load one file and then add the other channels one by one using the 'Layer' module.

1. Make sure the images are the same size and aligned the same as described in the Background Notes.
2. Start StarTools.
3. Open R file - indicate the data is not linear (even though it is) by selecting 'Modified and not linear' - This turns Tracking off.
4. Open the 'Layer' module.
5. Open G - Set blend to T2/(T1+T2)% - which is 30/60=50% in the example.
6. Copy, then Paste>Bg
7. Open B - Set blend to T3/(T1+T2+T3)% - which is 30/90=33% in the example.
8. Copy, then Paste>Bg
9. Open Ha - Set blend to T4/(T1+T2+T3+T4)% - which is 110/200=55% in the example.
10. Either: Save the image to a file and open as a linear file later - or turn Track on.

We now have our weighted average synthetic luminance frame.

Hubble palette synthetic luminance frame example

Ha:SII:OIII 110m:40m:40m T1:T2:T3

1. Make sure the images are the same size and aligned the same as described in the Background Notes.
2. Start StarTools.
3. Open the Ha File - indicate the data is not linear.
4. Open the 'Layer' module.
5. Open the SII file - Set blend to T2/(T1+T2)%=40/150=27%.
6. Copy, then Paste>Bg
7. Open the OIII file - Set blend to T3/(T1+T2+T3)%=40/190=21%.
8. Either: Save the image to a file and open as a linear file later - or turn Track on.

Color balancing of data that was captured using a light pollution filter

Data that was captured using a light pollution filter is difficult to color balance. In this technique the color data is captured without a light pollution filter. The luminance data is

a longer exposure using a light pollution filter. The method described here uses the 'Layer' Mode and a Brightness Mask to combine the two data sets.
See also Color balancing of data that was filtered with a light pollution filter.

1. Make sure the images are the same size and aligned the same as described in the Background Notes.
2. Load the color data.
3. Launch the 'Layer' module.
4. Select 'Layer Mode' 'Color' Extract fg' to extract the colors from the color image.
5. Click Copy to store the output in the buffer.
6. Open the luminance image.
7. Click Swap - swaps background and foreground (putting luminance in the background).
8. Click Paste->Fg to paste the extracted colors into the foreground layer.
9. Select 'Layer Mode' 'Color' Of fg' to use the luminance information from the background and the color of the foreground.
10. Select the Brightness Mask Mode 'Where fg is dark use bg'.
11. Use the Brightness Mask Power to subdue the coloring of the background.
12. Use the Blend Amount to increase/decrease saturation.
13. Use the 'Filter' Kernel Radius (with Gaussian filter) to mitigate color noise.
14. Keep the result.

Creating Bi-Color images using Steve Cannistra's method

This is a description of a method proposed by Steve Cannistra Modified Bicolor Technique for combining Ha and OIII images which can produce fairly 'natural' results. You may get better results if you preprocess the images separately and then combine them as described here.

1. Create Synthetic Green Image:
    - Make sure the Ha and OIII images are aligned and the same size in pixels.
    - Load Ha data - indicate the data is not linear (even though it is) by selecting 'Modified and not linear' - This turns Tracking off.
    - Open the 'Layer' Module.
    - Load OIII data into foreground.
    - Set 'Layer' Mode to 'Multiply'.
    - Set Blend Amount to 100%.
    - Keep.
    - Save as a Tiff file.
2. Restart StarTools.
3. Load the LRGB Module.
4. Load the Ha data into the Red Channel.
5. Load the saved Synthetic Green image into the Green Channel.
6. Load the OIII data into the Blue Channel.
7. Increase the Green Ratio to about 1.50 - or as you prefer - this increases the influence of the Synthetic Green data.
8. Keep.
9. When prompted for Type of Data select 'Linear, was not bayered or is whitebalanced'.
10. Process with other modules as normal - 'Crop', wipe etc.
11. When using the 'Color' Module with narrow-band data the automatic color balance that happens on loading may produce odd results. Set the Bias Sliders back to 1.00

and then adjust from there until you get the result you want. Adjust the Saturation Amount if you want.

Splitting a color image into Luminance and RGB before Processing

The technique was described by Kevin Bisher and based on Scott Rosens LLRGB method for DSLR image processing - which used PhotoShop. It allows us to process luminance and color separately which can produce superior results. The luminance is processed to enhance the detail, the RGB is processed to minimise noise - see M45 Advanced Processing In StarTools Part 2

1. Load the image.
2. 'Bin' and 'Crop' the Image as needed.
3. Save the image - this is the RGB version.
4. Create a Luminance (monochrome) version of the image (one way is: 'Layer' -> 'Layer' Mode=Desaturate fg (Luminance) -> Keep) and save it.
5. Load and Process the RGB version using heavy denoise - save it.
6. Load and Process the Luminance image - concentrating on bringing out the detail.
7. Merge images using the 'Layer' module:
    - Make sure the images are the same size and aligned the same as described in the Background Notes.
    - Load the 'Layer' module - Luminance image will be loaded into centre (foreground) and left (background) panel.
    - Load the RGB image - this puts the RGB image in foreground (Luminance stays in background).
    - Set the 'Layer' Mode to 'Blend'.
    - Set the Blend Amount - reduce from 100% as needed.
8. Do any further processing needed to the combined image.
9. Optionally - Load the Luminance image again, swap and put it in the background, 'Blend' again, and process the combined image. This is the main additional step in the LLRGB method.
10. You can repeat this process with the result to further improve it.
11. Keep the result.

Screen Mask Invert (SMI)

This is an old method used by PhotoShop/PixInsight users to increase contrast in the shadows. Normally other StarTools functions are better but, if needed, StarTools can achieve this as described in SMI

1. It is best run just before the 'Color' Module (while Tracking is still on!)
2. Launch the 'Layer' module.
3. Set 'Layer Mode' to 'Screen'.
4. Set 'Brightness Mask Mode' to 'Where foreground is light, use background'.
5. Keep the result.

Power of Inversed Pixel (PIP)

This is an old method used by PhotoShop/PixInsight users to increase contrast in the shadows. Normally other StarTools functions are better but, if needed, StarTools can achieve this as described in PIP.

It is best run just before the 'Color' Module (while Tracking is still on!).

1. Launch the 'Layer' module.
2. Set 'Layer Mode' to 'Power of Inverse'.
3. Set 'Brightness Mask Mode' to 'Where foreground is light, use background'.
4. Keep the result.

Using StarNet++ with StarTools on linear data

- The point of using this procedure is to obtain a star mask, using StarNet++'s star detection. All I/we know is that StarNet++ only works on (e.g. has been trained with) stretched, processed images.
- It is always good to ask yourself if removing stars is the best course of action - there may be better ways to achieve your aim.
- See the link above for further discussion on when and how to use this.
- The source image in this case, is the "processed" image you would normally feed to StarNet++. E.g. this source image has been stretched to bring out all the stars so that StarNet++'s neural net can "see" them.
- StarNet++ only works with "finished" images. However, with the guide below, we can use StarNet++'s smarts to successfully process a linear dataset in StarTools. The crux of the method, is the creation of a perfect star mask, by having StarTools look at which pixels were modified by StarNet++.
- Once that star mask has been extracted, StarTools' heal module should actually yield superior results in terms of removal, as the Heal module is agnostic when it comes to the replacement of stars. StarNet++'s neural net will be only as good as the stars it has been trained on when it comes to replacing them.

- Once the star mask is created you can use it in a number of ways - removing, healing, separating, recombining etc.
- If using LRGB data then do the process on the luminance or synthetic luminance only.

Creating the Starnet++ Starless Image:

1. Fully process the image from the raw data as normal in StarTools and save that (stretched) image
2. Use that image as input to StarNet++
3. Process
4. Save the resulting StarNet++ output (starless) image.

Creating the Star Mask using the StarNet++ output:

1. Load the (stretched) StarNet input (e.g. with stars still in place) image in StarTools. Don't turn Tracking on.
2. Launch the Layer module.
3. Load the (stretched) StarNet output (e.g. with stars removed) image in the Layer module.
4. Set Layer Mode to "Difference".
5. You could skip to step 9 here, but the following steps may be needed if the Starless image has been normalized.
6. Click Copy.
7. Click Paste Foreground.
8. Set Layer Mode to "Multiply Foreground Only".
9. Set Blend Amount to 500%.
10. Keep the result.
11. Launch the FilmDev module.
12. Set Skyglow to 49%.
13. Keep the result.
14. Save the image (e.g. starnet-mask.tif).

Now use the Star Mask to remove the stars from linear data using the Heal module:

1. Now load your original raw data into StarTools once more. Don't turn Tracking on - When prompted select 'Non-linear sRGB source' then 'Don't activate Tracking'.

2. In the Mask editor, load the image you just saved as a mask. StarTools will automatically convert it into a mask and all stars should be correctly selected at this point.
3. Launch the Heal module.Your stars will be healed out using the mask.
4. Keep the result.

Starting with a good dataset

StarTools will not work correctly (or work poorly) with an incorrectly stacked dataset. Getting a suitable dataset from your free or paid stacking solution, is extremely important.

Important pre-processing do's and don'ts

Important dataset preparation do's and don'ts

There are a few simple, but important, do's and don'ts to prepare your dataset for post-processing in StarTools.

Learning how to use a new application is daunting at the best of times. And if you happen to be new to astrophotography (welcome!), you have many other things, acronyms and jargon to contend with too. Even if you consider yourself an image processing veteran, there are some important things you should know. That is because some things and best practices play a bigger role in StarTools than in other applications.

Really, these all boil down to making sure your is as virgin as possible. Note that doesn't mean noise-free or even good, it just means you have adhered to all best the below conditions and practices to the best of your abilities.

If you are new to processing

When learning how to process astrophotography images, the last thing you want to do, is learning all sorts of post-processing tricks and techniques, just to work around issues that are easily avoidable during acquisition or pre-processing. Fixing acquisition and pre-processing issues during post-processing, will never look as good, while you will also not learn much from this; it is likely whatever you learn and do to fix a particular dataset learn is likely not applicable to the next.

Conversely, if your dataset is clean and well calibrated according to best practices, you will find workflows much more replicable and shorter. In short, it is just a much better use of your time and efforts! You will learn much quicker and you will start getting more confident in finding your personal vision for your datasets - and that is what astrophotography is all about.

If you are an image processing veteran

When we say StarTools requires the most virgin dataset you can muster, we really mean it. It means no procedures or modifications must be done by any software - no matter how well-meaning. No gradient reduction, no color balancing, not even normalization (if not strictly necessary). Signal evolution Tracking - the reason why StarTools achieves objectively better results than other software - absolutely requires it.

Checklist

- Make sure your dataset is as close to actual raw photon counts as possible.
- Make sure your dataset is linear and has not been stretched (no gamma correction, no digital development, no levels & curves)
- Make sure your dataset has not been normalized (no channel calibration or normalization) unless unavoidable due your chosen stacking algorithm
- Make sure all frames in your dataset are of the same exposure length and same ISO (if applicable)
- Make sure your dataset is the result of stacking RAW files (CR2, CR3, NEF, ARW, FITS, etc.) and not lossily compressed or low-bit depth formats (e.g. not JPEGs or PNGs).
- Make sure no other application has modified anything in your dataset; no stretching, no sharpening, no gradient reduction, no normalization

- If you can help it, make sure your dataset is not color balanced (aka 'white balanced'), nor has had any camera matrix correction applied
- Flats are really not optional - your dataset must be calibrated with flats to achieve a result that would be generally considered acceptable
- Dithering between frames during acquisition is highly recommended (a spiraling fashion is recommended, and if your sensor is prone to banding, you will want to use larger movements)
- If you use an OSC or DSLR, choose a basic debayering algorithm (such as bilinear or VNG debayering) in your stacker. Avoid 'sophisticated' debayering algorithms meant for single frames and terrestrial photography like AHD or any other algorithms that attempt to reconstruct detail.
- If using an mono CCD/CMOS camera, make sure your channels are separated and not pre-composited; use the Compose module to create the composite from within StarTools and specify exposure times where applicable.

Some common problems in StarTools caused by ignoring the check-list above

- Achieving results that are not significantly better than from other software
- Trouble getting any coloring
- Trouble getting expected coloring
- Trouble getting a good global stretch
- Halos around dust specks, dead pixels or stacking artifacts
- Faint streaks (walking noise)
- Vertical banding
- Noise reduction or other modules do not work, or require extreme values to do anything
- Ringing artifacts around stars
- 'Color' artifacts in highlights (such as star cores)
- Trouble replicating workflows as seen in tutorials and/or videos

Allowed / not allowed

Allowed in your dataset:

- Noise grain
- Light pollution
- Sky gradients

Ideally avoided:

- Vignetting
- Dust specks, dust donuts
- Smudges
- Amp glow
- Dead pixels, dead sensor columns
- Satellite trails
- Trees or buildings
- Walking noise and other correlated noise (e.g. noise that is not single-pixel speckles)

Do's and don'ts

Do:

- Process your dataset from start-to-finish in StarTools including compositing (LRGB, LLRGB, SHO, HOO, etc.)
- Use simple workflows
- Acquire and apply flats
- Dither between frames during acquisition
- 'Bin' your dataset if your dataset is oversampled
- Use deconvolution to restore detail if possible
- Practice with some publicly available datasets that are of reasonable quality to get a feel for what a module is trying to do under normal circumstances

Don'ts:

- Do not post-process any part of your image in any way, in other application before opening it in StarTools
- Do not make composites in any other application but StarTools
- Do not process any part of your subs any way, in other application before stacking them
- Do not visit the same modules many times
- Do not process your dataset at higher resolution than necessary
- Do not drizzle your dataset in your stacker if your dataset is already oversampled

Preprocessing Images for StarTools

StarTools works best with images that are Linear and without preprocessing apart from registration and stacking.

- No stretching, noise reduction, color balancing or deconvolution.
- When stacking, use median or some form of sigma clipping with more than 10 sub-frames.

If using DeepSkyStacker (DSS):

- In 'Register Checked Pictures...' - 'Stacking parameters...' button:
- 'Light' tab: Change calibration from 'Per Channel Color Calibration' or 'RGB Channels Calibration' to 'No Background Calibration'.
- 'Result' tab: Uncheck 'Align RGB Channels in Final Image'.
- In Options - Settings... - 'RAW/FITS DDP Settings...' - 'RAW Files' tab:
  - Uncheck 'Use Auto White Balance'.
  - Uncheck 'Use Camera White Balance'.
  - Uncheck 'Set black point to 0'.
- When saving the stacked image - in the save dialog choose:
- Save as type: 'TIFF Image' or 'FITS image', '16 bit/ch' or '32 bit/ch - integer' (a 32-bit file is twice the size of a 16-bit file).
- Compression:'None' and Options:'Embed adjustments...' - otherwise you will have problems processing using StarTools.

- For other settings see the topic DSS settings for StarTools.
- DSS always white balances - unless you preprocess with DCRAW using this method: Eliminating white balancing when using DSS.

If using PixInsight:

- Rejection algorithm: Sigma clipping or Winsorized Sigma Clipping.
- No white balance. Don't use 'BackgroundNeutralization' or 'Color Calibration' processes.

- Save as: TIFF: 16-bit, or FITS: either 16-bit unsigned integer or 32-bit unsigned integer.
- StarTools can take over before you use 'Dynamic Crop' or 'DynamicBackgroundExtension' (DBE) processes. See the topic Welcome, Pixinsight Users! for advice on how to use StarTools as an alternative to PixInsight for post-processing.

If using Nebulosity:

- Align and Combine: Standard Deviation based stacking: Sigma-clip.
- Don't do post processing such as Digital Development, Adjust 'Color' Background or Auto 'Color' Balance.
- Edit menu - Preferences:
  - - 'Save as compressed FITS' - unticked,
  - - 'Save in 32-bit floating point' - unticked,
  - - 'Scale to 15 bit (0-32767) at Save' - unticked.
  - - 'Color File Format' - 'RGB Fits: Maxim'
  - - 'Manually Override color reconstruction' - tick - thgis will stop white balancing - set Manual Demosaic Setup.
  - - 'Demosaic (Debayer) method' - Bilinear.
- Save as: TIFF: uncompressed 16-bit (File - Save 16 bit/color TIFF File), or FITS: uncompressed 16-bit integer (File - Save Current File).

If using MaximDL:

- During the Stack process:
  - 'Color' tab: Use Defaults (1:1:1) color balance, Don't do post processing such as 'Auto equalize background'.
  - Combine tab: Combine Method: Sigma-clip.
- Save as: TIFF: uncompressed 16-bit, or FITS: uncompressed, either 16-bit int or 32-bit int.

The importance of not White Balancing (Color balancing) when preprocessing.

With White Balanced ('Color' Balanced) images the R, G and B channels are scaled differently:

- This makes it impossible to derive an accurate luminance channel for the separate processing of luminance and color data.
- Noise levels are changed - Noise cannot be tracked as accurately.
- Scaling of R,G and B can cause clipping - impacting the ability to extract detail in the highlights.

File Formats for Preprocessed files

Notes on the choice of file format and bit-depth when saving files after preprocessing:

- To avoid any possible loss of data it is best to save as 32-bit Integer FITS if you have the memory/processing power/disk space. See the topic 'DSS save as choice'.
- However, most of the time you will not notice the difference if you save as 16-bit TIFF/FITS.
- Most applications will use the output from StarTools and not the preprocessed output. StarTools output is 16-bit TIFF (or JPEG).
- If you want to use the preprocessed output with applications other than StarTools - choose TIFF or FITS depending the applications you use.
  - FITS is standard for astronomical applications. It can be read by Startools, PixInsight and FITS Liberator but not MS Windows' own apps. Photoshop requires a Plug-in.

    - - TIFF files (uncompressed) can be viewed by MS Windows' own apps, StarTools, PixInsight and PhotoShop (but see below).
- Saving as 32-bit may capture a little more detail than 16-bit in some cases but further post processing (e.g.in StarTools) takes more processing power and memory, and the files are up to double the size.
- If using 32-bit - choose 32-bit integer over 32-bit rational unless it is not supported by one of the applications you need to use (see below).
- The file formats read by common applications are as follows:
    - - StarTools uses 64-bit integer arithmetic internally - using integer avoids mapping issues.
    - - StarTools and PixInsight read 16-bit, 32-bit integer and 32-bit rational files. - FITS Liberator reads 16-bit, 32-bit integer and 32-bit rational FITS format files.
    - - PhotoShop does not read 32-bit integer files.
    - - Nebulosity does not read 32-bit TIFF or 32-bit integer FITS files.
- Any intermediate preprocessed files are normally stored as 32-bit - don't change this.

Debayering/Interpolation Methods

A brief summary of major methods:

- Nearest neighbor - uses value of nearest pixel.
- Bi-linear interpolation - average of the 4 nearest diagonal pixels of same color (R,G or B) - a good choice if there is noise. For StarTools this is the best choice as it allows StarTools to do better noise control.
- VNG - Variable Number of Gradients - noise in each frame can cause the VNG algorithm to introduce artifacts and 'bleeds' the noise to surrounding pixels as described in the topic 'Using debayered integrated images'. This method limits StarTools' noise control.
- PPG - Patterned Pixel Grouping - fewer artifacts than VNG.
- AHD - Adaptive Homogeneity-Directed - causes artifacts especially if there is noise. Don't use unless the subs have a high SNR. This method limits StarTools' noise control.
- Bi-cubic interpolation - uses 16 neighbor pixels of the same color.

If you are unsure which to choose - select bi-linear interpolation if available.

Important Background Information

Tracking

Tracking in StarTools is the name given to the way in which StarTools gathers information about the signal and its evolution through different modules. It provides each module with as much information as possible to allow it to get the best results.

Each module can:

- Understand how each pixel has been modified by previous modules.
- Influence data earlier in the processing chain (e.g. linear data) and re-apply the modifications made since then.
- Use the information gained from previous modules to understand how the signal has been changed and where the noise is.

Some benefits of Tracking:

- Deconvolution only works on linear data - but the 'Decon' Module is used after the data has been stretched.
- With the 'Decon' module we use it after stretching and processing the data but apply deconvolution on the linear data and watch its effect on the processed image.
- Noise reduction is applied at the end of processing where Tracking has gained the most information about noise.
- This means Noise reduction can be automatically targeted most at the areas where it is needed.

StarTools Log File

- Every time you process an image, the processing steps are saved to a log file named StarTools.log in the StarTools program folder.
- It contains entries for:
  - StarTools version used.
  - Date processed.
  - Source file name and location.
  - Modules used along with the associated parameter values.
  - Saved file name and location.
- It is very useful if you want to recreate a processing workflow or to help with troubleshooting.
- For versions earlier than v1.3.5 it was saved in the folder where the source image file is located.

Debayering Artifacts

Artifacts are caused by errors in estimating the 'missing' colors.

Look out for the following effects in debayered color images:

- Zipper effect - at edges.
- False Color - usually at edges - fringing.
- Halo - at edges.
- Aliasing - interference pattern.
- Mazing - fine parallel lines confuse demosaic algorithm producing an effect that looks like a rectangular maze.
- Blurring

Configuration

Various basic setting may be configuered using files with the below naming in the StarTools folder:

- '**noblink**'
  - Stop Mask blinking 3 times as a reminder it is set: 'noblink'
  - file contents ignored
- '**highdpi**'
  - Make StarTools buttons etc. bigger
  - only uses the real-estate to achieve high DPI (4K), but not more
  - file contents ignored
- '**largeui**'
  - Make StarTools use the full real estate available to it and make buttons etc. bigger:
    file contents ignored
- '**opencldeviceindex.cfg**'
  - Override default OpenCL device selection:
  - For use in the StarTools GPU version to manually select the device for systems with multiple devices from same vendor
  - Put 0 or 1 in file
  - Putting 0 will force StarTools to use the "0th" device it finds, which should be your iGPU
- '**openclplatformindex.cfg**'
  - Override default OpenCL vendor selection: 'openclplatformindex.cfg'
  - For use in the StarTools GPU version to manually select platform/vendor/driver on systems with multi-vendor GPUs
  - Put 0 or 1 in file
- '**openclforcecpu.cfg**'
  - Force selection of CPU not GPU
  - For use in the StarTools GPU version when there is misbehaving OpenCL GPU hardware/drivers
  - On the bottom of the splash screen at startup it shows whether it is using your CPU and GPU, or just the CPU
  - file contents ignored

Image Analysis and Evaluation

This discussion covers analyzing the image quality and tackle its imperfections. Such an understanding will also help in guiding how best to improve your technique and (over time) equipment. At each of these steps in the imaging process, imbalances, noise, inefficiencies, etc. can arise. Identify and learn to spot them or their hallmarks, and find a software or hardware solution that eliminates or ameliorates them.

Checking sub-frames and stacking options

After an imaging session - as a minimum - start with a set of Lights, Darks, Flats & Bias (or Dark Flat) frames.

1. Check the exposure length of the sub-frames:
    - Only stack subs with the same exposure time and ISO setting - if necessary combine them later in Startools using the 'Layer' module- see creating a multi exposure length 'HDR' composite.
    - If combining subs using median or sigma modes (anything but additive) the sub exposure must be long enough to catch some signal photons on most of the light frames so that the signal is not considered as noise by the stacking process.
    - Light pollution can tempt you to reduce the exposure time but often the signal is lost. Stacking and post-processing techniques can help to reduce the effect of light pollution, and compensate for saturated areas, but there is no way to make up for lack of signal.
    - If there are some over-exposed stars they can be improved by using the 'Shrink' ('Shrink') module.
    - Look at the luminance histogram Is there a high count on left of the histogram and nothing on the right? If so, it may benefit from a longer exposure time.
    - Is there much saturation (a high count on the far right of histogram) showing the exposure time was perhaps too long?
    - If both left and right have high counts then the object has a high dynamic range ('HDR') - consider collecting sets of subs with different exposure times - stack them separately and combine later in StarTools.
2. Look at the light sub-frames:
    - If there are hot or stuck pixels or fixed pattern noise:
    - If you do dithering between sub-frames (which is recommended) then these will be managed if you stack using a form of Sigma clipping with at least 10 subs. See this external article on Dithering.
    - Without dithering, dark frames will get rid of the hot or stuck pixels.
    - If there is horizontal or vertical banding: If you do dithering this should be much reduced after stacking - also stacking with a set of bias frames may help with banding. If you don't currently dither between sub-frames then consider doing so in the future.
    - If there are satellite trails, cosmic rays and other small blemishes in individual frames: These will be managed if you stack using Median, or a form of Sigma clipping with at least 10 subs.
    - Otherwise remove the sub-frame with the defect - or, if you only have a few sub-frames, see if you can remove it using the 'Heal' module during post-processing.
    - If there are any that are low grade or have any other individual artifacts that don't appear in the majority of images (e.g. focus issues, cloud, mist - etc. - etc.) discard these sub-frames.
    - Discarding the subs is better than introducing errors which will be difficult to remove later - however, if you only have a few subs leaving in subs with minor

problems may improve the overall signal to noise ratio (SNR) more than leaving them out.

3. For OSCs and DSLRs:
    - When debayering, use Bilinear interpolation rather than VNG or AHD - This allows StarTools to control noise better.
    - Improve detail by avoiding debayering - if you have used spiral dithering during capture and have a large number of sub-frames you can stack using Bayer-drizzle (DSS) - See the topic Using debayered integrated images.
    - For noisy images - consider whether the image is sufficiently oversampled that you can afford to reduce the resolution by a factor of 4 - if so consider stacking using Super-pixels (DSS) or color binning (Nebulosity) after reading this note by Michael Covington.
4. If using DSS:
    - Consider preprocessing with DCRAW to remove the white balancing - as described DSS vs Nebulosity vs ? and Eliminating white balancing when using DSS.
    - Consider saving the output as a FITS (32 bit/ch - integer) rather than a TIFF file (if not already) - see discussion M13 Globular cluster.
    - You can also change the Autosave file format by setting Stacking Parameters - Intermediate Files - 'Intermediate and Final Image File Format' to FITS - See Difference between TIFF and FITs files.
5. For DSLRs, check whether the optimum ISO was used:
    - The optimum ISO for a camera is one where the sensor output is not artificially amplified to emulate improved sensitivity.
    - Above this setting you are not gaining any benefit in capturing detail and the dynamic range may be reduced.
    - See the external links: DSLR Exposure, Best ISO setting for Astrophotography, DSLR Camera Data.
    - If necessary, make a note to change it for next time.

Study the stacked image. First, look for the problems:

1. Check if image has been white balanced
    - RAW sub-frames will not be white balanced, JPEG will.
    - Stacking can introduce white balancing (Deep Sky Stacker (DSS) always white balances when using RAW subs, PixInsight and Regim do not).
    - White balanced images often show a red, brown or yellow bias. Non white-balanced images are teal or blue-green.
2. This will affect how you open the image in StarTools:
    - If it is a (stacked) RAW image - either monochrome or white-balanced color - select 'Linear, was not bayered or is white balanced'.
    - If it is a (stacked) RAW color image that is not white balanced - select 'Linear, was Bayered, is not white balanced'.
    - If it is a JPEG use 'Modified and not linear'.
3. Use the 'AutoDev' module and look at the initial image - this is specially designed to show up any color casts, vignetting and artificial patterns such as banding. See this example using M51.
    - Are the corners darker? - this is a sign of vignetting - using flats will help.
    - Can you see darker blotches? - These may be due to dust specks on the sensor - best way to remove these is to use flats. In the absence of flats try a method covered in the 'Heal' module description.

- Is there a Zipper or Checkerboard pattern - a possible indication of a debayering issue - see the topic 'Checkerboard Pattern in Image' on homepage.
- Zoom in to see the noise
    - Is there a lot of noise? - For images with low signal to noise - increase the 'Ignore Fine Detail' setting to stop 'AutoDev' responding to the noise. See also later references to noise.
    - Is the background quite uniform or are there colored or darker single-pixel spots? If there are dark spots then consider increasing the Dark Anomaly 'Filter' when using the 'Wipe' module.

4. Look for stacking artifacts - be sure to 'Crop' the image to remove these before using the 'Wipe' module.
5. Look for any banding - vertical or horizontal:
    - Ideally, if you can, restack using a set of bias frames.
    - Otherwise plan to try the 'Wipe' module to fix this.
    - In the future stacking with bias frames, or dithering during capture, may help.
6. Look at the background for Skyglow
    - Does the background have a color bias (skyglow or light pollution):
        - Red or yellow/brown cast - skyglow that has been white balanced.
        - Teal, blue or green cast - skyglow that has not been white balanced.
        - Bright blue-green cast - skyglow filtered using a light pollution filter.
        - Missing yellow (e.g. no yellow stars) - indicates use of light pollution filter.
    - Try the 'Wipe' Module - if this doesn't work it may be that you need more subs to increase the total integration time - or your subs aren't long enough - to get a reasonable SNR. See CN discussion Deep Sky Imaging, Aperture and Sky Fog Limits.
7. Look at the stars in the centre of the image:
    - Are the stars oval and all oriented in the same direction?
        - Look back at the sub-frames and see if some are worse than others - if so discard sub(s) and re-stack.
        - If all the same this may indicate guiding problems or the imaging plane is not perpendicular to the telescope axis. For now plan to try the 'Repair' module.
    - If round but blurred or doughnuts - see if some are worse than others - if so discard sub(s) and restack.
        - If all the same think about what might cause this (e.g. poor focus) and try and put it right for next time. For now plan to use the 'Decon' module.
    - If there are some stars (or other features) that are over-exposed then they can be improved by:
        - shrinking stars using the 'Shrink' ('Shrink') module or ...
        - using 'HDR' techniques to combine a stack of long exposures with a stack of short exposures - See Creating a multi exposure length 'HDR' composite.
8. Look at the stars at the edge of the image:
    - Are the stars elongated on axis towards the centre of the image (Coma)?
        - Make a note to use the 'Lens' module to reduce this before doing any cropping of the image. For the future a Coma Corrector will reduce this effect.
    - Are the stars at the edge more blurred than at the centre - or vice versa?

        - This is most likely caused by field curvature. For the future using a field flattener may help this.
9. Sampling - is the resolution better than the seeing? (over-sampled) - You can improve the signal to noise ratio by reducing the resolution with the 'Bin' module.
    - To check - look at the smallest stars - if they occupy more than a 3x3 block of pixels then the image is over-sampled.
    - When binning, reduce the resolution so this blurriness is gone. See 'Bin' module usage.
10. Zoom in and look at the background - Look for shadows and marks that could indicate dust, scratches, cosmic rays or satellite trails - make a note to try the 'Heal' module.
    - Remember to mask out dust specks (as well as the subject) when using the 'Wipe' module.
11. Look at noise:
    - Identify the type of noise
        - In most cases by far the biggest noise type is shot noise (aka Poisson noise). 'Correcting' this type of noise is not really possible as such - this noise is fundamentally 'uncertainty' in your signal.
        - We can, however, (in all sorts of clever ways) pretend that uncertainty is not there by making educated guesses about what the signal would look like if we modeled that uncertainty and then removed it.
        - The only unique thing that StarTools does versus any other software is that it makes sure to keep an exact handle on that modeled uncertainty at all times, even when your data is being processed, stretched and modified. Therefore noise reduction in ST is using the right model at all times, applicable to the image as you currently see it. This in turn makes for extremely targeted noise reduction across the image.
    - Is there a lot of noise in the background? - try using the 'Isolate' preset in the 'SuperStructure' module later using a full mask as described in the topic First Astro Photograph.
    - Zoom in - does the noise have a color bias (skyglow or light pollution?) or is it random colors (equipment generated)?
    - Are there signs hot pixels have not been eliminated by stacking? Go back and use Sigma or Kappa Sigma stacking to eliminate stuck pixels or use a bad pixel map if your stacker allows.
    - Is the noise level such that additional binning of the image would look better - in spite of the loss of resolution? See 'Bin' Module Use.
        - There is a balance: the noise level is always such that a binned image would look better. There is always more detail just over the horizon, waiting to be brought out, if only your data was that little bit better. At the end of the day, it's up to you to decide of course. If you want a deep, high-fidelity post stamp, that's your artistic prerogative. The cool thing is that binning gives you that freedom to make that trade-off.
    - If noisy data then the deconvolution ('Decon') module needs to be used with caution - it works best when the image has little noise.
    - Finally, the De-noise module (Wavelet De-Noise - when you switch Tracking off) will help eliminate noise - adjust the Grain Size until the noise grain can't be seen - don't worry about the signal - StarTools has been tracking it and will protect it.
    - If noise is predominantly color blotches - set 'Brightness Detail Loss' to 0% in the Wavelet De-Noise module.

12. Look at the stars and other features:
    - Signs of Chromatic aberration - try using 'Lens' or 'Filter' module later
        - See the Fringe and halo killer - or
        - Make a synthetic Luminance channel from the green channel to reduce bloat from CA in the blue channel as described in the topic Sh2-82 and reducing stars.
    - Signs of Purple fringing?
        - Try using 'Lens' or 'Filter' module later - or
        - Consider using one of the approaches for processing the blue channel separately as discussed in the topic fringe killer filter add blue back to central star.
        - This may throw automatic color calibration off - try using 'Max RGB' mode when color correcting.
13. Look at the RGB histogram:
    - Is there saturation/clipping of colors (high count on far right of histogram)?
    - Is there much green? This will also show up in 'Max RGB' tool in the 'Color' module - Plan to fix this using 'Color' Module - See 'Color' module - color balance.

Now look at the qualities of the image:

Are there any reasons to use something other than your 'standard' workflow?

1. Subject - modify process depending on the subject.
    - Imaging Planets - perhaps from a video source or jpeg - try using the process described Planetary Images and LRGB stacks.
    - Imaging Comets - remove the blurry stars in the background - either when stacking by using sigma clipping, or using the 'Heal' module as described in the topic Comet Lovejoy/2014 close up or as described in the Background Notes in 'Heal' Module use.
    - Lunar imaging - when using the 'Decon' module set 'Image Type' to 'Lunar/ Planetary' - see the discussion in the topic 'Colorful Aristarchus' on homepage.
2. Object characteristics (high dynamic range, low dynamic range, faintness, etc.) :
    - Is it very high dynamic range - are there saturated areas? - merge 2 images with different exposure times using 'Layer' Module (need to capture more data) - See Creating a multi exposure length 'HDR' composite.
    - Busy star field in background - try using the Isolate preset in the 'SuperStructure' Module to de-emphasize it.
3. Image framing - does the image need to be cropped - for framing or to get rid of edge effects?
4. Are there special non-celestial elements in the image? - Trees, mountains, dust specks - Remember to mask out these object(s) when using the 'Wipe' module.
5. Atmospheric conditions - e.g. seeing is not good - Using 'Decon' module will help here - see 'Decon' Module use.
6. Details - are there specific features you want to bring out?
    - Some nebulosity for example - See Using 'Heal' to process stars and background independently
    - Faint Detail - Try creating a synthetic luminance channel using this approach or this alternative approach.
    - Small to medium structures - use the 'HDR' module to locally adjust the stretch to bring out small to medium details.

    - Medium to large structures - use the 'Contrast' module to locally adjust the stretch to bring out medium to large details.
    - Large Structures - Try using the 'SuperStructure' module to emphasise large structures.
7. Colors - what colors are there?
    - Are they real?
        - Full range of star colors - Older (red, orange, yellow) and younger (white and blue)
        - Hot ionized gas clouds - HII regions Red (Ha) or Purple (Ha & Hb), SII and NII Red, OIII cyan
        - Reflection Nebulae - Dust and cool (non-ionized) gas - can scatter light from nearby source (e.g. star) - blue scattered better than red
        - How to ensure they are preserved during processing? See Color balancing techniques in 'Color' module description
    - Are they wrong - like green in the wrong place? - For green use Cap Green in the 'Color' module. For other colors try using the 'Filter' module or ...
    - Are you using the Hubble Palette or similar? - you may want to adjust the channel balance in the 'Color' module using the Bias Sliders.
8. Combining RGB and Ha data - consider using the approach described in the topic H,R,G,B data -> synthetic luminance.
9. Shooting at different times (moon out) - try using the techniques described above to handle Skyglow.
10. Shooting from a different location - do you have to work to control the effect of light pollution? Are you using a light pollution filter? - See Color Balancing using a light pollution filter which is more fully described in the topic 'Color' balancing of data....
11. Equipment change or retuning (e.g. mount, flexure, insertion of filter etc.).

Introduction

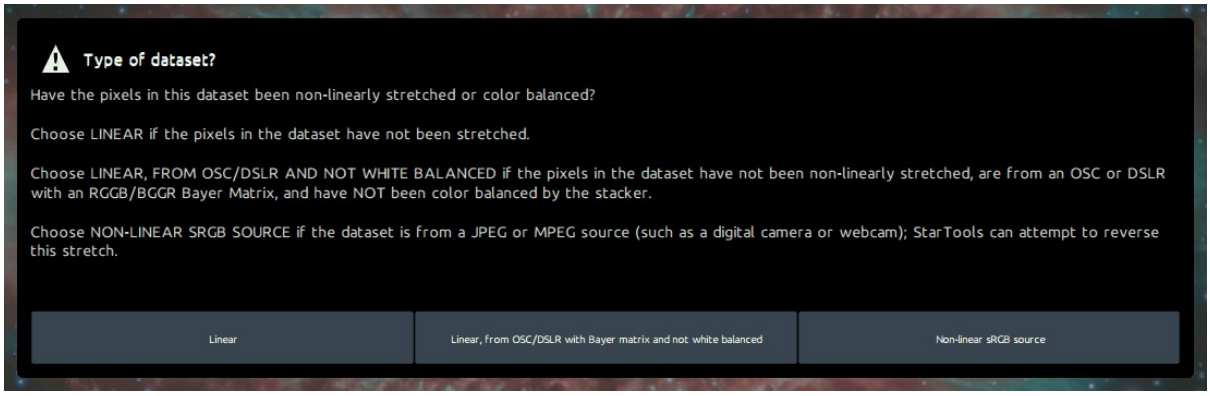

Tracking

Signal evolution Tracking starts as soon as you load your dataset.

Signal evolution Tracking data mining plays a very important role in StarTools and understanding it is key to achieving superior results with StarTools.

As soon as you load any data, StarTools will start Tracking the evolution of every pixel in your image, constantly keeping track of things like noise estimates, parameters you use and other statistics.

Tracking makes workflows much less linear and allows for StarTools' engine to 'time travel' between different versions of the data as needed, so that it can insert modifications or consult the data in different points in time as needed ('change the past for a new present and future'). It's the primary reason why there is no difference between linear and non-linear data in StarTools, and the reason why you can do things in StarTools that would have otherwise been nonsensical (like deconvolution after stretching your data). If you're not familiar with Tracking and what it means for your images, signal fidelity and simplification of the workflow & UI, please do read up on it!

'Signal evolution Tracking data mining plays a very important role in StarTools and understanding it is key to achieving superior results with StarTools.'

Tracking how you process your data also allows the noise reduction routines in StarTools to achieve superior results. By the time you get to your end result, the Tracking feature will have data-mined/pin-pointed exactly where (and how much) visible noise grain exists in your image. I therefore 'knows' exactly how much noise reduction to apply in each area of your image.

Noise reduction is applied at the very end, as you switch Tracking off, because doing it at the very last possible moment will have given StarTools the longest possible amount of time to build and refine its knowledge of where the noise is in your image. This is different from other software, which allow you to reduce noise at any stage, since such software does not track signal evolution and its noise component.

Tracking how you processed your data also allows the 'Color' module to calculate and reverse how the stretching of the luminance information has distorted the color information (such as hue and saturation) in your image, without having to resort to 'hacks'. Due to this capability, color calibration is best done at the end as well, before switching Tracking off. This too is different from other software, which wants you to do your color calibration before doing any stretching, since it cannot deal with color correction after the signal has been non-linearly transformed like StarTools can.

The knowledge that Tracking gathers is used in many other ways in StarTools, however, the nice thing about Tracking is that it is very unobtrusive. In fact, it actually helps get you

get better results from your data in less time by homing in on parameters in the various modules that it thinks are good defaults, given what Tracking has learnt about your data.

The best kept secret amongst signal processing purists

StarTools monitors your signal and its noise component, per-pixel, throughout your processing (time). It sports image quality and unique functionality that far surpasses other software. Big claim? Let us back it up.

Your signal and its noise component

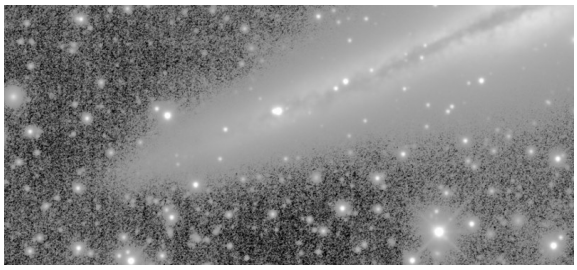

When you stretch your dataset, you will notice noise grain becoming visible quickly in the darker areas.

If you have ever processed an astrophotographical image, you will have had to non-linearly stretch the image at some point, to make the darker parts with faint signal visible. Whether you used levels & curves, digital development, or some other tool, you will have noticed noise grain becoming visible quickly.

You may have also noticed that the noise grain always seems to be worse in the darker areas than the in brighter areas. The reason is simple; when you stretch the image to bring out the darker signal, you are also stretching the noise component of the signal along with it.

And the former is just a simple global stretch. Now consider that every pixel's noise component goes through many other transformations and changes as you process your image. Once you get into the more esoteric and advanced operations such as local contrast enhancements or wavelet sharpening, noise levels get distorted in all sorts of different ways in all sorts of different places.

The result? In your final image, noise is worse in some areas, less in others. A 'one-noise-reduction-pass-fits-all' no longer applies. Yet that's all other software packages - even the big names - offer.

In this example, brightness has been deliberately modified locally (top right), enhancing contrast. Unfortunately, traditional software cannot know that area has a much higher signal-to-noise ratio. As a result noise reduction treats all areas the same, destroying much detail.

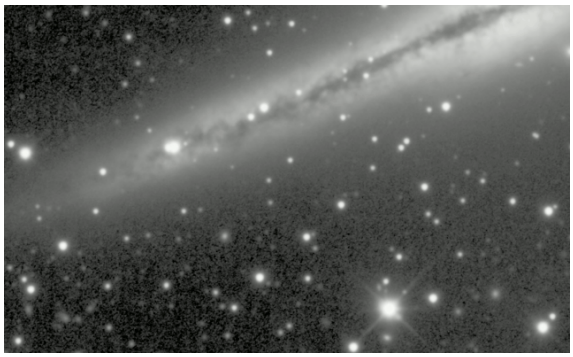

Current image processing software for astrophotography is fundamentally broken

Chances are you have used a noise reduction routine at some stage. In astrophotography, the problem with most noise reduction routines, is that they have no idea how much worse the noise grain has become in the darker parts. They have no idea how you stretched and processed your image earlier. And they certainly have no idea how you squashed and stretched the noise component locally with wavelet sharpening or local contrast optimization.

'The separation of image processing into dumb filters and objects, is one of the biggest problems for signal fidelity in astrophotographical image processing software today.'

In short, the big problem, is that separate image processing routines and filters have no idea what came before, nor what will come after when you invoke them. All pixels are treated the same, regardless of their history. Current image processing routines and filters are still as 'dumb' as they were in the early 90s. It's still 'input, output, next'.

Without knowing how signal and its noise component evolved to become your final image, trying to, for example, squash noise accurately is impossible. What's too much in one area, is too little in another, all because of the way prior filters have modified the noise component beforehand.

The separation of image processing into dumb filters and objects, is one of the biggest problems for signal fidelity in astrophotographical image processing software today. It is the sole reason for poorer final images, with steeper learning curves than are necessary. Without addressing this fundamental problem, 'having more control with more filters and tools' is an illusion. The IKEA effect aside, long workflows with endless tweaking do not make for better images.

But what if every tool, every filter, every algorithm could work backwards from the finished image, and trace signal evolution, per-pixel, all the way back to the source signal? That's Tracking.

Tracking is the solution

Tracking in StarTools makes sure that every module and algorithm can trace back how a pixel was modified at any point in time. It's the Tracking engine's job to allow modules and algorithms 'travel in time' to consult data and even change data (changing the past) and then forward-propagate the changes to the present.

The latter sees the Tracking module re-apply every operation made since that point in time, however with the changed data as a starting point; changing the past for a better future. This is effectively signal processing in three dimensions; X, Y and time (X, Y, t).

The power of 3D (X, Y, t) signal processing

Correct deconvolution of extremely noisy, stretched, locally contrast-enhanced (top left) data. No further masks (other than simple auto-generated star mask), local supports or selective processing was performed. Noise grain is correctly identified and ignored. Only areas with sufficient SNR are enhanced.

Deconvolution - an example

This remarkable feature is responsible for never-seen-before functionality that allows you to, for example, apply deconvolution to heavily processed data. The deconvolution module 'simply' travels back in time to a point where the data was still linear (normally deconvolution can only correctly be applied to linear data!). Once travelled back in time, deconvolution is applied and then Tracking forward-propagates the changes. The result is exactly what your processed data would have looked like with if you had applied deconvolution earlier and then processed it further.

Signal evolution Tracking allows for many more enhancements over traditional software. For example color constancy (right), effortlessly visualizing features with similar chemical/physical properties, regardless of brightness.

Sequence doesn't matter any more, allowing you to process and evaluate your image as you see fit. But wait, there's more!

Deconvolution - an example that gets even better

Time traveling like this is very useful and amazing in its own right, but there is another major, major difference in StarTools' deconvolution module.

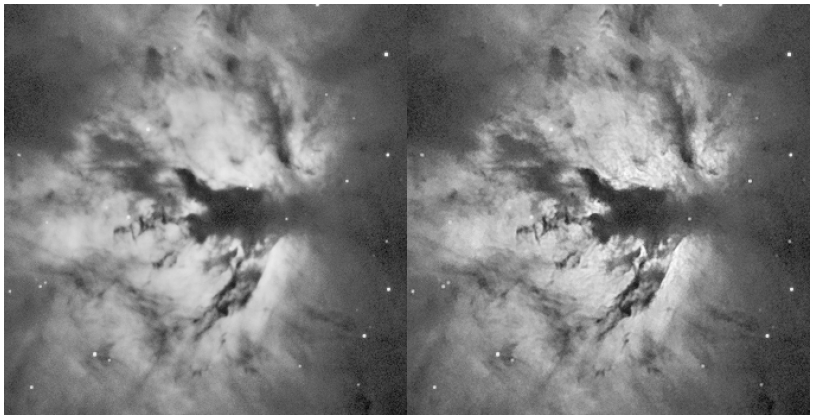

200% zoomed crop input and output. Left: original, right: 'Decon' result. Thanks to signal evolution Tracking, and despite stretching, local dynamic range optimization and noise presence, 'Decon' is able to recover the finest details without introducing artifacts, or the need for support masks or manual intervention.

The major difference, is that, because you initiated deconvolution at a later stage, the deconvolution module can take into account how you processed the image after the moment deconvolution should normally have been invoked (e.g. when the data was still linear). The deconvolution module now has knowledge about a future it normally is not privy to in any other software. Specifically, that knowledge of the future tells it exactly how you stretched and modified every pixel - including its noise component - after the time its job should have been done.

You know what really loves per-pixel noise component statistics like these? Deconvolution regularization algorithms! A regularization algorithm suppresses the creation of artifacts caused by the deconvolution of - you guessed it - noise grain. Now that the deconvolution algorithm knows how noise grain will propagate in the 'future', it can take that into account when applying deconvolution at the time when your data is still linear, thereby avoiding a grainy 'future', while allowing you to gain more detail. It is like going back in time and telling yourself the lottery numbers to today's draw.

What does this look like in practice? It looks like a deconvolution routine that just magically brings into focus what it can. No local supports, luminance masks, or selective blending needed. No exaggerated noise grain, just enhanced detail.

And all this is just what Tracking does for the deconvolution module. There are many more modules that rely on Tracking in a similar manner, achieving objectively better results than any other software, simply by being smarter with your hard-won signal.

Mathematical implementation of the signal evolution Tracking engine

In conventional processing engines, every pixel as-you-see-it is the result of the operation that was last carried out (some simple screen stretch capabilities excepted to visualize linear data). Operations are carried out one after the other and exist in some linear stack (typically accessible via an 'undo' history). The individual operations however, have no concept of what other operations preceded them, nor what operations will follow them, nor what the result was or will be. Signal flows one way in time: forward. Conventional software does not feed back signal, nor propagates it back and forth in order to refine the final result of the stack or 'undo' history.

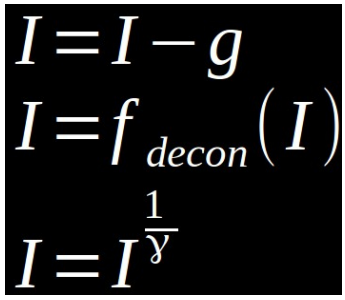

Image processing in a traditional application: every step uses the output of the previous step as its input. This workflow shows a sequence where the user subtracts a gradient "g", performs deconvolution on the linear data and then performs a non-linear gamma correction stretch.

Some software platforms even mistakenly implement astronomical signal processing in a formalized object oriented platform. An object oriented approach, by definition, implements strict decoupling of the individual operations, and formalises complete unawareness of the algorithms contained therein, with regard to where and when in the signal flow they are being invoked. This design completely destroys any ability of such algorithms to know what augmenting data or statistics may be available to them to do a better job. Worse, such software allows for entirely nonsensical signal flows that violate mathematical principles and the physics these principles are meant to model. The result is lower quality images through less sophisticated (but more numerous) algorithms, rounding errors, user-induced correction feedback loops (invoking another module to correct the output of the last), and steeper learning curves than necessary.

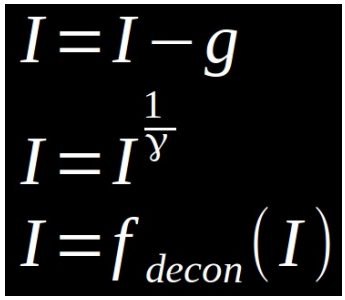

An "illegal" sequence where deconvolution is erroneously applied after a non-linear gamma correction stretch.

In contrast, StarTools works by constantly re-building and refining a single equation, for every pixel, that transforms the source data into the image-as-you-see-it. It means there is no concept of linear versus non-linear processing, there are no screen stretches with lookup tables, there is no scope for illegal sequences, there is no overcooking or noise grain/artifact propagation, there are no rounding errors. What you see is the shortest, purest transformation of your linear signal into a final image. And what you see is what you get.

In StarTools, your image as-you-see-it is the result of an ever changing equation that is mathematically sound and simplified. This equation shows a sequence where the user subtracts a gradient "g", performs a gamma correction, and performs deconvolution (noting that the deconvolution step is normally considered out-of-sequence). As a result of the equation building, non-linear vs linear processing is completely abstracted away.

$$I = I_{origin} - g$$

$$I = (I_{origin} - g)^{\frac{1}{\gamma}}$$

$$I = (f_{decon}(I_{origin} - g))^{\frac{1}{\gamma}}$$

Even more ground-breaking: substituting some of its variables for the equation itself (or parts thereof), allows complex feedback of signal to occur. This effectively provides, for example, standard algorithms like deconvolution or noise reduction, precise knowledge about a "future" or "past" of the signal. Such algorithms will be able to accurately calculate how the other algorithms will behave in response to their actions anywhere on the timeline. The result is that such algorithms are augmented with comprehensive signal evolution statistics and intelligence for the user's entire workflow. This lets these algorithms yield greatly superior results to that of equivalent algorithms in conventional software. Applying the latter innovation to - otherwise - standard, well known algorithms is, in fact, the subject of most of StarTools' research and development efforts.

$$I_{decon}^{n+1} = f_{decon}(I_{decon}^{n}, f_{regularization}((I_{decon}^{n} - g))^{\frac{1}{\gamma}})$$

An algorithmic tour de force; enhancing a standard Richardson & Lucy deconvolution iteration by basing the regularization step on a forward propagated version of the previous iteration, thereby taking into account artifact propagation in the user's full workflow in the stretched, processed domain; this makes an algorithm like deconvolution privy to statistics that are far outside its normal linear-only purview. Substitution and equation solving like this, encompassing the entire workflow, is impossible in conventional software.

The power of StarTools' novel engine, is not only expressed in higher signal fidelity and lifting of limitations of conventional engines; its power is also expressed in ease-of-use. Illegal or mathematically incongruent paths are closed off, while parameter tweaks always yield useful and predictable results. Defaults just work for most datasets, proving that the new engine is universally applicable, consistent and rooted in a mathematically sound signal processing paradigm.

Do it once, do it right

Physics and applied mathematics demand that some operations are done in a particular order. No ifs, no buts. Certain operations have one specific place in a sound signal flow, yet others have less rigid sequence requirements. Whichever your processing decisions, they are worked into the equation in a mathematically congruent way.

The most elegant equation is often the shortest one. In StarTools, you refine the final equation like a sculptor would refine a coarse piece of marble into a sculpture; from coarse gestures to fine tweaks. Module functionality does not overlap in StarTools; you will never be correcting one module's output with another module that does-the-same-thing-but-differently. I.e. the engine's goal is to "tack on" to the equation as little as possible, and to rather tweak its present form and variables as much as possible.

'In StarTools, you should never apply noise reduction to an unfinished image; any further processing will change your image's noise profile again, invalidating your previous noise reduction decisions and efforts'

Less is more. The shorter solution is the better solution. The best part is no part. Endless tweaking is not a thing in StarTools, and all decisions and module invocations are meant to be done with a clear direction, decisiveness and purpose. Feeling a sense of closure on an image is a feature, not a bug.

A good example of the "do it once, do it right" philosophy that StarTools' engine affords, is its approach to noise reduction. In StarTools you don't need to "help" any of the algorithms by noise reducing earlier in your workflow and passing them noise reduced versions of your datasets. All modules are fully noise-aware. As such, in StarTools, noise is an aesthetic consideration only. Noise grain only becomes a problem if it is visible and aesthetically objectionable. Therefore noise reduction is only applied at the very last moment, when it is at its most visible and most objectionable. In StarTools, you should never apply noise reduction to an unfinished image; any further processing will change your image's noise profile again, invalidating your previous noise reduction decisions and efforts. As such, there is only one noise reduction tool and one noise reduction moment; the one right tool at the one right moment. That is, a tool that models the noise profile in your image with pin-point accuracy, at the very end of your workflow.

TL; DR

In StarTools, your signal is processed (read and written) in a time-fluid way. Being able to change the past for a better future not only gives you amazing new functionality, changing the past with knowledge of the future also means a cleaner signal. Tracking always knows how to accurately estimate the noise component in your signal, no matter how heavily modified.

User-friendly by mathematical nature

For its unique engine to function, StarTools needs to be able to make mathematical sense of your signal flow. That's why it's simply unable to perform 'nonsensical' operations. This is great if you're a beginner and saves you from bad habits or sub-optimal decisions.

Just like in real life, in astrophotographical image processing, some things need to be done in a particular order to get the correct result. Folding, drying then washing your shirt, will achieve a markedly different result to washing, drying and folding it. Similarly, deconvolution will not achieve correct results if it is done after stretching, ditto for light pollution removal and color calibration. In mathematics, this is called the commutative property.

The 'Tracking' feature, constantly backward propagates and forward propagates your signal through processing 'time' as needed. This means that 'nonsensical' signal paths (e.g. signal paths that get sequences wrong) would break Tracking's ability. Therefore, such signal paths are closed off. For this reason, it is neigh-impossible in StarTools to perform catastrophically destructive operations on your data; it simply wouldn't be sound mathematics and the code would break.

'The 'Tracking' feature, constantly backward propagates and forward propagates your signal through processing 'time' as needed.'

For example, the notion of processing in the linear domain vs non-linear (stretched) domain is completely abstracted away by the engine because it needs to do that. If you didn't know the difference between those two yet, you can get away with learning about this later. Even without knowing the ins-and-outs of astronomical signal processing, you can still produce great images from the get-go; StarTools takes care of the correct sequence.

So, whereas other software will happily (and incorrectly!) allow you to perform light pollution removal, color calibration or deconvolution after stretching, StarTools will...
...actually also let you do that, but with a twist!

Tracking will rewind and/or fast-forward to the right point in time, so that the signal flow to makes sense and is mathematically consistent. It inserts the operation in the correct order and recalculates what the result would have looked like if your decision had always been the case. It's time travelling for image processing, where you can change the past to affect the present and future.

For an in-depth explanation of Tracking, see the Tracking section.

GPU Acceleration and Your Hardware

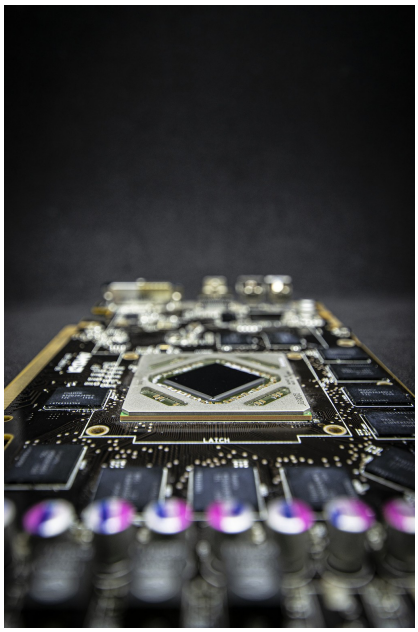

As of version 1.7, heavy arithmetic in StarTools is offloaded to your Graphics Process Unit (GPU). GPUs offer enormous advantages in compute power, under the right circumstances.

Depending on your hardware configuration and module, speed-ups versus the CPU-only version can range from 3x - 20x.

Compatibility

StarTools supports virtually all modern GPUs and iGPUs on all modern Operating Systems.

StarTools is compatible with any GPU drivers that support OpenCL 1.1 or later. Almost all GPU released after ~2012 should have drivers available that expose this API.

StarTools GPU acceleration has been successfully tested on Windows, macOS and Linux with the following GPU and iGPU solutions:

- Nvida GT/GTS/GTX 400, 500, 600, 700, 800M, 900 and 1000 series
- Nvidia RTX series
- AMD HD 6700 series, HD 7800 series, HD 7900 series,R7 series, R9 series, RX series
- Intel HD 4000, HD 5000, UHD 620, UHD 630

Please note that if your card's chipset is not listed, StarTools may still work. If it does not (or does not do so reliably), please contact us.

If you run into instabilities

Not all GPUs, operating systems and GPU drivers are created equal. Some more consumer-oriented operating systems (e.g. Windows), by default, assume the GPU is only used for graphics processing and not for compute tasks. If some compute tasks do not complete quickly enough, some drivers or operating systems may assume a GPU hang, and may reset the driver. This can particularly be an issue on systems with a relatively underpowered GPU (or iGPU) solution in combination with larger datasets. Please see the FAQ section on how to configure your operating system to mimimise this problem.

StarTools' algorithms push hardware to the limit and your GPU is no exception. If your GPU or power supply is aging, StarTools will quickly lay bare weaknesses in thermal and power management. Similarly, laptops with iGPUs or discrete GPUs will have to work harder to rid themselves of waste heat.

If your GPU or iGPU appears to be unstable on your Windows operating system in StarTools, and you think it may be struggling with any larger datasets you give it, then the issue may be caused an unsuitable Timeout Detection and Recovery (TDR) allowance.

TDR is a feature that is meant to prevent GPU "hangs". If a task "hangs" the GPU for longer than 2 seconds, the TDR kicks in and will reset the GPU driver.

This Windows default behavior is not suited for compute-heavy tasks as found in StarTools. Fortunately, it can be corrected by making modifications to the default 2 second timeout value. More information may be found here:

https://www.pugetsystems.com/labs/hpc/Working-around-TDR-in-Windows-for-a-better-GPU-computing-experience-777/

Burst loads versus sustained loads

Depending on your GPU monitoring application, it may appear your GPU is only used partially. This is not the case; your GPU solution is used and loaded up 100% where possible. However, as opposed to other tasks like video rendering or gaming, GPU usage in image processing tends to happen in short, but very intense bursts.

'Depending on your GPU monitoring application, it may appear your GPU is only used partially'

Depending on how your monitoring application measures GPU usage, these bursts may be too short to register. Spikes are usually averaged out over time (usually 1000ms) by your monitoring application (with CPU intermittently doing its thing, leaving GPU momentarily unused). With the GPU loaded only for short times, but the load averaged out over longer periods, many monitoring applications make it appear only partial usage is happening.

If your monitoring application can show maximum values (on Windows you can try GPU-Z or Afterburner, on Linux the Psensor application), you should immediately see the GPU being maxed out. For examples of heavy sustained GPU activity, try the Deconvolution module with a high number of iterations or the Super Structure module.

Only tasks that:

- can be parallelized
- are rather "dumb" in terms of logic (with few if-then-else branches)
- perform a lot of complex calculations
- AND process large amounts of data
- complete in milliseconds (up to a couple of seconds or so)

...are suitable for momentary GPU acceleration. As a result, during processing, you should see processing switch back and forth between CPU and GPU.

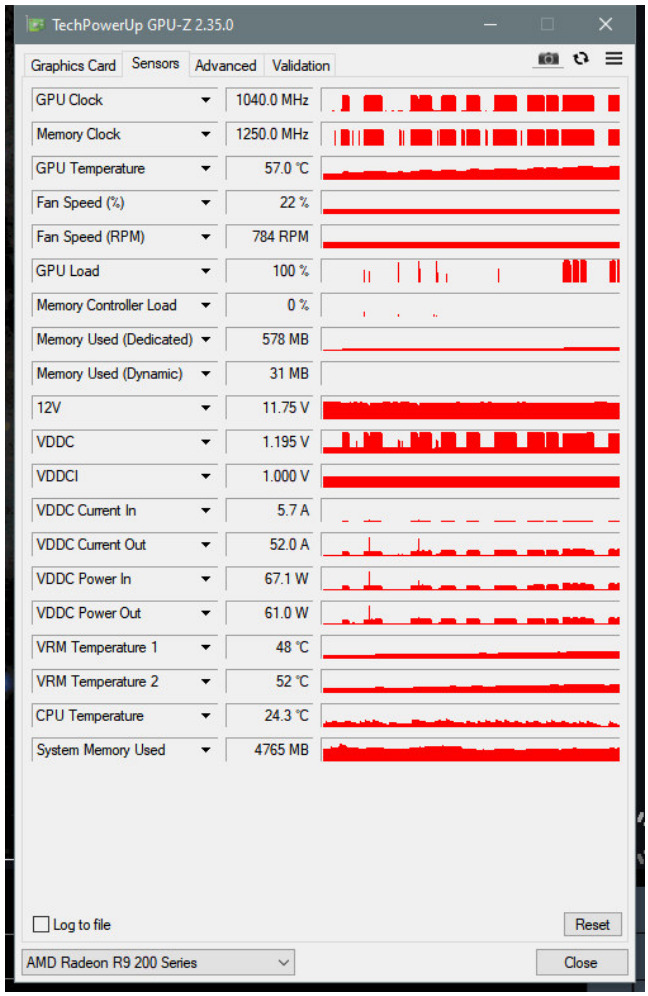

StarTools loads your GPU 100% in bursts

Modules & Features

Log

StarTools keeps a detailed log of what modules and parameters you used. This log file is located in the same folder as the StarTools executable and is named StarTools.log.

As of the 1.4 beta versions, this log also includes the mask you used, encoded in base64 format. See the documentation on masks on how to easily decode the base64 if needed.

Masks

The Mask feature is an integral part of StarTools. Many modules use a mask to operate on specific pixels and parts of the image, leaving other parts intact.

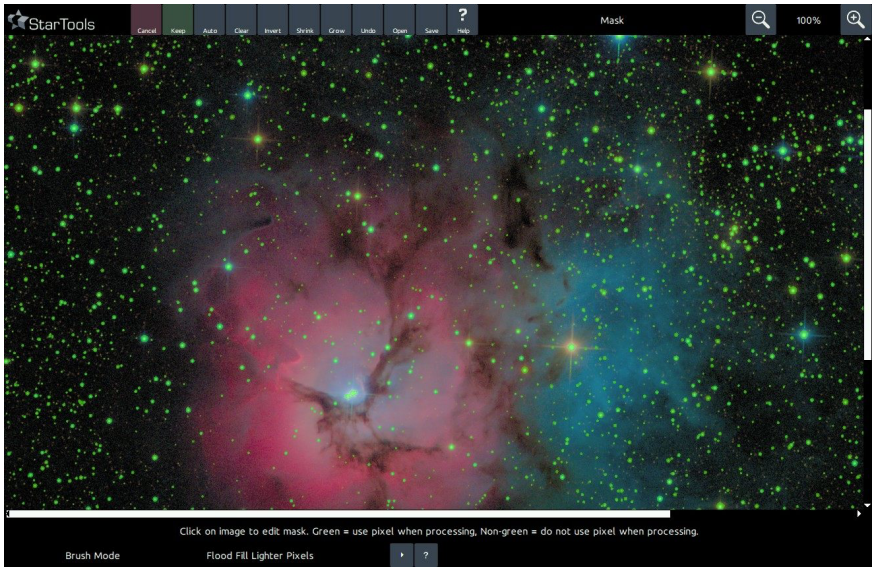

Masking is an integral part of working with StarTools.

Importantly, besides operating only on certain parts of the image, it allows the many modules in StarTools to perform much more sophisticated operations.

You may have noticed that when you launch a module that is able to apply a mask, the pixels that are set in the mask will flash three times in green. This is to remind you which parts of the image will be affected by the module and which are not. If you just loaded an image, all pixels in the whole image will be set in the mask, so every pixel will be processed by default. In this case, when you launch a module that is able to apply a mask, the whole image will flash in green three times.

'You may have noticed that when you launch a module that is able to apply a mask, the pixels that are set in the mask will flash three times in green.'

Green colored pixels in the mask are considered 'on'. That is to say, they will be altered/ used by whatever processing is carried out by the module you choose. 'Off' pixels (shown in their original color) will not be altered or used by the active module. Again, please note that, by default all pixels in the whole image are marked 'on' (they will all appear green).

For example, an 'on' pixel (green colored) in the 'Sharp' module will be sharpened, in the 'Wipe' module it will be sampled for gradient modelling, in 'Synth' it will be scanned for being part of a star, in 'Heal' in will be removed and healed, in 'Layer' it will be layered on top of the background image, etc.

To recap:

- If a pixel in mask is 'on' (colored green), then this pixel is fed to the module for processing.
- If a pixel in mask is 'off' (shown in original color), then tell the module to 'keep the pixel as-is, hands off, do not touch or consider'.

Usage

The Mask Editor is accessible from the main screen, as well as from the different modules that are able to apply a mask. The button to launch the Mask Editor is labelled 'Mask'. When launching the Mask Editor from a module, pressing the 'Keep' or 'Cancel' buttons will return StarTools to the module you pressed the 'Mask' button in.

As with the different modules in StarTools, the 'Keep' and 'Cancel' buttons work as expected; 'Keep' will keep the edited Mask and return, while 'Cancel' will revert to the Mask as it was before it was edited and return.

As indicated by the 'Click on the image to edit mask' message below the image, clicking on the image will allow you create or modify a Mask. What actually happens when you click the image, depends on the selected 'Brush mode'. While some of the 'Brush modes' seem complex in their workings, they are quite intuitive to use.

Apart from different brush modes to set/unset pixels in the mask, various other functions exist to make editing and creating a Mask even easier:

- The 'Save' button allows you to save the current mask to a standard TIFF file that shows 'on' pixels in pure white and 'off' pixels in pure black.
- The 'Open' button allows you to import a Mask that was previously saved by using the 'Save' button. Note that the image that is being opened to become the new Mask, needs to have the same dimensions as the image the Mask is intended for. Loading an image that has values between black and white will designate any shades of gray closest to white as 'on', and any shades of gray closest to black as 'off'.
- The 'Auto' button is a very powerful feature that allows you to automatically isolate features.
- The 'Clear' button turns off all green pixels (i.e. it deselects all pixels in the image).
- The 'Invert' button turns on all pixels that are off, and turns off all pixels that were on.
- The 'Shrink' button turns off all the green pixels that have a non-green neighbor, effectively 'shrinking' any selected regions.
- The 'Grow' button turns on any non-green pixel that has a green neighbor, effectively 'growing' any selected regions.
- The 'Undo' button allows you to undo the last operation that was performed.

NOTE: To quickly turn on all pixels, click the 'clear' button, then the 'invert' button.

Brush Modes

StarTools' mask editor showing the brush modes selection menu.

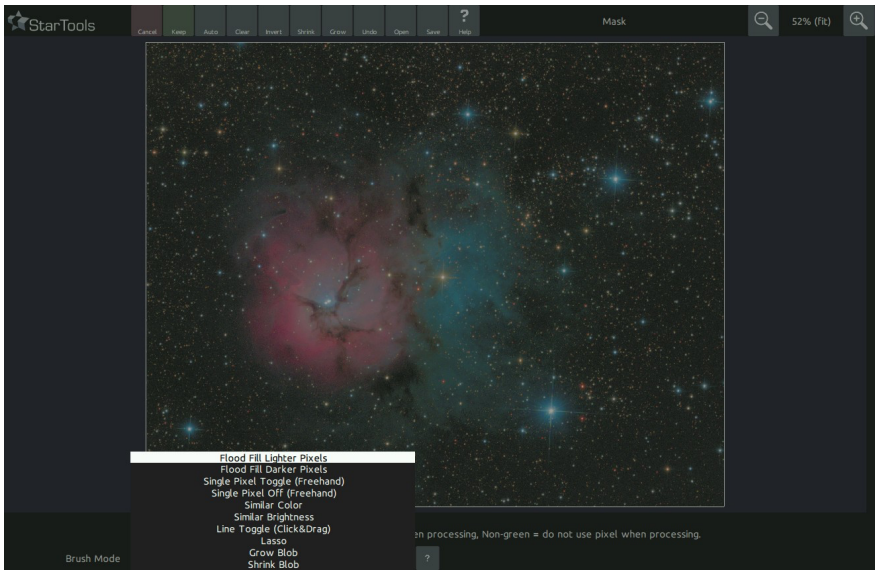

10 different brush modes are at your disposal.

Different 'Brush modes' help in quickly selecting (and de-selecting) features in the image. For example, while in 'Flood fill lighter pixels' mode, try clicking next to a bright star or feature to select it. Click anywhere on a clump of 'on' (green) pixels, to toggle the whole clump off again.

The mask editor has 10 'Brush modes':

- 'Flood fill lighter pixels': use it to quickly select an adjacent area that is lighter than the clicked pixel (for example a star or a galaxy). Specifically, Clicking a non-green pixel will, starting from the clicked pixel, recursively fill the image with green pixels until it finds that; either all neighboring pixels of a particular pixel are already filled (on/green), or the pixel under evaluation is darker than the original pixel clicked. Clicking on a green pixel will, starting from the clicked pixel, recursively turn off any green pixels until it can no longer find any green neighboring pixels.
- 'Flood fill darker pixels': use it to quickly select an adjacent area that is darker than the clicked pixel (for example a dust lane). Specifically, clicking a non-green pixel will, starting from the clicked pixel, recursively fill the image with green pixels until it finds that; either all neighboring pixels of a particular pixel are already filled (on/green), or the pixel under evaluation is lighter than the original pixel clicked. Clicking on a green pixel will, starting from the clicked pixel, recursively turn off any green pixels until it can no longer find any on/green neighboring pixels.
- 'Single pixel toggle': clicking a non-green pixel will make a non-green pixel turn green. Clicking a green pixel will make green pixel turn non-green. It is a simple toggle operation for single pixels.
- 'Single pixel off (freehand)': clicking or dragging while holding the mouse button down will turn off pixels. This mode acts like a single pixel 'eraser'.
- 'Similar color': use it to quickly select an adjacent area that is similar in color.
- 'Similar brightness': use it to quickly select an adjacent area that is similar in brightness.

- 'Line toggle (click & drag)': use it to draw a line from the start point (when the mouse button was first pressed) to the end point (when the mouse button was released). This mode is particularly useful to trace and select satellite trails, for example for healing out using the 'Heal' module.
- 'Lasso': toggles all the pixels confined by a convex shape that you can draw in this mode (click and drag). Use it to quickly select or deselect circular areas by drawing their outline.
- 'Grow blob': grows any contiguous area of adjacent pixels by expanding their borders into the nearest neighboring pixel. Use it to quickly grow an area (for example a star core) without disturbing the rest of the mask.
- 'Shrink blob': shrinks any contiguous area of adjacent pixels by withdrawing their borders into the nearest neighboring pixel that is not part of a border. Use it to quickly shrink an area without disturbing the rest of the mask.

The Auto Feature

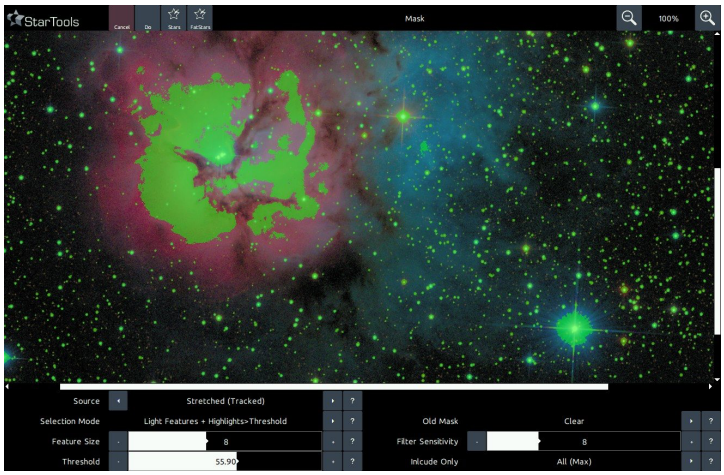

The Auto Mask Generator is indispensable when, for example, dealing with star mask, as required for many of the modules in StarTools.

The powerful 'Auto' function quickly and autonomously isolates features of interest such as stars, noise, hot or dead pixels, etc.

For example, isolating just the stars in an image is a necessity for obtaining any useful results from the 'Decon' and 'Shrink' module.

The type of features to be isolated are controlled by the 'Selection Mode' parameter

- 'Light features + highlight > threshold': a combination of two selection algorithms. One is the simpler 'Highlight > threshold' mode, which selects any pixel whose brightness is brighter than a certain percentage of the maximum value (see the 'Threshold' parameter below). The other selection algorithm is 'Light features' which selects high frequency components in an image (such as stars, gas knots and nebula edges), up to a certain size (see 'Max. feature size' below) and depending on a certain sensitivity (see 'Filter sensitivity' below). This mode is particularly effective for selecting stars. Note that if the 'Threshold' parameter is kept at 100%, this mode produces results that are identical to the 'Light features' mode.
- 'Light features': selects high frequency components in an image (such as stars, gas knots and nebula edges), up to a certain size (see 'Max feature size') and depending on a certain sensitivity (see 'Filter sensitivity').
- 'Highlight > threshold': selects any pixel whose brightness is brighter than a certain percentage of the maximum (e.g. pure white) value. If you find this mode does not select bright stars with white cores that well, open the 'Levels' module and set the

'Normalization' a few pixels higher. This should make light features marginally brighter and dark features marginally darker.

- 'Dead pixels color/mono < threshold': selects dark high frequency components in an image (such star edges, halos introduced by over sharpening, nebula edges and dead pixels), up to a certain size (see 'Max feature size' below) and depending on a certain sensitivity (see 'Filter sensitivity' below') and whose brightness is darker than a certain percentage of the maximum value (see the 'Threshold' parameter below). It then further narrows down the selection by looking at which pixels are likely the result of CCD defects (dead pixels). Two versions are available, one for color images, the other for mono images.
- 'Hot pixels color/mono > threshold': selects high frequency components in an image up to a certain size (see 'Max feature size' below) and depending on a certain sensitivity (see 'Filter sensitivity' below). It then further narrows down the selection by looking at which pixels are likely the result of CCD defects or cosmic rays (also known as 'hot' pixels). The 'Threshold' parameter controls how bright hot pixels need to be before they are potentially tagged as 'hot'. Note that a 'Threshold' of less than 100% needs to be specified for this mode to have any effect. Noise Fine - selects all pixels that are likely affected by significant amounts of noise. Please note that other parameters such as the 'Threshold', 'Max feature size', 'Filter sensitivity' and 'Exclude color' have no effect in this mode. Two versions are available, one for color images, the other for mono images.
- 'Noise': selects all pixels that are likely affected by significant amounts of noise. This algorithm is more aggressive in its noise detection and tagging than 'Noise Fine'. Please note that other parameters such as the 'Threshold', 'Max feature size', 'Filter sensitivity' and 'Exclude color' have no effect in this mode.
- 'Dust & scratches': selects small specks of dusts and scratches as found on old photographs. Only the 'Threshold' parameter is used, and a very low value for the 'Threshold' parameter is needed.
- 'Edges > Threshold': selects all pixels that are likely to belong to the edge of a feature. Use the 'Threshold' parameter to set sensitivity where lower values make the edge detector more sensitive.
- 'Horizontal artifacts': selects horizontal anomalies in the image. Use the 'Max feature size' and 'Filter sensitivity' to throttle the aggressiveness with which the detector detects the anomalies.
- 'Vertical artifacts': selects vertical anomalies in the image. Use the 'Max feature size' and 'Filter sensitivity' to throttle the aggressiveness with which the detector detects the anomalies.
- 'Radius': selects a circle, starting from the centre of the image going outwards. The 'Threshold' parameter defines the radius of the circle, where 100.00 covers the whole image.

Some of the selection algorithms are controlled by additional parameters:

- 'Include color': tells the selection algorithms evaluate specific color channels only when looking for features. This is particularly useful if you have a predominantly red, purple and blue nebula with white stars in the foreground and, say, you'd want to select only the stars. By setting 'Include only' to 'Green', you are able to tell the selection algorithms to leave red and blue features in the nebula alone (since these features are most prominent in the red and blue channels). This greatly reduces the amount of false positives.
- 'Max feature size': specifies the largest size of any feature the algorithm should expect. If you find that stars are not correctly detected and only their outlines show up, you may want to increase this value. Conversely, if you find that large features are being

inappropriately tagged and your stars are small (for example in wide field images), you may reduce this value to reduce false positives.

- 'Filter sensitivity': specifies how sensitive the selection algorithms should be to local brightness variations. A lower value signifies a more aggressive setting, leading to more features and pixels being tagged.
- 'Threshold': specifies a percentage of full brightness (i.e. pure white) below, or above which a selection algorithm should detect features.

Finally, the 'Source' parameter selects the source data the Auto mask generator should use. Thanks to StarTools' Tracking functionality which gives every module the capability to go 'back in time', the Auto mask generator can use either the original 'Linear' data (perfect for getting at the brightest star cores) or the data as you see it right now.

Advanced techniques

The mask editor and its auto-mask generator are very flexible tools. These more advanced techniques will allow you to create specialized masks for specific situations and purposes.

Object protection

Sometimes, it is desirable to keep an object or area from being included in an auto-generated mask. It is possible to have the auto-mask generator operate only on designated areas;

1. Clear the mask, and select the part of the image you wish to protect with the Flood Fill Lighter or Lasso tool, then click Invert.
2. In the Auto mask generator, set the parameters you need to generate your mask. Be sure to set 'Old Mask' to 'Add New Where Old Is Set'.
3. After clicking 'Do'. The auto-generator will generate the desired mask, however excluding the area we specified earlier.

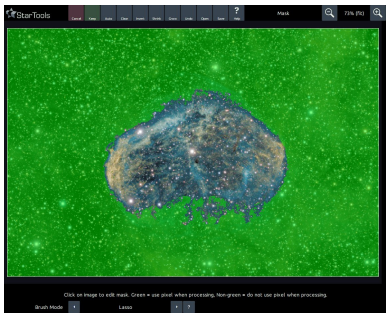
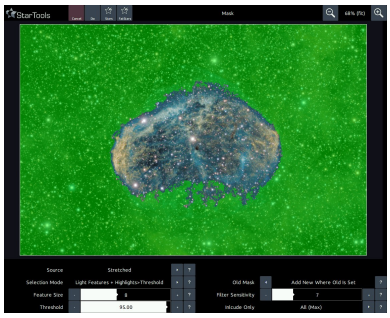
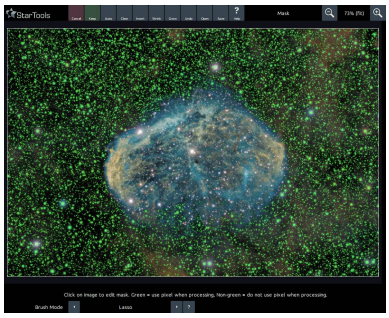

The ethics of using masks and selective processing

Where documentary photography is concerned, selective manipulation by hand is typically frowned upon, unless the practice of it is clearly stated when the final result is presented.

However, in cases where a mask is algorithmically derived, purely from the dataset itself, without adding any outside extra information, masking is common practice even in the realm of documentary photography. Examples of such use cases are range masks (for example, selecting highlights only based on brightness), star mask (selecting stars only based on stellar profile), color masks (selecting features based on color), etc.

In some modules in StarTools specifically, masks are used for the purpose of selective sampling to create internal parameters for an operation that is applied globally to all pixels. This too is common practice in the realm of documentary photography. Examples of such use cases are gradient modeling (selecting samples to model a global gradient on) and color balancing (selecting samples to base a global white balance on).

Finally, it is also generally permitted to mask out singularities (pixels with a value that is unknown) by hand, in order to exclude this from some operations that may otherwise

generate artifacts in response to encountering these. Examples may be over-exposing star cores, dead or hot pixels, stacking artifacts, or other data defects.

As a courtesy, when in doubt, it is always good to let you viewers know how you processed an image, in order to avoid confusion.

Using masks from startools.log

StarTools stores the masks you used in your workflow in the StarTools.log file itself. This StarTools.log file is located in the same folder as the executables. The masks are encoded as BASE64 PNG images. To convert the BASE64 text into loadable PNG images, you can use any online (or offline) BASE64 converter tool.

The part to copy and paste, typically starts with:

iVBOR.....

Online BASE64 converter by Motobit

One online tool for BASE64 is Motobit Software's BASE64 encoder/decoder.

To use it to convert StarTools masks back into importable PNG files;

Paste the BASE64 code into the text box

Select the 'decode the data from a Base64 string (base64 decoding)' radio button

Select the 'export to a binary file, filename:' radio button.

Name the file for example "mask.png"

Click the convert the source data button.

This should result in a download of the mask as a PNG file which can be imported into the StarTools mask editor, as well as other applications.

Wrap Up

Description

To help in the setting of a mask that can be used by many modules in StarTools to allow them to act only on the selected parts of an image.
Pixels set in the mask are shown in green. These pixels will be processed by the module.
Pixels not set in the mask are left in their original color. These pixels will not be processed.
If a selective mask is set the Mask will flash three times when you load a module.

When to use

Masks can be used in the following Modules:
'Color', 'Decon', 'Denoise', 'Filter', 'Flux', 'Heal', 'Layer', 'SuperStructure', 'Shrink', 'Repair', 'Sharp', 'Synth', 'Wipe'

What result to look for

- The mask set should ensure that only the selected pixels are modified.
- If parts of a DSO are wrongly included in a star mask try excluding the predominant colors of the DSO from the mask selection by using the 'Exclude Color' setting.

After Use

Continue with normal workflow. Keep or Clear the mask as required.

Description of Controls:

Open

The 'Open' button allows you to import a Mask that was previously saved by using the 'Save' button. Note that the image that is being opened to become the new Mask, needs to have the same dimensions as the image the Mask is intended for. Loading an image that has values between black and white will designate any shades of gray closest to white as 'on', and any shades of gray closest to black as 'off'

It works for any TIFF file so you can create masks using other programs like PhotoShop or GIMP

Save

The 'Save' button allows you to save the current mask to a standard TIFF file that shows 'on' pixels in pure white and 'off' pixels in pure black.

Auto

The 'Auto' button is a very powerful feature that allows you to automatically isolate features.

- See detailed description below

Clear

The 'Clear' button turns off all green pixels (i.e. it deselects all pixels in the image)

Invert

The 'Invert' button turns on all pixels that are off, and turns off all pixels that were on.

'Shrink'

The 'Shrink' button turns off all the green pixels that have a non-green neighbor, effectively 'shrinking' any selected regions.

Grow

The 'Grow' button turns on any non-green pixel that has a green neighbor, effectively 'growing' any selected regions.

Undo

The 'Undo' button allows you to undo the last operation that was performed.

Brush Mode

Specifies how clicking on the image affects the toggling on or off of pixels in the mask.
For a description of the different brush modes see Brush Modes.

Auto button description

Stars Preset

Creates a star mask where all the stars are selected (green).

FatStars Preset

Creates a star mask where a larger area around a star is selected (green) for all the stars.

- FatStars requires Tracking to be on as it requires a 'Linear' Source type. The preset is greyed out when Tracking is off.

AltStars Preset

uses a star detection algorithm that is much better at detecting stars with full stellar profile included.

- AltStars requires Tracking to be on as it requires a 'Linear' Source type. The preset is greyed out when Tracking is off.
- change Star Aberrations and Star Size as needed

Source

Sets which data the mask generator should derive the mask from.

- Linear - The mask generator should use the data when it is linear and has not been stretched.
  - This is particularly suited for detecting bright elements such as stars and over exposed elements.
  - This is perfect for creating star masks for deconvolution.
- Stretched - The mask generator should use the data that has been stretched to the current extent.
  - Use this if you need to select the detail as it is now.
- Default is Stretched [Stars], Linear [FatStars]
- The Linear setting is only available when Tracking is on.

Selection Mode

Specifies which type of features are to be isolated by the Auto algorithm.
For a description of the different Selection Modes.

Feature Size

Sets the maximum feature size Auto algorithm should attempt to isolate.

- The appropriate setting depends largely on the image size versus angular size of the object image.
- Increasing this setting may help in the correct detection of bigger stars.
- Used by Selection Modes: 'Light Features + Highlights', 'Light Features', 'Horizontal Artifacts', 'Vertical Artifacts' and 'Noise'.
- Default is 8. Range is 0 to 20.

Threshold

Sets a brightness threshold as a percentage of 'full brightness' to be used by the Selection Mode to decide which pixels to select.

- Used by 'Light Features + Highlights', 'Highlights', 'Dead Pixels', 'Hot Pixels', 'Dust & Scratches', 'Noise', 'Edges', 'Horizontal Artifacts', 'Vertical Artifacts' and 'Radius'
- Default is 100.00. Range is 0.00 to 100.00

Old Mask

Specifies what to do with any mask that already exists. New and old masks can be added in a number of ways - allowing complex masking:

- Clear - Clears the old mask - only uses the new mask.
- Add New To Old - Adds to the old mask any new pixels that are selected. The new mask is overlayed over the old mask.
- Subtract New From Old - Any pixels that are selected in the new mask will be unselected in the resulting mask. This is useful for deselecting stars in a specific part of an image.
- Add New Where Old Is Set - Only selects a pixel if both the old and the new mask have that pixel selected. You could use this to create a star mask in only a part of the image.
- Default is 'Clear'.

'Filter' Sensitivity

Sets how sensitive the feature detector should be when detecting features.

- Lower values increase the number of features detected (more sensitive), while bigger values reduce the number of features detected (less sensitive) - this may be counter-intuitive.
- For hot or dead pixel detection a higher sensitivity is appropriate.
- Default is 5 [Stars], 0 [FatStars]. Range is 0 to 30.

Exclude 'Color'

Sets which color channels to ignore when detecting features.

- This is particularly useful when detecting stars by minimizing false positives caused by, for example, bright gas knots in DSO's.
- Used by Selection Modes: 'Light Features + Highlights' and 'Light Features'
- Range is: None, Red, Blue, Yellow (Red + Green), Purple (Red + Blue)
- Default is None.

Associated Controls

Sometimes there are other controls in a module which are associated with using a mask.

Grow Mask

If a mask is used, sets the amount to temporarily grow the mask by.

- This allows a common mask to be used amongst several modules (e.g. a star mask) making small temporary adjustments to suit individual modules

Mask Fuzz

If a mask is used, Mask Fuzz controls the blending of the transition between masked and non-masked parts of the image.

- Increase this value to make a more gradual transition

Compose: Effortless, Signal Evolution-Tracked Composite Creation and Processing

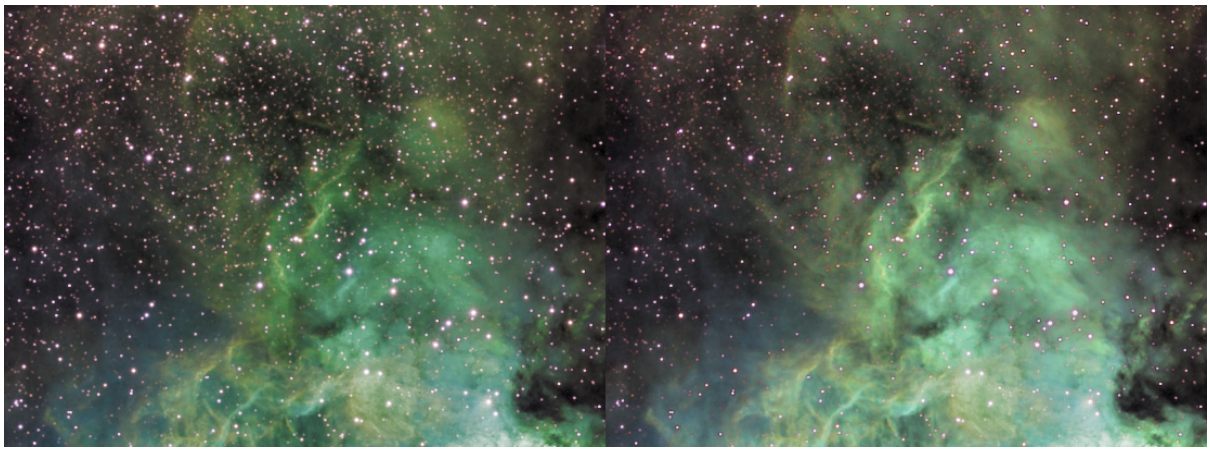

In conjunction with the Compose module, the 'Entropy' module can be used to boost detail in, for example, Synthetic L+SHO datasets, while availing of signal evolution Tracking. Here O-III was boosted, while Tracking kept noise propagation under control.

The Compose module is easy-to-use, but extremely flexible compositing and channel extraction tool. As opposed to all other software, the Compose module allows you to effortless process, LRGB, LLRGB, or narrowband composites like SHO, LSHO and more composites, as if they were simple RGB datasets.

In traditional image processing software, composites with separate luminance information (for example acquired through a luminance filter, created by a synthetic luminance frame, or a combination of both), require lengthy processing workflows; luminance (detail) and color information needs (or should!) be processed separately and only combined at the end to produce the final image.

Through the Compose module, StarTools is able to process luminance and color information separately, yet simultaneously.

This has important ramifications for your workflow and signal fidelity;

Your workflow for a complex composite is now virtually the same as it is for a simple DSLR/OSC dataset; Modules like 'Wipe' and 'Color' automatically consult and manipulate the correct dataset(s) and enable additional functionality where needed.

Because everything is now done in one Tracking session, you get all the benefits from signal evolution tracking until the very end, without having to end your workflow for luminance and start a new one for chroma/color; all modules cross-reference luminance and color information as needed until the very end, yielding vastly cleaner results.

The 'Entropy' module can consult the chroma/color information to effortlessly manipulate luminance as you see fit, while Tracking monitors noise propagation.

Synthetic luminance datasets will be created by simply specifying the total exposure times for each imported dataset. With a click of a button, synthetic luminance datasets can be added to an existing luminance dataset, or can be used as a (synthetic) luminance dataset in its own right.

Finally, the Compose module can be used to create bi-color composites, or to extract individual channels from color images.

Usage

Creating a composite is as easy as loading the desired datasets into the desired slots, and optionally setting the desired composite scheme and exposure lengths.

The 'Luminance' button loads a dataset into the 'Luminance File' slot. The 'Lum Total Exposure' slider determines the total exposure length in hours, minutes and seconds. This value is used to create the correct weighted synthetic luminance dataset, in case the 'Luminance, 'Color'' composite mode is set to create a synthetic luminance form the loaded channels. Loading a Luminance file will only have an effect when the 'Luminance, Color' parameter is set to a compositing scheme that incorporates a luminance dataset (e.g. 'L, RGB', 'L + Synthetic L From RGB, RGB' or 'L + Synthetic L From RGB, Mono') .

The Red, Green and Blue buttons load a dataset in the 'Red File', 'Green File' and 'Blue File' slots respectively. The 'Red Total Exposure', 'Green Total Exposure', 'Blue Total Exposure' sliders determine the total exposure length in hours, minutes and seconds for each of the three slots. These values are used to create the correct weighted synthetic luminance dataset (at 1/3 weighting of the 'Lum Total Exposure'), in case the 'Luminance, Color' composite mode is set to create a synthetic luminance from the loaded channels.

Loading an dataset into the 'Red File', 'Green File' or 'Blue File' slots will see any missing slots be synthesized automatically if the 'Color Ch. Interpolation' parameter is set to 'On'. Loading a color dataset into the 'Red File', 'Green File' or 'Blue File' slots will automatically extract the red, green and blue channels of the color dataset respectively.

There are a number of compositing schemes available, some of which will put StarTools into 'composite' mode (as signified by a lit up 'Compose' label on the Compose button on the home screen). Compositing schemes that require separate processing of luminance and color will put StarTools in this special mode. Some module may exhibit subtly different behavior, or expose different functionality while in this mode.

The following compositing schemes are selectable;

- 'RGB, RGB' simply uses red + green + blue for luminance and uses red, green and blue for the color information. No special processing or compositing is done. Any loaded Luminance dataset is ignored, as are total exposure settings.
- 'RGB, Mono' simply uses red + green + blue for luminance and uses the average of the red, green and blue channels for all channels for the color information, resulting in a mono image. Any loaded Luminance dataset is ignored, as are total exposure settings.
- 'L, RGB' simply uses the loaded luminance dataset for luminance and uses red, green and blue for the color information. Total exposure settings are ignored. StarTools will be put into 'composite' mode, processing luminance and color separately but simultaneously. If not Luminance dataset is loaded, this scheme functions the same as 'RGB, RGB' with the exception that StarTools will be put into 'composite' mode, processing luminance and color separately yet simultaneously.
- 'L + Synthetic L from RGB, RGB' creates a synthetic luminance dataset from Luminance, Red, Green and Blue, weighted according to the exposure times provided by the 'Total Exposure' sliders. The color information will consists of simply the red, green and blue datasets as imported. StarTools will be put into 'composite' mode, processing luminance and color separately yet simultaneously.
- 'L + Synthetic L from RGB, Mono' creates a synthetic luminance dataset from Luminance, Red, Green and Blue, weighted according to the exposure times provided by the 'Total Exposure' sliders. The color information will consists of the average of the red, green and blue channels for all channels, yielding a mono image. StarTools is not put into 'composite' mode, as no color information is available.

On synthetic luminance generation

For practical purpose, synthetic luminance generation assumes that, besides possibly varying total exposure lengths, all other factors remain equal. E.g. it is assumed that bandwidth response is exactly equal to that of the other filters in terms of width and transmission, and that only shot noise from the object varies (either due to differences in signal in the different filter band from the imaged object, or due to differing exposure times).

When added to a real (non synthetic) luminance filter source, the synthetic luminance's three red, green and blue channels are assumed to contribute exactly one-third to the added synthetic luminance. E.g. it is assumed that the aggregate filter response of the individual three red, green and blue channels, exactly match that of the single luminance channel.

Channel assignment and coloring and narrowband datasets

Unique to StarTools, channel assignment does not dictate final coloring. In other words, loading, for example, a SHO dataset as RGB, does not lock you into using precisely that channel mapping. Thanks to the signal evolution Tracking engine, the Color module allows you to completely remap the channels at will for the purpose of colouring, even far into your processing.

As is common practice in astronomy, StarTools assumes channels are imported in order of descending wavelength. E.g. the dataset with the longest wavelength (e.g. the light with the highest nm or Å comes first). In other words, the reddest light comes first, and the bluest light comes last.

In practice this means that:

When using visual spectrum datasets, load red into the red channel, green into the green channel, and blue into the blue channel.

When using triple channel narrowband datasets such as Hubble-like S-II + H-alpha + O-III (aka "SHO" datasets), load S-II as red, H-alpha as green and O-III as blue.

When using a duo/tri/quad band filtered dataset, load H-alpha (which is possibly combined with the neighbouring S-II line depending on the filter) as red, and load O-III (which is possibly combined with the neighbouring H-beta line depending on the filter) as green.

In any case, you should not concern yourself with the colouring until you hit the Color module in your workflow; as opposed to other software, this initial channel assignment has no bearing at all on the final colouring in your image. Please note that failing to import channels correctly in the manner and order described above, will cause the Color module to mis-label the many colouring and blend options it offers.

Narrowband accents

With the introduction of the NBAccent module in StarTools 1.8, a third parallel datastream type has been introduced; that of narrowband accents for visual spectrum augmentation. Adding narrowband accents to visual spectrum datasets has traditionally been a daunting, difficult and laborious process, involving multiple workflows. The NBAccent module is a powerful module that starts its work as soon as your load your data in the Compose module. Crucially it adds only a single, easy step to an otherwise standard workflow, while yielding superior results in terms of color fidelity/preservation.

Dedicated functionality for visual spectrum narrowband accents are part of StarTools' integrated workflow.

By making narrowband accents an integral part of the complete workflow and signal path, results are replicable, predictable and fully tracked by StarTools' unique signal evolution Tracking engine, yielding perfect noise reduction every time.

Enabling narrowband accents in your wofklow, is as easy as loading the file containing the signal you wish to add as narrowband accents, and specifying the type of accents the file contains. Three possible types are selectable;

H-alpha or S-II from a narrowband filter

O-III or H-beta from a narrowband filter

A combination of narrowband signals across multiple channels from a duo, tri or quadband filter (such as the Optolong L-Extreme or L-eNhance) or a combined single narrowband filter

Be sure to specify the correct type before continuing.

Popular coloring

Popular narrowband composite coloring

Hubble / HST / SHO

The Hubble Space Telescope palette (also known as 'HST' or 'SHO' palette) is a popular palette for color renditions of the S-II, Hydrogen-alpha and O-III emission bands. This palette is achieved by loading S-II, Hydrogen-alpha and O-III ('SHO') as red, green and blue respectively. A special 'Hubble' preset in the 'Color' module provides a shortcut to color rendition settings that mimic the results from the more limited image processing tools from the 1990s.

Top: SHO (HST) palette rendering showing 3 emission concentrations (S-II as red, H-alpha as green, O-III as blue). Bottom: HOO bi-color rendering showing 2 emission concentrations (Ha-alpha as red, O-III as cyan).

H-alpha and O-III bicolor

A popular bi-color rendition of H-alpha and O-III is to import H-alpha as red and O-III as green as well as blue. A synthetic luminance frame is then created that only gives red and blue (or green instead of blue, but not both!) a weighting according to the two datasets' exposure lengths. The resulting color rendition tends to be close to these bands' manifestation in the visual spectrum with H-alpha a deep red and O-III appearing as a teal green.

Wrap Up

Description

- To allow the loading of separate files representing Luminance and Red, Green and Blue color channels to enable complex composites. The Compose module is used to load and combine multiple files into a single image. Either standard luminance, Red, Green and Blue data or a combination of narrow-band data using different filters.
- This allows simple creations of complex composites such as LLRGB
- It can also be used to extract the Red, Green or Blue channel from a color image.
- If a channel is missing it can interpolate from the other channels to create the missing channel.
- The balancing (weighting) of the different channels is done automatically based on the total exposure time for each channel.
- While Tracking is on the 'Compose' button is lit up to show processing is being done in a special 'Compose' mode.
- In the Compose Mode luminance and color datasets are processed separately in parallel in subsequent modules.
- 'Bin', 'Wipe' etc until Tracking is turned off.

When to use

At the start of processing if you need to load files representing different channels - either traditional LRGB or other combinations such as Hubble Palette (SII, Ha, OIII)

What result to look for

- Make sure all the channels are properly aligned. In rare cases it is possible for the registration to be out.

Description of Controls

Luminance button

Allows you to navigate to and load the Luminance file.

Red Button

Allows you to navigate to and load the Red channel file.

- If a luminance file has been loaded then any Red, Green or Blue file subsequently loaded can be either the same size or exactly a quarter of the resolution of the luminance file. This allows luminance channels to be 1x1 binned while R,G and B files can be 2x2 binned. In all other cases all loaded files must be the same dimensions.
- If loading a color (RGB) image - extracts the Red channel.

Green Button

Allows you to navigate to and load the Green channel file

- If loading a color (RGB) image - extracts the Green channel.

Blue Button

Allows you to navigate to and load the Blue channel file

- If loading a color (RGB) image - extracts the Blue channel.

Luminance, 'Color'

Defines how to combine the data from each channel.

- RGB, RGB - Uses the luminance from the R,G and B channels, and the color from the R,G and B channels.
- RGB, Mono - - Uses the luminance from the R,G and B channels, does not output color.
- L, RGB - Uses the luminance from the L channel, and the color from the R,G and B channels.
- L + Synthetic L from RGB, RGB - Creates synthetic luminance data from the R,G and B channels and adds it to the luminance channel, and the color from the R,G and B channels. Also known as LLRGB.
- L + Synthetic L from RGB, Mono - Creates synthetic luminance data from the R,G and B channels and adds it to the luminance channel, does not output color.
- L + Synthetic L from RGB, R(GB)(GB) (BiColor) - Creates synthetic luminance data from the R,G and B channels combined, and the colour from the R and either G or B channels. Useful for creating Bi-colour images from 2 narrowband filtered datasets (v1.7.420+).
- L + Synthetic L from R(2xG)B, R(GB)(GB) (BiColor from OSC/DSLR) - Creates synthetic luminance data from the R,2xG and B channels combined, and the colour from the R and G+B and G+B channels (Standard HOO palette). Palette can be changed in Color Module later. Useful for creating Bi-colour images from 2 duo/tri/quad-band filtered datasets. (v1.7.420+).

Channel Interpolation

The Compose module can be set to interpolate any missing channels.

- Just set Channel Interpolation On and load what channels you have. This feature can be used to:
  - Generate a missing green channel in the case of an Ha/Hb composite
  - Generate a greyscale from a single Ha, Hb, OIII or SII frame which may later be turned into a false color image using the 'Color' module.
- Default is On

Luminance File

Shows the path of any file loaded in the Luminance channel.

- Default is None.

Red File

Shows the path of any file loaded in the Red channel.

- Default is None.

Green File

Shows the path of any file loaded in the Green channel.

- Default is None.

Blue File

Shows the path of any file loaded in the Blue channel.

- Default is None.

Lum Total Exposure

Sets the total exposure time of the luminance data.

- Default is Not Set.
- Range is 1 minute to 720 minutes.

Red Total Exposure

Sets the total exposure time of the red data.

- Default is Not Set.
- Range is 1 minute to 720 minutes.

Green Total Exposure

Sets the total exposure time of the green data.

- Default is Not Set.
- Range is 1 minute to 720 minutes.

Blue Total Exposure

Sets the total exposure time of the blue data.

- Default is Not Set.
- Range is 1 minute to 720 minutes.

'AutoDev'

Top: traditional Digital Development curve (via 'FilmDev' module)

Bottom: 'AutoDev'. Notice the vastly better dynamic range allocation, with more detail visible in the shadows and highlights, while not compromising on detail in midtones or blowing out stars. The 'AutoDev' image is the perfect starting point for enhancing local detail.

'AutoDev' is an advanced image stretching solution that relies on detail analysis, rather than the simple non-linear transformation functions from yesteryear.

To be exact, in StarTools, Histogram Transformation Curves (DDP, Levels and Curves, ArcSinH stretch, MaskedStretch etc.) are considered obsolete an non-optimal; 'AutoDev' uses robust, controllable image analysis to achieve better, more objective results in a more intuitive way.

When data is acquired, it is recorded in a linear form, corresponding to raw photon counts. To make this data suitable for human consumption, stretching it non-linearly is required.

Historically, simple algorithms were used to emulate the non-linear response of photographic paper by modeling its non-linear transformation curve. Later, in the 1990s because dynamic range in outer space varies greatly, 'levels and curves' tools allowed imagers to create custom histogram transformation curves that better matched the object imaged so that the most amount of detail became visible in the stretched image.

Creating these custom curves was a highly laborious and subjective process. And, unfortunately, in many software packages this is still the situation today. The result is almost always sub-optimal dynamic range allocation, leading to detail loss in the shadows (leaving recoverable detail unstretched), shrouding interesting detail in the midtones (by not allocating it enough dynamic range) or blowing out stars (by failing to leave enough dynamic range for the stellar profiles). Working on badly calibrated screens, can exacerbate the problem of subjectively allocating dynamic range with more primitive tools.

'StarTools' 'AutoDev' module however uses image analysis to find the optimum custom curve for the characteristics of the data.'

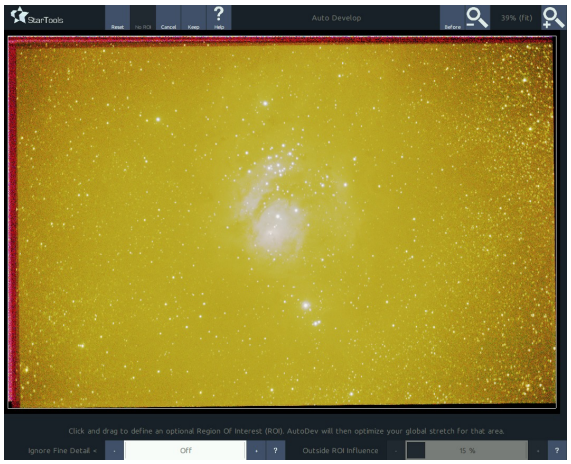

Not a bug, but a feature! Don't let a first result like this scare you. 'AutoDev' is doing you a favor by showing you exactly what is wrong with your data. In this we can see heavy light pollution, gradients and stacking artifacts that need taking care of before we can go any further

StarTools' 'AutoDev' module uses image analysis to find the optimum custom curve for the characteristics of the data. By actively looking for detail in the image, 'AutoDev' autonomously creates a custom histogram curve that best allocates the available dynamic range to the scene, taking into account all aspects and detail. As a consequence, the need for local 'HDR' manipulation is minimized.

'AutoDev' is, in fact, so good at its job that it is also one of the most important tools in StarTools for initial data inspection; using 'AutoDev' as one of the first modules on your data will see it bring out problems in the data, such as stacking artifacts, gradients, bias, dust donuts, etc. Upon removal and/or mitigation of these problems, 'AutoDev' may then be used to stretch the cleaned up data.

Upon removal and/or mitigation of these problems, 'AutoDev' may then be used to stretch the cleaned up data, bringing out detail across the entire dynamic range equally.

Usage

To be able to detect detail, 'AutoDev' has a lot of smarts behind it. Its main detail detection algorithm analyses a Region of Interest ("RoI") - by default the whole image - so that it can find the optimum histogram transformation curve based on what it "sees".

Understanding 'AutoDev' on a basic level is pretty simple really; its goal is to look at what's in your image and to make sure as much as possible is visible, just as a human would (try to) look at what is in the image and approximate the optimal histogram transformation curves using traditional tools.

The problem with a histogram transformation curve (aka 'global stretch') is that it affects all pixels in the image. So, what works in one area (bringing out detail in the background), may not necessarily work in another (for example, it may make a medium-brightness DSO core harder to see). Therefore it is important to understand that - fundamentally - globally stretching the image is always a compromise. 'AutoDev"s job then, is to find the best-compromise global curve, given what detail is visible in your image and your preferences. Of course, fortunately we have other tools like the 'Contrast', 'Sharp' and 'HDR' modules to 'rescue' all detail by optimizing for local dynamic range on top of global dynamic range.

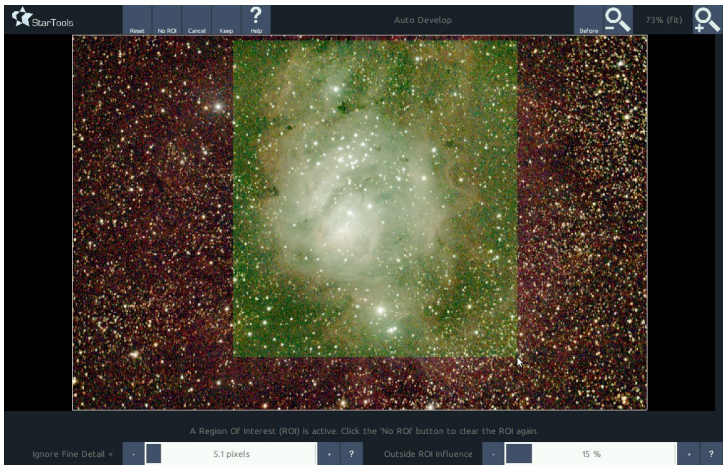

Great allocation of dynamic range by 'AutoDev' after taking care of the stacking artifacts, gradients and light pollution using the 'Wipe' module.

Being able to show all things in your image equally well, is a really useful feature, as it is also very adept at finding artifacts or stuff in your image that is not real celestial detail but requires attention. That is why 'AutoDev' is also extremely useful to launch as the first thing after loading an image to see what - if any - issues need addressing before proceeding. If there are any, 'AutoDev' is virtually guaranteed to show them to you. After fixing such issues (for example using 'Crop', 'Wipe', Band or other modules), we can go on to use 'AutoDev"s skills for showing the remaining (this time real celestial) detail in the image.

'AutoDev finds the best compromise global curve, given what detail is visible in your image and your preferences.'

If most of the image consists of a background and just a small object of interest, by default 'AutoDev' will weigh the importance of the background higher (since it covers a much larger part of the image vs the object). This is understandable and neatly demonstrates its behavior. It will always look for the best compromise stretch to show the entire Region of Interest ("RoI" - by default the entire image). This also means that if the background is noisy, it will start digging out the noise, taking it as "fine detail" that needs to be "brought out". If this behavior is undesirable, there are a couple of things you can do in 'AutoDev'.

1. Change the 'Ignore Fine Detail <' parameter, so that 'AutoDev' will no longer detect fine detail (such as noise grain).
2. Simply tell it what it should focus on instead by specifying an ROI and not regard the area outside the ROI just a little bit ('Outside ROI influence').

You will find that, as you include more background around the object, 'AutoDev', as expected, starts to optimise more and more for the background and less for the object. To use the RoI effectively, give it a "sample" of the important bit of the image. This can be a whole object, or it can be just a slice of the object that is a good representation of what's going on in the object in terms of detail. You can, for example, use a slice of a galaxy from the core, through the dust lanes, to the faint outer arms. There is no shame in trying a few different ROIs in order to find one you're happy with. What ever the case, the result will be more optimal and objective than pulling at histogram curves.

There are two ways of further influencing the way the detail detector "sees" your image;

- The 'Detector Gamma' parameter applies - for values other than 1.0 - a non-linear stretch to the image prior to passing it to the detector. E.g. the detector will "see" a darker or brighter image and create a curve that suits this image, rather than the real image.

- This makes the detector proportionally more (< 1.0) or less (> 1.0) sensitive to detail in the highlights.
- Conversely it makes the detector less (<1.0) or more (> 1.0) sensitive to detail in the shadows. The effect can be though of as a "smart" gamma correction. Note that tweaking this parameter will, by virtue of its skewing effect, cause the resulting stretch to no longer be optimal.

- The 'Shadow Linearity' parameter specifies the amount of linearity that is applied in the shadows, before non-linear stretching takes over. Higher amounts have the effect of allocating more dynamic range to the shadows and background.

Understanding 'AutoDev''s behavior

In 'AutoDev', you're controlling an impartial and objective detail detector, rather than a subjective and hard to control (especially in the highlights) bezier/spline curve.

Having something impartial and objective taking care of your initial stretch is very valuable, as it allows you to much better set up a "neutral" image that you can build on with the other local detail-enhancing tools in your arsenal (e.g. 'Sharp', 'HDR', 'Contrast', 'Decon', etc.). For example, when using Autodev, it will quickly become clear that point lights and over-exposed highlights, such as the cores of bright stars, remain much more defined. The dreaded "star bloat" effect is much less pronounced or even entirely absent, depending on the dataset.

Notice how over-exposed highlights do not "bloat" at all. The cores stay in their place and do not "bleed" into the neighboring pixels.

However, knowing how to effectively use Region of Interests ("RoI") is crucial to making the most of 'AutoDev'. Particularly if the object of interest is not image-filling, a Region of Interest will often be necessary. Fortunately the fundamental workings of the RoI are easy to understand.

Detail inside the ROI

Let's say our image is of galaxy, neatly situated in the center. Then confining the RoI progressively to the core of the galaxy, the stretch becomes more and more optimized for the core and less and less for the outer rim. Conversely, if we want to show more of the outer regions as well, we would include those regions in the RoI.

Detail outside the ROI

Shrinking or enlarging the RoI, you will notice how the stretch is optimized specifically to show as much as possible of the image inside of the RoI. That is not to say that anything outside the RoI will be invisible. It just means that any detail there will not (or much less) have a say in how the stretch is made. For example, if we would have an image of a galaxy, cloned it, put the two image side by side to create a new image, and then specified the RoI perfectly over just one of the cloned galaxies, the other one, outside the RoI would be stretched precisely the same way (as it happens to have exactly the same

detail). Whatever detail lies outside the RoI, is simply forced to conform to the stretch that was designed for the RoI.

It is important to note that 'AutoDev' will never clip your blackpoints outside the RoI, unless the 'Outside RoI Influence' parameter is explicitly set to 0% (though it is still not guaranteed to clip even at that setting). Detail outside the RoI may appear very dark (and approach 0/black), but will never be clipped.

Bringing up the 'Outside RoI Influence' parameter will let 'AutoDev' allocate the specified amount of dynamic range to the area outside the RoI as well, at the expense of some dynamic range inside the RoI. If 'Outside RoI Influence' set 100%, then precisely 50% of the dynamic range will be used to show detail inside the RoI and 50% of the dynamic range will be used to show detail outside the RoI. Note that, visually, this behavior is area-size dependent; if the RoI is only a tiny area, the area outside the RoI will have to make do with just 50% of the dynamic range to describe detail for a much larger area (e.g. it has to divide the dynamic range over many more pixels), while the smaller RoI area has much fewer pixels and can therefore allocate each pixel more dynamic range if needed, in turn showing much more detail.

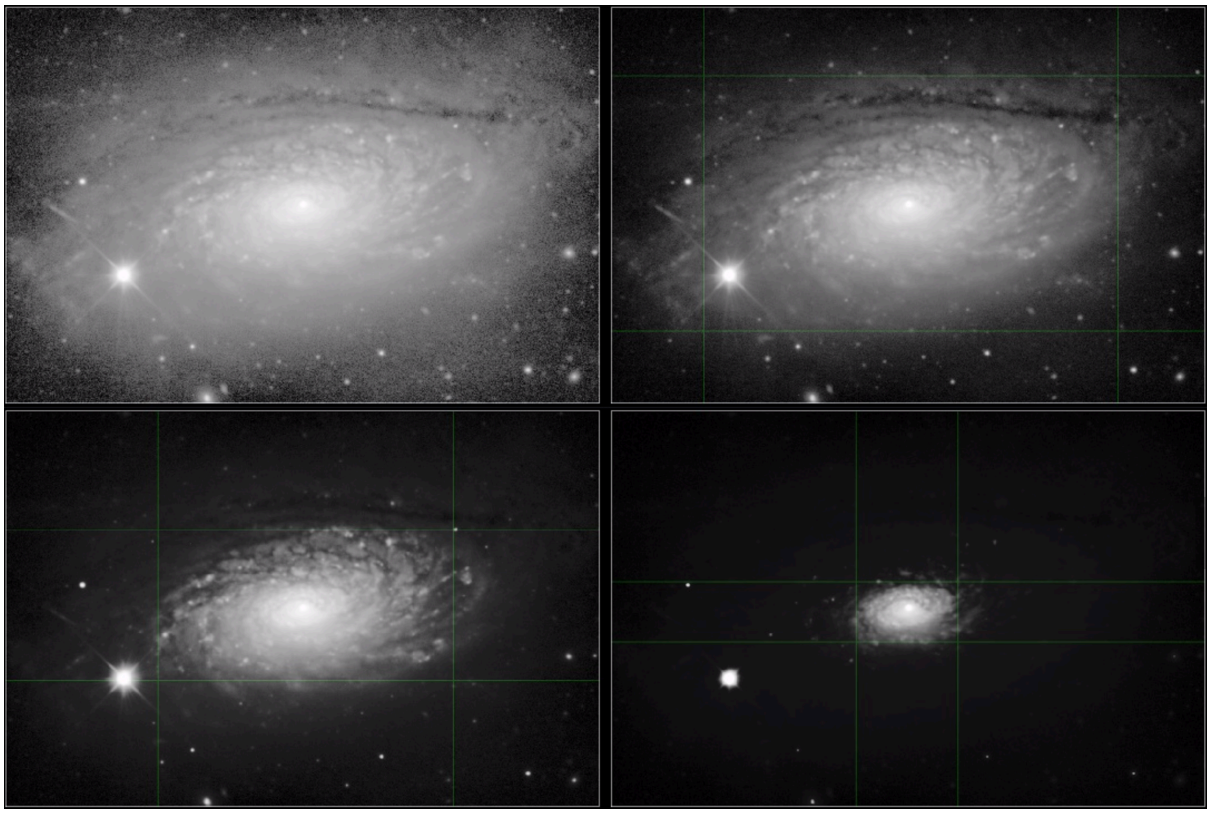

Confining the Region of Interest ("RoI") progressively to the core of this galaxy, the stretch becomes more and more optimized for the core and less and less for the outer regions.

'Color' retention

Please note you should completely disregard the coloring in 'AutoDev' (if coloring is even at all visible).

Non-linearly stretching an image's RGB components causes its hue and saturation to be similarly stretched and squashed. This is often observable as 'washing out' of coloring in the highlights.

Traditionally, image processing software for astrophotography has struggled with this, resorting to kludges like 'special' stretching functions (e.g. arcsinh) that somewhat minimize the problem, or even procedures that make desaturated highlights adopt the colors of neighboring, non-desaturated pixels.

'While other software continues to struggle with color retention, StarTools Tracking feature allows the 'Color' module to go back in time and completely reconstruct the RGB ratios as recorded, regardless of how the image was stretched'

While other software continues to struggle with color retention, StarTools Tracking feature allows the 'Color' module to go back in time and completely reconstruct the RGB ratios as recorded, regardless of how the image was stretched.

This is one of the major reasons why running the 'Color' module is preferably run as one of the last steps in your processing flow; it is able to completely negate the effect of any stretching - whether global or local - may have had on the hue and saturation of the image.

Because of this, 'AutoDev"s performance is not stymied like some other stretching solutions (e.g. arcsinh) by a need to preserve coloring. The two aspects - color and luminance - of your image are neatly separated thanks to StarTools' signal evolution Tracking engine.

Wrap Up

Description

- 'AutoDev' is designed to show up the faults in an image - that is why most images look bad when doing the initial global stretch.
- When given a suitably cropped and 'Wiped image 'AutoDev' is designed to find the optimal global stretch.
- 'AutoDev' employs an enhanced algorithm that yields optimal global dynamic range assignment - without masking artifacts - while being easy to control.

When to use

- For the initial global stretch - Just after loading the image in StarTools - to allow initial viewing of the features and problems associated with the image.
- For the final global stretch - After using the 'Wipe' module.

What result to look for (Initial global stretch)

'AutoDev' will highlight any of the following:

- Stacking Artifacts - remove later with the 'Crop' module.
- Color bias (remove this later with the 'Wipe' module)
    - Red or yellow/brown cast - skyglow that has been white balanced.
    - Teal, blue or green cast - skyglow that has not been white balanced.
    - Bright blue-green cast - skyglow filtered using a light pollution filter.
    - Missing yellow (no yellow stars) - indicates light pollution filter used
- Vignetting - darkening towards the corners, Amp Glow - remove later with the 'Wipe' module.
- Dust specks - remember to mask out when using the 'Wipe' module.
- Noise.
- Banding - remove with the 'Wipe' module later
- Debayering Problems - checkerboard pattern.
- Coma - fix later with the 'Lens' module.

What result to look for (Final global stretch)

Look for:

- Good dynamic range in area of interest, showing all major structures of interest. Don't worry about detail within the structures - they will be targeted later.

- Controlled noise in the background - if not go back and do a less aggressive stretch (by changing ROI selection). The 'Denoise' module can handle a reasonable amount of background noise but has problems when it is excessive.
- Color Bias, Vignetting, Dust specks, etc. should be well controlled or removed by this stage. If not you may need to go back and redo the 'Wipe' or 'Crop'.

Ways of getting better results

- Use Flats and Darks (or Dither) when creating the original image
- Use the 'Bin' module at the start of processing if the original image is oversampled.
- If you continue to have problems with background noise you may want to try the 'FilmDev' module instead. It may provide the level of control you need.

Description of Controls

Presets

- Reset - Resets controls 'Ignore Fine Detail' to 'Off' and 'Outside ROI Influence' to 0%.
- No ROI - Clears any Region of Interest (ROI) set.

Ignore Fine Detail <

Allows 'AutoDev' to ignore small features such as noise when allocating dynamic range.

- Particularly important when the image has a low signal to noise ratio (SNR).
- Default is Off.
- Range is Off then 1.1 to 50.0 pixels.
- Increase to exclude noise - usually until the image doesn't darken any more.

Detetor Gamma

Allows you to apply a Gamma correction before the linear data is analysed by the detail detection algorithms

- Increasing the value above 1.00 will make bright details stand out more - so more dynamic range will be allocated to bright detail.
- Decreasing the value below 1.00 will make bright details stand out less.
- The effect is like having an intelligent Gamma correction.
- Keep at 1.00 for no change.
- Default is 1.00.
- Range is 0.00 to 10.00

Outside ROI Influence

Defines how much dynamic range to reserve for outside the Region of Interest (ROI).

- Defaults to 15%. Range is 0% to 100%.
- Reduce if you have a high dynamic range subject and you want to allocate more dynamic range to the subject.
- Avoid 'AutoDev' clipping the data by never reducing this value to 0%.
- Increase if there is some detail outside the ROI which you want to show up.

Shadow Linearity

Sets how much linearity is applied to the shadows - before non-linear stretching takes over.

- Values above 50% allocate more dynamic range to the shadows and background.
- Values below 50% allocate more dynamic range to the highlights and foreground.
- Default is 50%.
- Range is from 0% to 100%.

'FilmDev': Stretching with Photographic Film Emulation

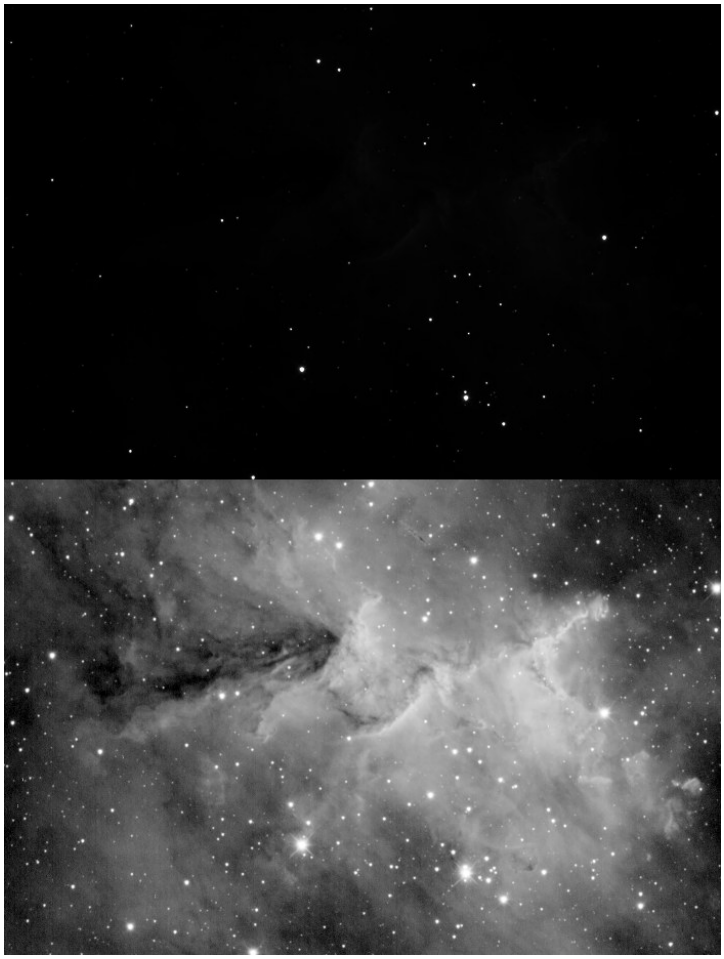

Top: linear image

Bottom: image developed by photographic film curve using Develop module by 'homing in'.

Notice the lack of star bloat, courtesy of the automatic black and white point detection.

The 'FilmDev' module was created from the ground up as an alternative the classic Digital Development algorithm that attempts to emulate classic film response when first developing a raw stacked image. It effectively functions as a classic digital 'dark room' where your prized raw signal is developed and readied for further processing.

The module can also be used as Swiss pocket knife for gamma correction, normalization and channel luminance contribution remixing.

Usage

First off, please note that this module emulates many aspects of photographic film, including its shortcomings. These shortcomings include photographic film's tendency to "bloat" stellar profiles. If your goal is to achieve a non-linear stretch that shows as much detail as possible, the far more advanced 'AutoDev' will always do an objectively better job for that purpose. Please note that the edge-enhancing qualities of photographic film are not emulated by this module, as this step is best done through other means.

'This module emulates many aspects of photographic film, including its shortcomings'

Enhancements over the classic Digital Development algorithm (Okano, 1997), are the introduction of an additional gamma correction component, the removal of the edge enhancement component, and the introduction of automated black and white point detection. The latter ensures your signal never clips, while making histogram checking a thing of the past.

'Central to the module, is the 'Digital Development' parameter, which controls the strength of the development and resulting stretch'

Central to the module, is the '**Digital Development**' parameter, which controls the strength of the development and resulting stretch. A semi-automated 'homing in feature' attempts to find the optimal settings that bring out as much detail as possible, while still adhering to the Digital Development curve. This feature can be accessed by clicking on the 'Home In' button until the image does not change much further. A simple '**Gamma**' correction can also be applied.

A 'Dark Anomaly Filter' helps the automatic black point detector ignore any dead pixels. Any dead or darker-than-real-background pixels are caught by the filter, they are re-allocated a reduced amount of dynamic range as set by the 'Dark Anomaly Headroom' parameter.

Automatic white point detection ('White Calibration') uses any over-exposing stars or other highlights in your image, however it can also be switched to use the 'Dark Anomaly 'Filter'' setting to filter out any bright anomalies (e.g. hot pixels) that are not stars or real highlights.

An artificial pedestal value can be introduced through the 'Skyglow' parameter. This parameter specifies how much of the dynamic range (up to 50%) should be taken up by the artificial pedestal.

Finally, a luminance mixer allows for re-mixing of the contribution of each color channel to brightness.

'Color' Retention

Non-linearly stretching an image's RGB components causes its hue and saturation to be similarly stretched and squashed. This is often observable as "washing out" of coloring in the highlights.

Traditionally, image processing software for astrophotography has struggled with this, resorting to kludges like "special" stretching functions (e.g. ArcSinH) or 'Color' enhancement extensions to the DDP algorithm (Okano, 1997) that only attempt to minimize the problem, while still introducing color shifts.

While other software continues to struggle with color retention, StarTools Tracking feature allows the 'Color' module to go back in time and completely reconstruct the RGB ratios as recorded, regardless of how the image was stretched.

This is one of the major reasons why running the 'Color' module is preferably run as one of the last steps in your processing flow; it is able to completely negate the effect of any stretching - whether global or local - may have had on the hue and saturation of the image.

Because of this, the digital development color treatment extensions as proposed by Okano (1997) has not been incorporated in the Develop module. The two aspects - color and luminance - of your image are neatly separated thanks to StarTools' signal evolution Tracking engine.

Wrap Up

Description

'FilmDev' allows a more manually controlled global assisted stretching than 'AutoDev'.

'FilmDev' automatically detects the white and black point and ensures the signal never clips.

When to use

- In place of 'AutoDev' where you have difficulty getting the results you want with 'AutoDev'. Principal causes are:

    - There is no region of interest with sufficient detail for 'AutoDev' to use
    - There is a lot of noise which disrupts 'AutoDev' operation
- For the initial global stretch - after loading the image
- For the final global stretch - after using the 'Wipe' module
- For re-adding an artificial skyglow - at the end of processing

What result to look for

Initial global stretch - 'AutoDev' will highlight any of the following:

- Stacking Artifacts - remove later with 'Crop' module
- Color bias (remove this later with the 'Wipe' module)
    - Red or yellow/brown cast - skyglow that has been white balanced
    - Teal, blue or green cast - skyglow that has not been white balanced
    - Bright blue-green cast - skyglow filtered using a light pollution filter
    - Missing yellow (e.g. no yellow stars) - indicates use of light pollution filter
- Vignetting, Amp Glow - remove these later with the 'Wipe' module
- Dust specks - remember to mask these out when using the 'Wipe' module
- Noise
- Banding - use the Band module next - but don't 'Keep' the Develop module result.
- Debayering Problems - checkerboard pattern. See description here
- Coma - fix later with the 'Lens' module

Final global stretch - look for:

- Good dynamic range in the area of interest
- Make sure you bring out the faintest detail you want to see - the brighter details will be brought out by other modules
- Look for controlled noise in the background - if not, go back and do a less aggressive stretch (by reducing the Digital Development setting)
- Color Bias, Vignetting, Dust specks, etc. should be well controlled or removed by this stage.

Ways of getting better results

Use Flats and Darks when creating the original image

After Use

- After the initial global stretch - use the Band module if needed - then use the 'Lens', 'Bin', 'Crop' and 'Wipe' modules as needed.
- After the final global stretch - Consider using the 'Decon', 'Sharp', 'Contrast', 'HDR' modules as needed.

Description of Controls

Digital Development:

Emulates the response of film

- Defaults to Off. Range is Off(0.00%) to 100.00%.
- Use the 'Home In' button or change the slider to adjust the black point of the global stretch of the image.

Home In:

Intelligently homes in on the optimum setting for Digital Development. Repeat until you get the result you want.

Gamma:

Use this control to do a simple Gamma non-linear stretch. Used as an alternative to Digital Development control.

- Defaults to 1.00.
- Range is 0.00 to 10.00.

- Leave at 1.00 if using Digital Development.
- If not using Digital Development - set to 2.2 for correct conversion of linear data so that it appears linear on-screen. Set to approx. 0.46 to undo a JPEGs gamma curve.

Dark Anomaly 'Filter':

Controls the filtering of Dark Anomalies so that they don't interfere with the Develop operation. Described further here.

Dark Anomalies are anything that is darker than the galactic background. Examples are: Dead pixels, stacking artifacts. Dust specks, scratches, trees, mountains.

- The Dark Anomaly 'Filter' can remove small dark anomalies. For larger dark anomalies you will need to use the 'Heal' or 'Crop' modules.
- Defaults to 'Off'.
- Range is Off to 20.0.
- Increase the value until the image stops darkening. Maximum 20.0 pixels. This ensures that the Dark Anomalies are represented below the darkest point of real data.

Dark Anomaly Headroom:

Defines the percentage of the original value a dark anomaly now occupies. This frees up more dynamic range for subject data.

- Defaults to 5%.
- Range is 0% to 100%.
- Keep above 0% - to avoid clipping.

White Calibration:

Specifies how the Develop module decides on the white point of the global stretch of the image.

- 'Use Stars' - look for big bright stars
- 'Use Dark Anomaly 'Filter' (May White Clip)' – (use if there are no bright stars and the brightest features are hot pixels). Uses the Dark Anomaly 'Filter' to filter out these white anomalies as well (requires the Dark Anomaly 'Filter' to be active).
- Defaults to 'Use Stars'

Color Channel Luminance Contribution:

- For color images Develop uses the RGB channels to create luminance data on which to do the stretch.
- If one of the channels is unreliable (due to Chromatic Aberration or lack of Infra Red) then its influence should be reduced.
- These controls have no effect on monochrome images and so in this case are disabled.
- Defaults to 100%.
- Range is 0% to 500%.

Luminance Contribution Settings:

Red Luminance Contrib. - Sets the relative influence of the Red channel.

Green Luminance Contrib. - Sets the relative influence of the Green channel.

Blue Luminance Contrib. - Sets the relative influence of the Blue channel.

Skyglow:

- Re-introduces an artificial sky glow.
- Best to leave this to just before saving the image.
- Defaults to 0%.
- Range is 0% to 50%.

'Bin': Trade Resolution for Noise Reduction

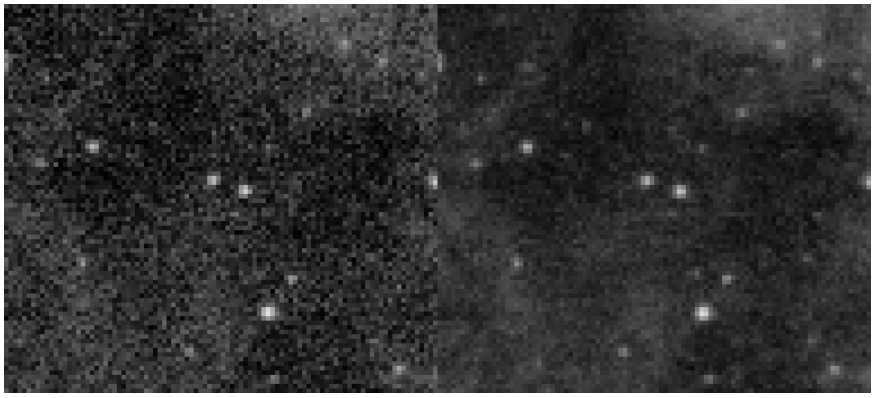

400% zoomed crop of an image. Left: scaled down to 25% of its original size using nearest neighbor sampling (retaining noise). Right: same image binned down to 25% of its original size. A significant amount of noise reduction has occurred. Further deconvolution is now an option. Notice real structural detail is not compromised, but any non-structural detail (noise) has been removed.

The 'Bin' module puts you in control over the trade-off between resolution, resolved detail and noise.

With today's multi-megapixel imaging equipment and high density CCDs, oversampling is a common occurrence; there is only so much detail that seeing conditions allow for with a given setup. Beyond that it is impossible to pick up fine detail. Once detail no longer fits in a single pixel, but instead gets 'smeared out' over multiple pixels due to atmospheric conditions (resulting in a blur), binning may turn this otherwise useless blur into noise reduction. 'Binning your data may make an otherwise noisy and unusable data set usable again, at the expense of 'useless' resolution.

The 'Bin' module was created to provide a freely scalable alternative to the fixed 2×2 (4x reduction in resolution) or 4×4 (16x reduction in resolution) software binning modes commonly found in other software packages or modern consumer digital cameras and DSLRs (also known as 'Low Light Mode'). As opposed to these other binning solutions, the StarTools' 'Bin' module allows you to bin your data (and gain noise reduction) by the amount you want – if your data is seeing-limited (blurred due to adverse seeing conditions) you are now free to bin your data until exactly that limit and you are not forced by a fixed 2×2 or 4×4 mode to go beyond that.

Similarly, deconvolution (and subsequent recovery of detail that was lost due to atmospheric conditions) may not be a viable proposition due to the noisiness of an initial image. Binning may make deconvolution an option again. The StarTools 'Bin' module allows you to determine the ratio with which you use your oversampled data for binning and deconvolution to achieve a result that is finely tuned to your data and imaging circumstances of the night(s).

Core to StarTools' fractional binning algorithm is a custom built anti-aliasing filter that has been carefully designed to not introduce any ringing (overshoot) and, hence, to not introduce any artifacts when subsequent deconvolution is used on the binned data.

Usage

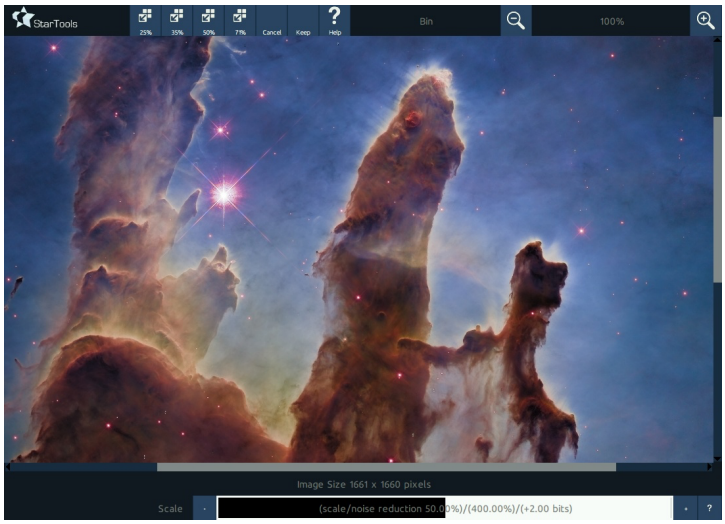

Operating the 'Bin' module couldn't be easier with just one slider doing all the work.

The 'Bin' module is operated with just a single parameter. This parameter controls the amount of binning that is performed on the data. The new resolution is displayed ('New Image Size X x Y') , as well the single axis scale reduction, the Signal-to-Noise-Ratio improvement and the increased bit-depth of the new image.

When to 'Bin'?

Data binning is a data pre-processing technique used to reduce the effects of minor observation errors. Many astrophotographers are familiar with the virtues of hardware binning. The latter pools the value of 4 (or more) CCD pixels before the final value is read. Because reading introduces noise by itself, pooling the value of 4 or more pixels reduces this 'read noise' also by a factor of 4 (one read is now sufficient, instead of having to do 4). Of course, by pooling 4 pixels, the final resolution is also reduced by a factor of 4. There are many, many factors that influence hardware binning and Steve Cannistra has done a wonderful write-up on the subject on his starrywonders.com website. It also appears that the merits of hardware binning are heavily dependent on the instrument and the chip used.

Most OSCs (One-Shot-'Color') and DSLR do not offer any sort of hardware binning in color, due to the presence of a Bayer matrix; binning adjacent pixels makes no sense, as they alternate in the color that they pick up. The best we can do in that case is create a grayscale blend out of them. So hardware binning is out of the question for these instruments.

So why does StarTools offer software binning? Firstly, because it allows us to trade resolution for noise reduction. By grouping multiple pixels into 1, a more accurate 'super pixel' is created that pools multiple measurements into one. Note that we are actually free to use any statistical reduction method that we want. Take for example this 2 by 2 patch of pixels:

7 7
3 7

A 'super pixel' that uses simple averaging yields (7 + 7 + 3 + 7) / 4 = 6. If we suppose the '3' is anomalous value due to noise and '7' is correct, then we can see here how the other 3 readings 'pull up' the average value to 6; pretty darn close to 7.

We could use a different statistical reduction method (for example taking the median of the 4 values) which would yield 7, etc. The important thing is that grouping values like this tends to filter out outliers and make your super pixel value more precise.

Sensor resolution may be going up, but the atmosphere's resolution will forever remain the same - buying a higher resolution instrument will do nothing for the detail in your data in that case!

'Binning and the loss of resolution

But what about the downside of losing resolution? That super high resolution may have actually been going to waste! If for example your CCD can resolve detail at 0.5 arcsecs per pixel, but your seeing is at best 2.0 arcsecs, then you effectively have 4 times more pixels than you need to record one 1 unit of real resolvable celestial detail. Your image will be 'oversampled', meaning that you have allocated more resolution than the signal really will ever require. When that happens, you can zoom in into your data and you will notice that all fine detail looks blurry and smeared out over multiple pixels. And with the latest DSLRS having sensors that count 20 million pixels and up, you can bet that most of this resolution will be going to waste at even the most moderate magnification.

'Sensor resolution may be going up, but the atmosphere's resolution will forever remain the same - buying a higher resolution instrument will do nothing for the detail in your data in that case!'

This is also the reason why professional CCDs are typically much lower in resolution; the manufacturers rather use the surface area of the chip for coarser but more deeper, more precise CDD wells ('pixels') than squeezing in a lot of very imprecise (noisy) CCD wells (it has to be said the latter is a slight oversimplification of the various factors that determine photon collection, but it tends to hold).

'Binning to undo the effects of debayering interpolation

There is one other reason to bin OSC and DSLR data to at least 25% of its original resolution; the presence of a bayer matrix means that (assuming an RGGB matrix) after applying a debayering (aka 'demosaicing') algorithm, 75% of all red pixels, 50% of all green pixels, and another 75% of all blue pixels are completely made up!

Granted, your 16MP camera may have a native resolution of 16 million pixels, however it has to divide these 16 million pixels up between the red, green and blue channels! Here is another very good reason why you might not want to keep your image at native resolution. 'Binning to 25% of native resolution will ensure that each pixel corresponds to one real recorded pixel in the red channel, one real recorded pixel in the blue channel and two pixels in the green channel (the latter yielding a 50% noise reduction in the green channel).

There are, however, instances where the interpolation can be undone if enough frames are available (through sub-pixel dithering) to have exposed all sub-pixels of the bayer matrix to real data in the scene (drizzling).

Fractional binning

StarTools' binning algorithm is a bit special in that it allows you to apply 'fractional' binning; you're not stuck with pre-determined factors (ex. 2x2, 3x3 or 4x4). You can bin exactly the amount that achieves a single unit of celestial detail in a single pixel. In order to see what that limit is, you simply keep reducing resolution until no blurriness can be detected when zooming into the image. Fine detail (not noise!) should look crisp. However, you may decide to leave a little bit of blurriness to see if you can bring out more detail using deconvolution.

Wrap Up

Description

The 'Bin' algorithm trades off resolution for an improved Signal-to-Noise Ratio (SNR) - it doesn't just change the scale.

The algorithm yields correct results even at arbitrary sizes (not just powers of 2) by applying an anti-aliasing filter at the proper cutoff frequency corresponding to the new image size.

When to use

- After initial global stretch ('AutoDev').
- When you want to improve the SNR - if the image is oversampled.
- If the data is noisy you may want to bin to improve the SNR - even if the image is not oversampled and you will lose some detail.

What result to look for

- Is the image still oversampled? For a description of oversampling see below.
- Zoom in and look at smallest stars when unbinned - are they spread over 3 or more pixels in any direction? - if they are the image is oversampled - the combination of the seeing, optics, focus and camera have lead to this spreading.
- Zoom out - make sure the reduction in resolution hasn't caused the more prominent stars to lose their roundness and other detail in the brighter parts of the image to become angular and 'boxy'.

Description of Controls

Presets

Define a preset amount of binning.

- 25% - Reduce image resolution to 25% of what it was - Improve SNR by approx. 4x
- 35% - Reduce image resolution to 35% of what it was - Improve SNR by approx. 2.9x
- 50% - Reduce image resolution to 50% of what it was - Improve SNR by approx. 2x
- 71% - Reduce image resolution to 71% of what it was - Improve SNR by approx. 1.7x

These SNR figures assume negligible read noise.

Scale

- Scale sets the reduction in resolution (e.g. 25% - the number of pixels along one axis has been reduced to 25%)
- Also shown is the corresponding noise reduction factor (1600.00%) and the bit depth improvement (+4.00 bits)
- 100% = no reduction in scale, 0% noise reduction factor, 0 bit depth improvement
- Default is (scale/noise reduction 50%)/(400.0%)/(+2.00 bits)

'Crop': Express cropping Tool with Switchable Luminance and Chroma Preview

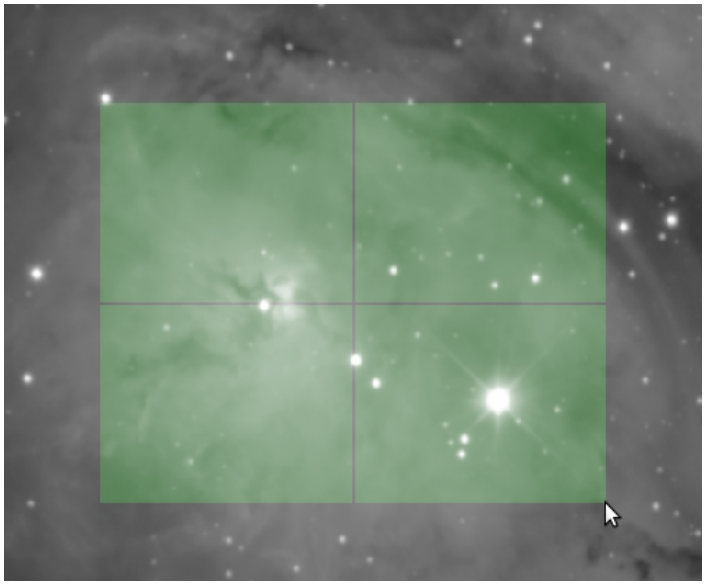

'cropping is as easy as clicking and dragging with the mouse

The crop module is an easy-to-use image cropping tool with quick aspect ratio presets and switchable luminance and chroma preview modes.
The module was designed to quickly find and eliminate stacking artifacts across luminance and chrominance data, as well as help with framing your object(s) of interest.

Usage

The 'Crop' module has a simple interface with 8 popular aspect ratio presets.

Using the crop module is fairly straightforward. The desired crop is created by clicking and dragging with the mouse the area to retain. Fine-tuning can be accomplished by changing the X1, Y1 and X2, Y2 coordinate pair parameters.
8 preset crops are available. Their names ('3:2', '2:3', '16:9', 9:16') denote the aspect ratio, while the minus or plus sign denotes their behavior:

- Presets with the greater-than ('—') sign will grow the current selection to achieve the selected aspect ratio
- Presets with the smaller-than ('++') sign will shrink the current selection to achieve the selected aspect ratio

A 'Color'/'NBAccent' button is available, which functions much like the 'Color'/'NBAccent' button in the Wipe module. Like in the Wipe module, it is only available when Compose mode is engaged (e.g. when luminance, chrominance and/or narrowband accents are being processed separately, yet simultaneously). The button allows you to switch the view between the luminance, chrominance and narrowband accent datasets that are being processed in parallel. The later is useful if, for example, you need to crop stacking artefacts that only exist in the chroma dataset and/or narrowband accent dataset, but not in the luminance dataset. Because chrominance data always remains linear and is never stretched like the luminance dataset, a courtesy (non-permanent) AutoDev is applied, so you can better see what is in the chrominance dataset. Likewise, a courtesy temporary AutoDev is applied to any narrowband accent data for that same purpose.

Wrap Up

Description

The 'Crop' module allows you to remove stacking artifacts and reframe your image.

When to use

- Before the 'Wipe' module. Normally just before or after the 'Bin' module (if used).
- After the 'Lens' module, if used. The 'Lens' module relies on the image being centered relative to the camera sensor - and can also warp the image when correcting distortion - so it is important this is done before using the 'Crop' module.

What result to look for

- Check the stacking artifacts have been removed.
- Check the image framing is how you want it.

Description of Controls:

- Presets: crop the image to selected ratio
    - 3:2< (v1.6) - adjust to 3:2 aspect ratio - by reducing the dimensions in one direction
    - 2:3< (v1.6) - adjust to 2:3 aspect ratio - by reducing the dimensions in one direction
    - 16:9< (v1.6) - adjust to 16:9 aspect ratio - by reducing the dimensions in one direction
    - 9:16< (v1.6) - adjust to 9:16 aspect ratio - by reducing the dimensions in one direction
    - 3:2> (v1.6) - adjust to 3:2 aspect ratio - by increasing the dimensions in one direction
    - 2:3> (v1.6) - adjust to 2:3 aspect ratio - by increasing the dimensions in one direction
    - 16:9> (v1.6) - adjust to 16:9 aspect ratio - by increasing the dimensions in one direction
    - 9:16> (v1.6) - adjust to 9:16 aspect ratio - by increasing the dimensions in one direction

- X1:
    - Sets the Left hand edge crop position - measured from the left.
    - Default is 0 pixels. Range is 0 pixels to max width.
- X2:
    - Sets the Right hand edge crop position - measured from the left (right).
    - Default is max width in pixels. Range is 0 pixels to max width.
- Y1:
    - Sets the Top edge crop position - measured from the top.
    - Default is 0 pixels. Range is 0 pixels to max height.
- Y2:
    - Sets the Bottom edge crop position - measured from the top (bottom).
    - Default is max height in pixels. Range is 0 pixels to max height.

'Wipe': Light Pollution, Vignetting and Gradient Removal

The 'Wipe' module detects, models and removes any source of unwanted light bias.

Usage

'Wipe' is able to detect - and correct for - various complex calibration problems and unwanted artificial signal sources. In addition to a gradient removal routine, it is to detect and model vignetting issues (including over-correction), as well as bias/darks issues.

Common calibration issues include:

- Vignetting manifests itself as the gradual darkening of a dataset towards the corners. It is ideally addressed through flat frame calibration when stacking.
- Amp glow is caused by circuitry heating up in close proximity to the CCD, causing localised heightened thermal noise (typically at the edges). On some older DSLRs and Compact Digital Cameras, amp glow often manifests itself as a patch of purple fog near the edge of the image.

Unwanted or artificial signal may include:

- Light pollution, moon glow, airglow, zodiacal light and gegenschein gradients are usually prevalent as gradual increases (or decreases) of background light levels from one corner of the image to another. Most earth-based acquisitions contain a gradient of some form, as even under pristine skies such gradients are prevalent.
- Signal bias is a fixed background levels which, contrary to a gradient, affects the whole image evenly. Most non-normalized datasets exhibit this.
- Amp glow is faint "glow" near one or more edges caused by local thermal noise from heat-dissipating electronics.

2 sources of unwanted light; a gradient starting at the upper right corner, and light pollution in the form of the typical yellow/ brown light. Also visible is vignetting, as seen in the darkening of the corners. Image courtesy of Charles Kuehne.

While highly effective, it is important to stress that 'Wipe's capabilities should not be seen as a replacement or long-term alternative to calibrating your datasets with calibration frames; calibrating your dataset with flats, darks and bias masters will always yield superior results. Flats in particular are the #1 way to improve your datasets and the detail you will be able to achieve in your images.

Preparing data for the 'Wipe' Module

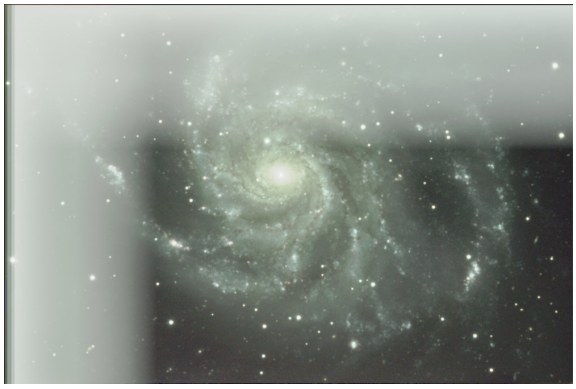

Leaving stacking artifacts in will cause 'Wipe' to interpret the anomalous data as true background, causing it to back off near the location of the artifacts.

It is of the utmost importance that 'Wipe' is given the best artifact-free, linear data you can muster.

Because 'Wipe' tries to find the true (darkest) background level, any pixel reading that is mistakenly darker than the true background in your image (for example due to dead pixels on the CCD, or a dust speck on the sensor) will cause 'Wipe' to acquire wrong readings for the background.

Halo around a simulated dust speck dark anomaly.

When this happens, 'Wipe' can be seen to 'back off' around the area where the anomalous data was detected, resulting in localised patches where gradient (or light pollution) remnants remain. These can often look like halos. Often dark anomalous data can be found at the very centre of such a halo or remnant.

The reason 'Wipe' backs off is that 'Wipe' (as is the case with most modules in StarTools) refuses to clip your data. Instead 'Wipe' allocates the dynamic range that the dark anomaly needs to display its 'features'. Of course, we don't care about the 'features' of an anomaly and would be happy for 'Wipe' to clip the anomaly if it means the rest of the image will look correct.

Fortunately, there are various ways to help 'Wipe' avoid anomalous data:

- A 'Dark anomaly filter' parameter can be set to filter out smaller dark anomalies, such as dead pixels or small clusters of dead pixels, before passing on the image to 'Wipe' for analysis.
- Larger dark anomalies (such as dust specks on the sensor) can be excluded from analysis by, simply by creating a mask that excludes that particular area (for example by 'drawing' a 'gap' in the mask using the Lasso tool in the Mask editor).
- Stacking artifacts can be cropped using the 'Crop' module.

Bright anomalies (such as satellite trails or hot pixels) do not affect 'Wipe'.

Edge located dark anomalies

Stacking artifacts are the most common dark anomalies located at the edges of your image. Failing to deal with them will lead to a halo effect near the edges of your dataset.

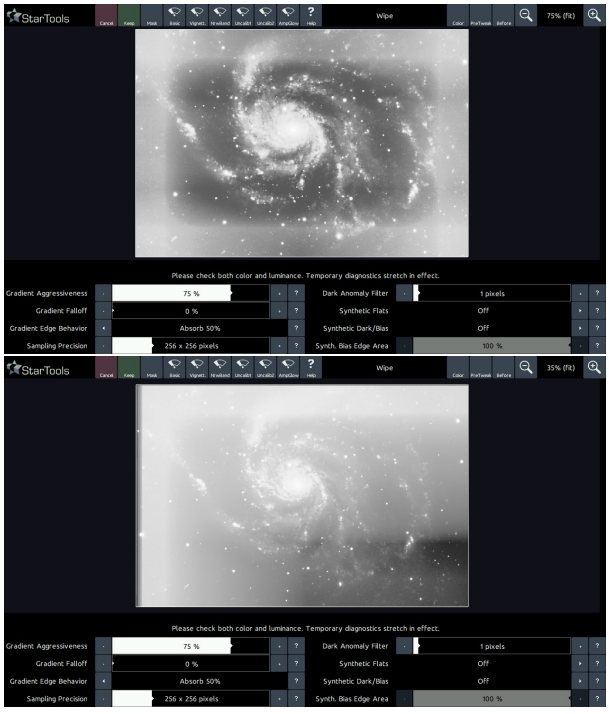

Beware of single-pixel artifacts around the edges; they will cause edge-located halos like these. Zoom-in to find them and use the 'Crop' module to eliminate them before using 'Wipe'.

Please remove any stacking artifacts before launching the 'Wipe' module. Failing to do so, will result in edge-located halos, like these.

Non-edge located dark anomalies

Dust specks, dust donuts, and co-located dead pixels all constitute dark anomalies and will cause halos around them if not taken care of. These type of dark anomalies are taken care of by masking them out so that 'Wipe' will not sample their pixels.

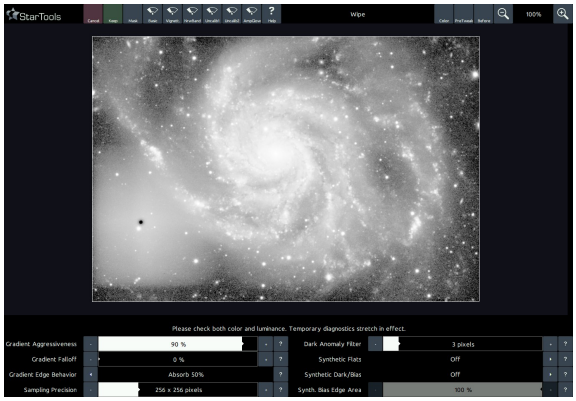

'Wipe' will generate halos around dark anomalies (e.g. darker than real background pixels), like this simulated dust speck

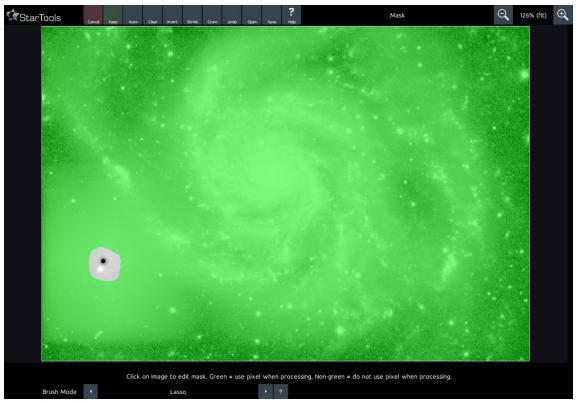

In cases where dark anomaly is too big for the 'Dark Anomaly 'Filter" parameter to filter out the pixels, you should mask such large dark anomalies out

'Wipe' no longer samples the pixels that are masked out, now allowing the dust speck to clip rather than elevating the local background to accommodate the dust speck in the dynamic range (causing the halo around it). The diagnostic stretch is doing its job, highlighting its presence

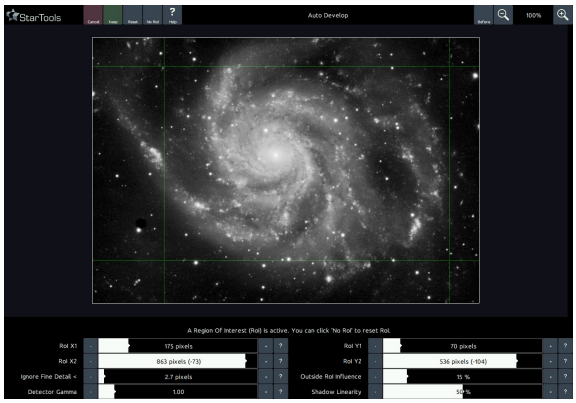

A subsequent global stretch in 'AutoDev' makes the dust speck a lot more inconspicuous

Operating the 'Wipe' Module

Once any dark anomalies in the data have successfully been dealt with, operating the 'Wipe' module is fairly straightforward.

To get started quickly, a number of presets cover some common scenarios;

- 'Basic' is the default for the 'Wipe' module and configures parameters that work with most well calibrated datasets.
- 'Vignetting' configures additional settings for vignetting modeling and correction.
- 'Narrowband' configures 'Wipe' for narrowband datasets which usually only need a light touch due to being less susceptible to visual spectrum light pollution.
- 'Uncalibrated 1' configures 'Wipe' for completely uncalibrated datasets, for cases where calibration frames such as flats were - for whatever reason - not available. This preset should be used as a last resort.
- 'Uncalibrated 2' configures 'Wipe' for poor quality, completely uncalibrated datasets. The settings used here are even more aggressive than 'Uncalibrated 1'. This preset too should only be used as a last resort.

Internally, the module's engine models three stages of calibration similar to an image stacker's calibration stages:

1. synthetic bias/darks modeling and correction (subtraction)
2. synthetic flats modeling and correction (division)
3. gradient modeling and correction (subtraction).

Any issues specified and/or detected are modeled during the correct stage and its results feeds into the next stage.

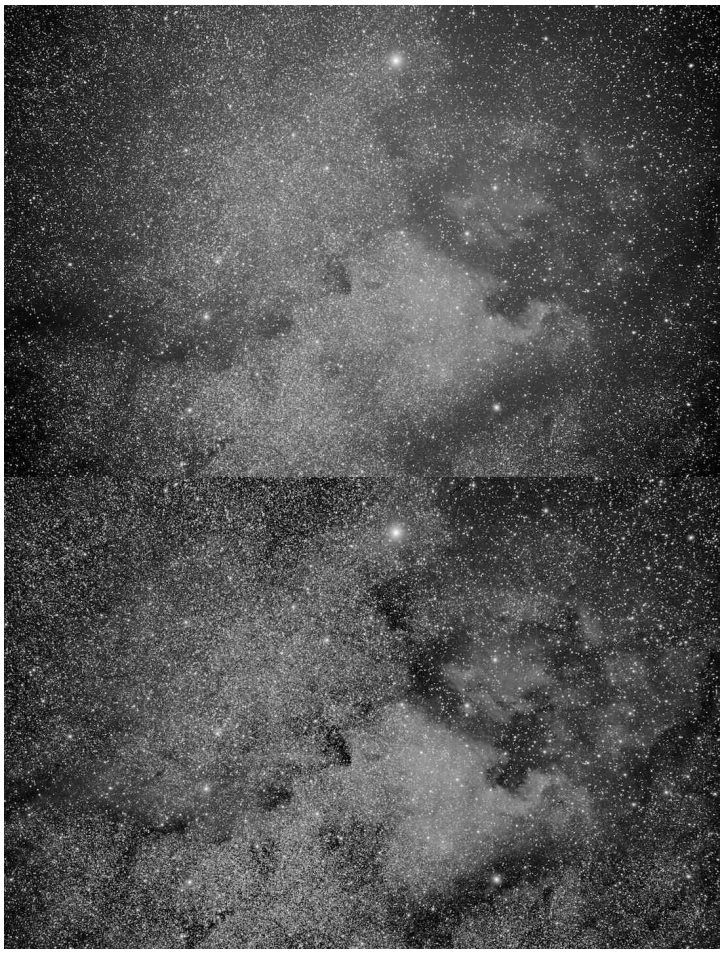

The 'Uncalibrated' presets model - and correct for - vignetting, as well as gradients.

Synthetic bias/darks modeling

The 'Wipe' module is able to detect horizontal or vertical banding and correct for this. Multiple modeling algorithms are available to detect and mitigate banding. A defective

sensor column repair feature is also available that attempts to recover data that was transposed but not lost, rather than interpolating or 'healing' it using neighboring pixels.

'Wipe', as part of the three stages of operation, is able to recover data from defective sensor rows and columns.

Any issues specified and/or detected are modeled during the correct stage and its results feeds into the next stage.

Synthetic flats modeling

The 'Wipe' module is able to quickly match and model a natural illumination falloff model to your dataset with correction for cropping and off-axis alignment.

Fixed pattern noise and correlated artifact filtering

The 'Correlation Filtering' parameter specifies the size of correlation artifacts that should be removed. This feature can ameliorate correlation artifacts that are the result of dithering, debayering or fixed pattern sensor cross-talk issues. Correlated noise and artifacts can look like detail to both humans and algorithms. By pre-emptively filtering out these artifacts, modules will be able to better concentrate on the real detail in your dataset and image, rather than attempting to preserve these artifacts.

The usage of this filter is most effective on oversampled data where the artifacts are clearly smaller than the actual resolved detail.

Gradient modeling and subtraction

'Wipe' discerns gradient from real detail by estimating undulation frequency. In a nut shell, real detail tends to change rapidly from pixel to pixel, whereas gradients do not. The 'Aggressiveness' specifies the undulation threshold, whereby higher 'Aggressiveness' settings latch on to ever faster undulating gradients. At high 'Aggressiveness' settings, be mindful of 'Wipe' not 'wiping' away any medium to larger scale nebulosity. To 'Wipe', larger scale nebulosity and a strong undulating gradients can look like the same thing. If you are worried about 'Wipe' removing any larger scale nebulosity, you can designate an area off-

limits to its gradient detection algorithm, by means of a mask that masks out that specific area. See the 'Sample revocation' section for more details.

After 'Wipe'

Because 'Wipe's impact on the dynamic range in the image is typically very, very high, a (new) stretch of the data is almost always needed. This is so that the freed up dynamic range, previously occupied by the gradients, can now be put to good use to show detail. 'Wipe' will return the dataset to its linear state, however with all the cleaning and calibration applied. In essence, this makes a global re-stretch using 'AutoDev' or 'FilmDev' is mandatory after using 'Wipe'. From there, the image is ready for further detail recovery and enhancement, with color calibration preferably done as one of the last steps.

The diagnostics stretch

Because 'Wipe' operates on the linear data (which is hard to see), a new, temporary automatic non-linear stretch is reapplied on every parameter change, so you can see what the module is doing. The diagnostics stretch is designed to show your dataset in the worst possible light on purpose, so you can diagnose issues and remedy them. The sole purpose of this stretch is to bring out any latent issues such as gradient, dust donuts, dark pixels. That is, it is entirely meant for diagnostics purposes inside the 'Wipe' module and in no way, shape or form should be regarded as a suggested final global stretch.

'The diagnostics stretch is designed to show your dataset in the worst possible light on purpose, so you can diagnose issues and remedy them'

Automatically separated luminance and chrominance datasets

If Compose mode is engaged (see Compose module), 'Wipe' processes luminance (detail) and chrominance (color) separately, yet simultaneously. If you process in Compose mode (which is recommended), you should check both the results for the luminance and chrominance portion of your image. Before keeping the result, the 'Wipe' module will alert you to this once, if you have not done so.

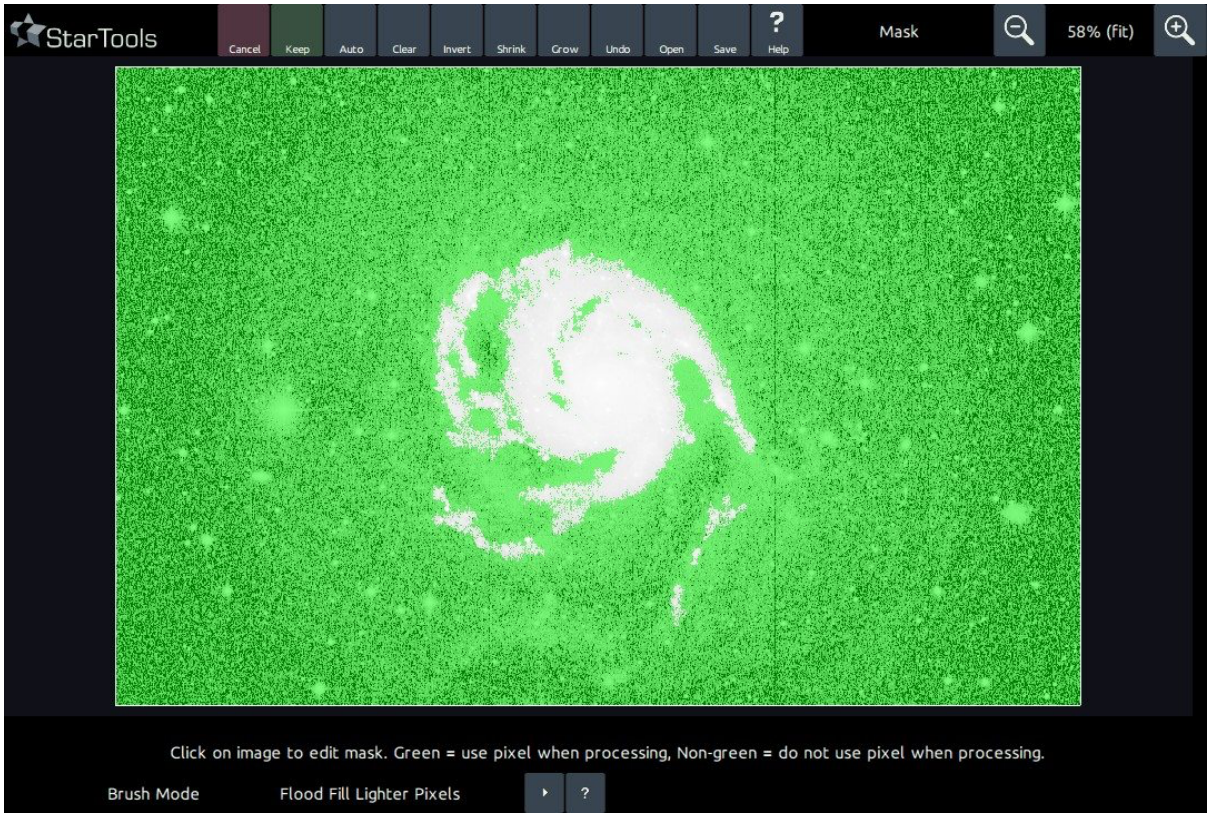

Sample revocation

At very high Aggressiveness settings to deal with extremely challenging data, you can use sample revocation to tell 'Wipe' where it should NOT look for background. This may help protect areas of detail you are certain are real, and should achieve superior results.

With the exception of the previously mentioned larger "dark anomalies" (such as dust donuts or clumps of dead pixels), it is typically unnecessary to provide 'Wipe' with a mask. However if you wish to give 'Wipe' specific guidance as to which areas of the image to include in the model of the background, then you may do so with a mask that describes which where background definitely does not exist.
This is a subtle but important distinction from background extraction routines in less sophisticated software, where the user must "guess" where background definitely exists. The former is easy to determine and is readily visible, whereas the latter is usually impossible to see, precisely because the background is mired in gradients. In other words, StarTools' 'Wipe' module works by sample revocation ("definitely nothing to see here"), rather than by the less optimal (and possibly destructive!) sample setting ("there is background here").
Analogous to how sample setting routines yield poor results by accidentally including areas of faint nebulosity, the opposite is the case in 'Wipe': accidentally masking out real background will yield the poorer results in 'Wipe'. Therefore, try to be conservative with what is being masked out. If in doubt, leave an area masked in for 'Wipe' to analyse.

Design philosophy and limitations

As with all modules in StarTools, the 'Wipe' module is designed around robust data analysis and algorithmic reconstruction principles. The data should speak for themselves and manual touch-ups or subjective gradient model construction is avoided as much as possible.
In general, StarTools' 'Wipe' module should yield superior results, retaining more faint detail and subtle large-scale nebulosity, compared to traditional manual gradient model construction routines. However, exceptions arise where gradients undulate (e.g. rise or fall) faster than the detail in the image due to atypical acquisition issues (incorrect flat frames, strongly delineated localised light pollution domes). Human nor machine will be able to discern detail objectively or with certainty. As a result 'Wipe' will, likewise, struggle in such cases.

Wrap Up

Description

To remove background color casts and gradients such as unwanted light (light pollution, moon, street lights), vignetting, amp glow and color casts. This does not mean that the background becomes a nice neutral grey. Rather it makes sure that the background is luminance bias-free across all color channels. Subsequent color calibration will neutralize the chrominance (color) information. It also compensates for sensor column and row defects.

When to use

After using Initial 'AutoDev', 'Lens', 'Crop' and 'Bin' as needed.

Only use 'Wipe' once - although you can use 'Undo' on main screen to immediately go back and redo the 'Wipe' Module if you aren't happy with the results.

What result to look for

- After applying 'Wipe' the stretched image should have a uniform and dark background.

- There should be no remaining background hue or color cast. If there is check the masking and redo, and try increasing the Gradient Aggressiveness.
- Look at the stretched image background for halos and other casts - this may be due to dark features such as dead pixels, dust, trees etc. which have not been removed. These can be removed using the Mask or Dark Anomaly 'Filter' as described below.
- If some detail has been wiped reduce the Gradient Aggressiveness - or Mask the area affected.
- If casts exist try increasing the Gradient Aggressiveness.
- After use:
- The 'Wipe' module should make the background darker. It is important to redo the global stretch (using Develop or 'AutoDev') to take advantage of the extra dynamic range made available.

Ways of getting better results

- Use Flats when stacking so that the Vignetting, dust specks, etc. will have already been accurately compensated for. This will simplify the 'Wipe' to work specifically on light gradients.
- If 'Wipe' does not reduce the color cast well it may be that the image has a lot of noise. In this case, increasing the Dark Anomaly filter can make a difference.

Description of Controls

Mask

For general instructions on using masks see Mask.

- Mask out (make non-green) the main subject and any dust specks, mountains, trees etc, using the 'Lasso' Brush Mode.
- This will stop 'Wipe' trying to use them in analyzing how to wipe the image - the whole image will still be wiped.
- If using the 'Amp Glow' preset - select (make green) the pixels that are affected by amp glow.
- You must use Develop or 'AutoDev' to bring out the image details before using 'Wipe' if you want to see them in the Mask.

Presets

The following buttons pre-set the control settings to values most suited to the intended purpose:

- Basic (v1.7+) - Simple wipe - gradual changes in background light level across the image - e.g. from light pollution or the moon.
- Vignetting - for removing any background light level that falls away at the corners - use flats instead if you can.
- Narrow Band - optimized wipe for narrow band images - Light pollution and gradients are less of a problem with narrow band data.
- Uncalibrated (v1.7+) - use when no calibration has been done - tries to correct for vignetting etc.
- Uncalibrated 2
- Amp Glow - a brightness normally at one edge of the image due to an uneven heating effect from the camera electronics.

'Color'/Luminance

In 'Compose' mode the 'Wipe' module works on the luminance and color datasets separately but in parallel. You can see the results of each by pressing this button. You will be reminded if you haven't checked the color view.

Gradient Aggressiveness

Sets how aggressive 'Wipe' should be:

- 'Wipe' searches for a constant background level of brightness around each pixel in the image - the search area is governed by the 'Aggressiveness' setting - the higher the value the smaller the area searched to establish the background level. See the description in the topic 'Wipe' feature - when to use Mask.
- Default is 75% for Basic and Uncalibrated presets, 0% for NrwBand, and not set for Vignetting or Amp Glow.
- Mask out any objects of interest if using a particularly aggressive setting.
- Change if gradients remain - up to 95% in extreme cases. Reduce from the default if detail is being wiped.

Dark Anomaly 'Filter'

Use this to stop dark spots in the image (such as dead pixels, small dust specks and small scratches) affecting the 'Wipe'.

- Default is 1 pixel. Range is 'Off'(0) to 30 pixels.
- Set to the size of the dark area. Larger dark areas (including larger dust specks) should be deselected using the mask instead.
- Try Increasing to 2-10 pixels - look at the effect - continue until no further benefit.

Gradient Falloff

Sets how gradient undulation changes closer to the corners.

- This can help get a flat field in cases where there is uneven lighting towards the corners but the basic vignetting modeling set by Synthetic Flats can't model light correctly or sufficiently.
- Default is 0% for Basic and NrwBand, 75% for Vignetting and Uncalibrated, and not set for Amp Glow preset.

Synthetic Flats

Specifies what standard vignetting model to try to fit to the dataset.

- Off - no model used
- Vignetting
- Inverse Vignetting (Overcorrection)
- Default is 'Off' for Basic preset, 'Vignetting' for Vignetting and Uncalibrated presets, and not set for Narrow Band and Amp Glow presets.

Gradient Edge Behavior

Defines how 'Wipe' should assume the gradient falls off beyond the edge of the image.

- Grow opposite axis - Best for poorly calibrated datasets and those that have a lot of changes in bias.
- Bounce Back - Best for well calibrated datasets and datasets that contain image-filling nebulosity.
- Absorb 50% - Good choice for datasets that fall between the two types described above.
- Defaults are 'Absorb 50%' for 'Basic' preset, 'Bounce Back' for Narrow Band, 'Grow opposite axis' for Uncalibrated, and not set for Vignetting and Amp Glow presets.

Synthetic Dark/Bias

Sets what darks and bias issues 'Wipe' should try to detect and repair.
Tries to detect and repair the effects of:

- Off - None
- Defective Sensor Columns - One or more adjacent sensor columns giving a strong positive or negative bias.
- Defective Sensor Rows - One or more adjacent sensor rows giving a strong positive or negative bias.

- Vertical Bias - Sensor columns having subtly different biases compared to their neighbors across the entire height of the image.
- Horizontal Bias - Sensor rows having subtly different biases compared to their neighbors across the entire width of the image.
- Adaptive Vertical Bias - Sensor columns having subtly different biases compared to other columns across the height of the image.
- Adaptive Horizontal Bias - Sensor rows having subtly different biases compared to other rows across the entire width of the image.
- Adaptive Multi-Axis Bias Aggressive - Sensor rows and columns having subtly different biases compared to other sensors across the entire width of the image - Aggressive correction.
- Adaptive Multi-Axis Bias Conservative - Sensor rows and columns having subtly different biases compared to other sensors across the entire width of the image - Conservative correction.
- Defaults are 'Off' for 'Basic' preset, 'Adaptive Multi-Axis Bias Conservative' for Uncalibrated, 'Adaptive Multi-Axis Bias Aggressive' for Amp Glow preset, and not set for Vignetting and Narrow Band presets.

Sampling Precision

Defines the image 'resolution' that 'Wipe' uses when analyzing the image prior to wiping.
Range is 128x128, 256x256, 512x512, 1024x1024 or 2048x2048.

- Default is 256x256 pixels for 'Basic', 'Vignetting' and 'Uncalibrated' presets, 128x128 for Narrow Band, not set for 'Amp Glow' preset.
- Change from the default to a larger number if the image has rapidly changing gradients and some background gradient has not been compensated for.

'Synth'. Bias Edge Area

Sets how far from the edge towards the image centre 'Wipe' should detect and correct dark/bias issues (i.e. those issues defined by the Synthetic Dark/Bias setting).

- Applies only when Synthetic Dark/Bias is used.
- Use this setting to target the correction of issues that happen near the edge of the image (e.g. thermal noise/bias) without affecting the central part of the image.
- Default is 33% for Uncalibrated and Amp Glow presets.

'Contrast': Local 'Contrast' Optimization

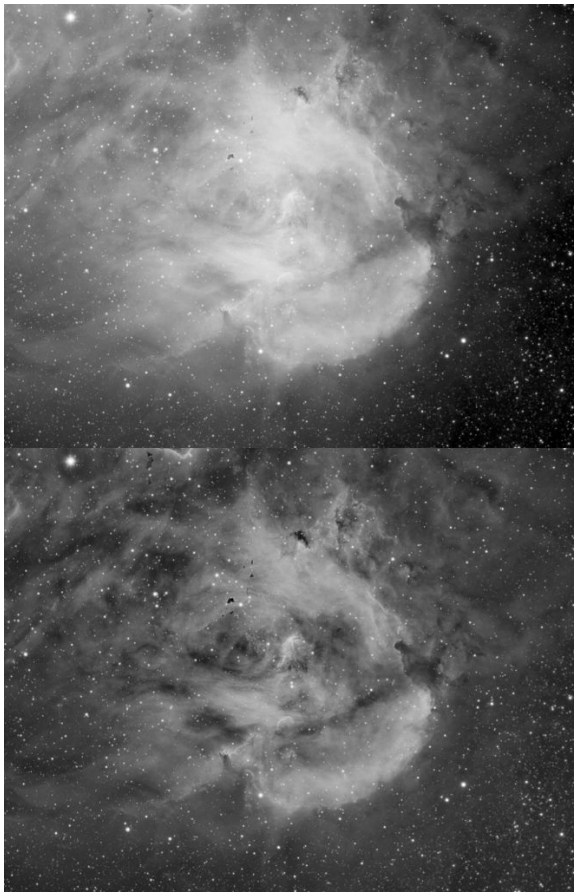

Top: globally stretched data without further local dynamic range optimization.

Bottom: Large to medium scale local dynamic range optimization with the 'Contrast' Module.

The 'Contrast' module optimizes local dynamic range allocation, resulting in better contrast, reducing glare and bringing out faint detail.

It operates on medium to large areas, and is especially effective for enhancing contrast and detail unobtrusively in image-filling nebulas, globular clusters and galaxies.

Usage

We will use this Hydrogen Alpha dataset of Meloitte 15, acquired by Jim Misti to demonstrate the 'Contrast' module. Pre-processing consisted of a simple 'Wipe' and 'AutoDev'.

The 'Contrast' module works by evaluating minimum and maximum brightness in a pixel's local area, and using these statistics to adjust the pixel's brightness. The size of the local areas is controlled by the 'Locality' parameter. In essence, the 'Locality' parameter controls how 'local' the dynamic range optimization is allowed to be. You will find that a higher 'Locality' value with all else equal, will yield an image with areas of starker contrast. More generally, you will find that changing the 'Locality' value will see the 'Contrast' module

take rather different decisions on what (and where) to optimise. The rule of thumb is that a higher 'Locality' value will see smaller and 'busier' areas given priority over larger more 'tranquil' areas.

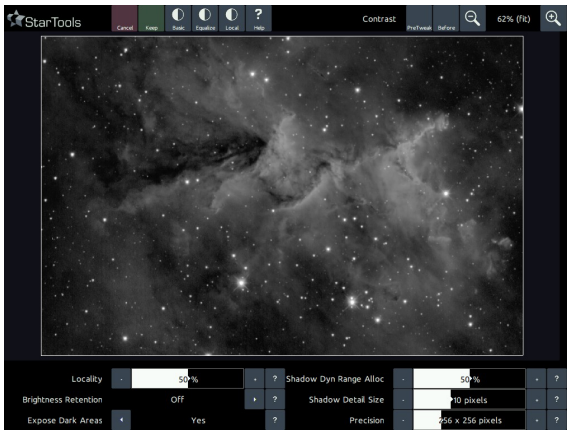

Default settings corresponding with the 'Basic' preset. The image is locally darkened, removing glare.

As alluded before, The 'Shadow Dynamic Range Allocation' parameter controls how heavily the 'Contrast' module "squashes" the dynamic range of dark, smaller scale features it deems "unnecessary"; by de-allocating dynamic range that is used to describe larger features and re-allocating it to interesting local features, the de-allocation necessarily involves reducing the larger features' dynamic range, hence "squashing" that range. Very low settings may appear to clip the image in some extreme cases (though this is not the case). For those familiar with music production, the 'Contrast' module is analogous to a compressor, but for your images instead.

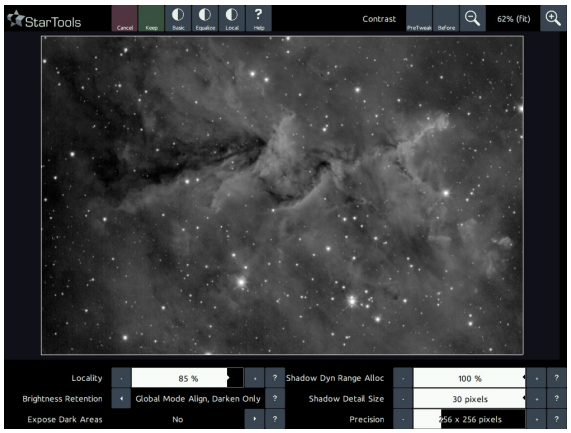

The 'Local' preset uses a higher Locality setting, thereby optimizing dynamic range at a local level much more aggressively.

The 'Brightness Retention' feature attempts to retain the apparent brightness of the input image. It does so through calculating a non-linear stretch that aligns the histogram peak (statistical 'mode') of the old image with that of the new image. An optional 'Darken Only' operation only keeps pixels from the resulting image that are darker than the input image.

The 'Expose dark areas' option can help expose detail in the shadows by normalizing the dynamic range locally; making sure that the fully dynamic range is used at all times.

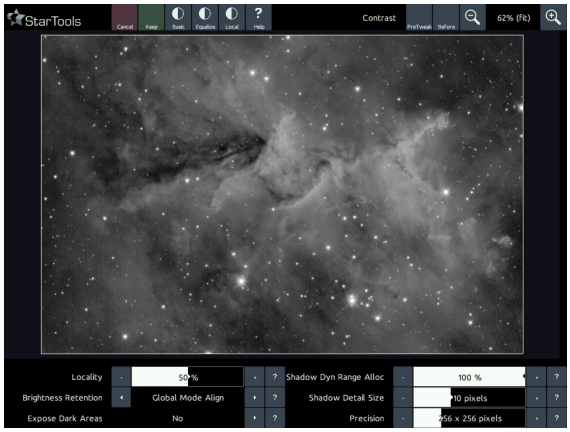

The 'Equalize' presets uses settings that darkens bright areas and lightens dark areas. The result is a more 'tranquil' image with small detail standing out well.

Wrap Up

Description

To optimise medium-to-large local contrast (dynamic range) by doing local stretching.

- Complements the 'HDR' module which optimizes medium-to-small local contrast, and the 'Sharp' module for detail enhancements.
- Adjusts stretch locally on medium-to-large elements to improve visibility.
- May Compress Dark Anomalies for extended shadow range.
- Similar in approach to the 'Wipe' module but without the masking capability.

When to use

- After final global stretch (Develop or 'AutoDev' modules) and optionally after using the 'Decon' module to reduce atmospheric effects and the 'Heal' or 'Crop' modules to remove larger Dark Anomalies.
- Use before the 'Color' module.
- It may be that if you have used the 'Wipe' module the scope for further enhancement is limited.

What result to look for

- Increased contrast in areas of medium to large detail within the image.
- If halos appear it is likely that there are hot or dead pixels that have not been taken into account.
- If detail is being lost - try increasing the 'Locality' setting.

Ways of getting better results

- If you can't get the results you want with the 'Contrast' module, try the 'HDR' module or the 'Sharp' module to highlight medium to small detail.

Description of Controls:

Presets

Following presets are available to create different contrast characters:

- 'Basic' (default) will darken the image locally reducing glare
- 'Local' will aggressively optimize dynamic level locally
- 'Equalize' will darken highlights and lighten shadows creating a tranquil look

Locality

A higher setting will enhance contrast for busy regions, A lower setting for more tranquil regions

- Default Value is 50% (Basic, Equalize), 85% (Local).
- Range is Off (0) to 100%.

Brightness Retention

Tries to retain the Brightness of the original image by applying a non-linear stretch aligning the histogram peaks.

- Default Value is 'Off' (Basic), 'Global Mode Align' (Equalize), 'Global Mode Align, Darken Only' (Local). Options are 'Off', 'Global Mode Align', 'Global Mode Align, Darken Only'
- 'Darken Only' will only keep pixels of new image which are darker than in original.

Expose Dark Areas

Allows the algorithm to brighten parts of the image to improve contrast.

- Default Value is Yes (Basic), No (Equalize, Local).
- Options are Yes, No

Shadow Dynamic Range Allocation

Defines how much dynamic range is given to any shadows below the 'Shadow Detail Size' threshold which are considered not relevant.

- Default Value is 50% (Basic), 100% (Equalize, Local).
- Range is Off (0) to 100%.
- This option is greyed out if 'Shadow Detail Size' is set to 'Off'

Shadow Detail Size

Sets the size below which 'Shadow Dynamic Range Allocation' kicks in.

- Increasing this parameter will let larger dark areas squashed, appearing with less detail
- Default Value is 10 pixels (Basic, Equalize), 30 pixels (Local) .
- Range is Off (0) to 100%.
- Common values are 2-30 pixels.

Precision

Specifies the size and amount of samples the module takes in calculating the background level.

- Default Value is 256x256 pixels (all presets).
- Possible values are 128x128, 256x256, 512x512, 1024x1024, or 2048x2048.
- Where the gradients change quickly a higher precision (smaller sample size) may be needed.
- Change from the default is not normally needed.

'HDR': Automated Local Dynamic Range Optimization

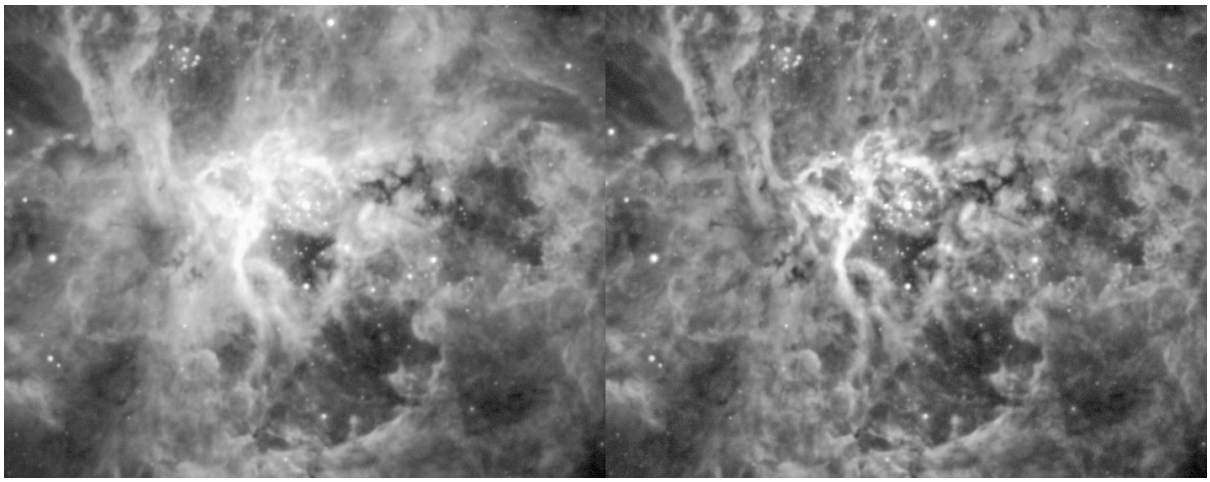

The HDR module puts you in full, intuitive control of local dynamic range allocation, without introducing artifacts or making the image look unnatural. Data acquisition by Marc Aragnou.

The HDR (High Dynamic Range) module optimises local dynamic range, recovering small to medium detail from your image. The module intuitively and effortlessly lets you resolve detail in bright galaxy cores, faint detail in nebulas and works just as well on solar, lunar and planetary images.

'The HDR (High Dynamic Range) module optimises local dynamic range, recovering small to medium detail from your image.'

Local histogram remapping, solves for the "ideal" luminance value per-pixel, based on its place in a local histogram, taking into account maximum spatial(!) contrast values.

Signal evolution Tracking-driven noise grain rejection ensures that - normally - noise-prone local histogram equalization (LHE), yields more robust estimates for signal/detail vs noise grain.

The HDR module operates exclusively on the luminance component of your image, retaining any coloring from the input image.

Launching the HDR module

Depending on the size (X * Y resolution) of the dataset at hand, the once-off initial processing/analysis may take some time, particularly at high resolution datasets and high 'Context Size' settings. Note that this processing/analysis is repeated every time the 'Context Size' parameter is changed, or when a new preview area is specified. Processing times may be cut by opting for a lower precision local gamma correction solving stage via the 'Quality' parameter.

However, once this initial processing/analysis has completed any parameter modification that does not involve 'Context Size', will complete virtually in real-time.

Presets

From left to right; Original, 'Reveal' preset, 'Tame' preset, 'Optimize' preset, 'Equalize' preset. Data acquisition by Jim Misti.

As with most modules in StarTools, the HDR module comes with a number of universally applicable presets that demonstrate settings for various use case:

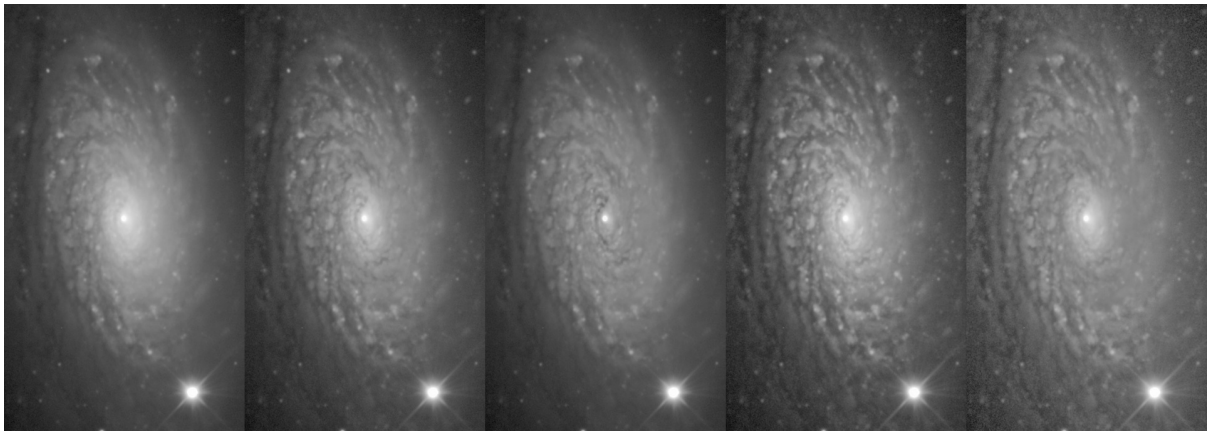

Reveal:
corresponds to the default settings, and combines moderate local gamma correction for the highlights with moderate local detail enhancement in both the shadows and highlights. This preset (and default setting) tends to be a generally applicable example.

Tame:
targets detail recovery in the highlights by applying aggressive local gamma correction in larger highlight areas. This presets demonstrates the HDR module's excellent ability to bring larger areas in the highlights under control and reveal any detail they might contain. This preset is, for example, useful to bring bright galaxy cores under control and reveal their detail.

Optimize:
targets and accentuates smaller detail in both shadows and highlights equally.

Equalize:
pulls both dim and bright larger contiguous areas into the midtones equally

Parameters

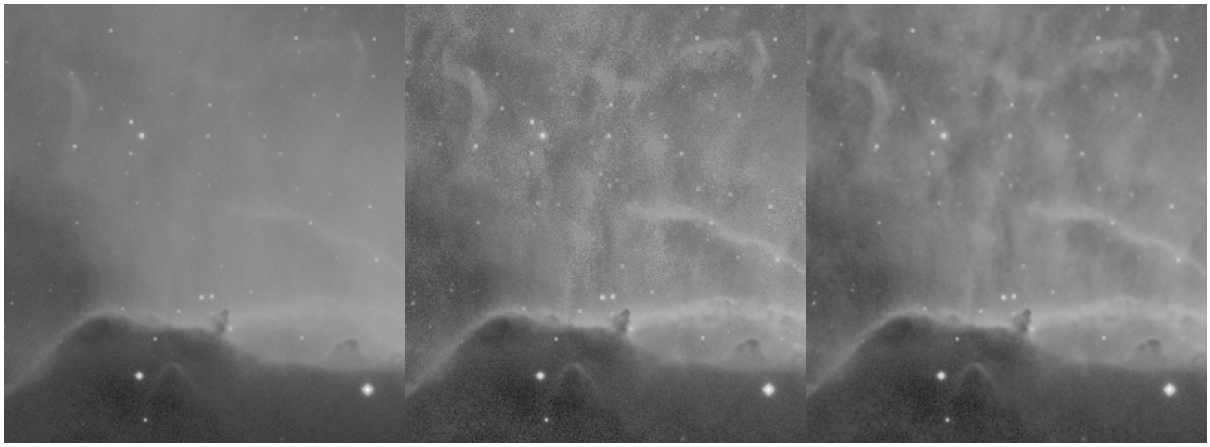

200% zoom image, showing the effects of signal evolution Tracking on image quality. Left: original, Middle: 'Optimize' preset detail recovery with 'Signal Flow' set to 'Visual As-Is', Right: 'Optimize' preset detail recovery with 'Signal Flow' set to 'Tracked'.

Evaluating the effect of the above presets, the intuitive nature of the parameters become clear:

- The 'Highlights Detail Boost' and 'Shadows Detail Boost' parameters generally provide a means to accentuate existing detail without affecting the brightness of larger contiguous areas, preserving that context.
- The 'Gamma Highlights' and 'Gamma Shadows' parameters generally provide a great dynamic range management solution for larger contiguous areas that are very bright (or dim), however contain smaller scale detail.
- The 'Gamma Smoothen' parameter controls the smoothness of the transition between differently locally stretched areas. Though the default value tends to be applicable to most situation, you can increase this value if any clear boundaries can be seen, or you decrease this value to get a clearer idea of which areas are modified (and how).
- The 'Signal Flow' parameter specifies the signal sources for the algorithm stack.;
- Tracked; uses a version of the signal that fully takes into account noise grain propagation in the signal. This allows the module to disregard recovered 'detail' in low-SNR areas that can be attributed to stretching the noise component of the signal, rather than the signal itself. Using this setting is highly recommended if you use HDR as part of a larger workflow, and plan on further detail recovery processing, particularly with algorithms like deconvolution.
- Visual As-Is; uses the stretched image (exactly as visible before launching the HDR module), without further noise propagation compensation.
- The 'Context Size' parameter controls the upper size of the detail/structures that may provide context for smaller detail. For example, reducing this parameter will see increasingly smaller detail being accentuated, with less and less concern for larger detail. A smaller 'Context Size' value may be appropriate in cases where resolving small detail is of higher priority and larger scale context is ideally ignored (for example globular clusters). The previously mentioned caveats for changing this parameter apply; high values tend to help preserving large scale context well, but may incur longer initial processing times. Processing times may be cut by opting for a lower precision local gamma correction solving stage via the 'Quality' parameter.

Artifacts

Results from the HDR module are generally artifact-free, unless using rather extreme values. This third iteration of the module was specifically engineered to further minimise the artifacts of alternative implementations (such as HDRWT and AHE/CLAHE). Star "bloat" or ringing artifacts should be negligible under normal operating conditions, while noise-induced "detail" development is suppressed through the incorporation of signal evolution Tracking statistics. Highlights vs Shadows manipulations are available independently, and applying just one or the other should not yield any detectable sharp transitions.

More caution should be exercised when using extreme values far outside of the defaults or presets.

Wrap Up

Description

To optimise an images medium-to-small local contrast. The 'HDR' module digs out small detail hidden in the image. It provides an automated curve manipulation. 'HDR' can operate on either the image detail brightness, color, or both.

When to use

After the final global stretch (Develop or 'AutoDev'), after 'Contrast' and prior the 'Decon' module (if used). Can be re-used if you have a number of different 'HDR' problems - at different scales for instance.

What result to look for

- Enhanced detail in the area you targeted.

After Use

- Re-use the module to target other areas as needed.
- Use the other optional modules as shown in the workflow. Then use the 'Color' module.

Description of Controls

Presets

Sets commonly used sets of parameter values.

- Reveal - reveals details in almost-overexposed cores. Default Setting.
- Tame - brings out detail in bright areas such as galaxy cores, like Equalize but doesn't brighten faint detail.
- Optimize - accentuates detail it can find in both shadows and highlights.
- Equalize - brings out detail in both shadows and highlights of larger areas.Algorithm
- Select the main algorithm 'HDR' will use - linked to the presets:

Context Size

Adjust the Context Size to target the size of detail you want to bring out.

- Changing context size starts a lot of processing so may take time. The larger the size the more processing.
- If you set this first to get it out of the way - adjusting the other values after this is quicker.
- Slider displays the percentage of the image size width and height.
- Controls the upper limit to the size of the detail being accentuated.
- A context size greater than 30 will likely take longer than the previous version of the HDR module.
- Selecting a sample area speeds up processing while you are finding the right 'Context Size' and 'Quality' settings. Then press 'All'.
- Avoid extreme values.
- Default is 50x50 pixels. Range is from 1 to 101.

Quality

- Adjusts the quality of the result. This has a significant effect on processing time.
- The Quality setting will affect the results in certain cases:
- You will see greatest effect where there are many small areas of dynamic range changes where there is violent, turbulent, detail.
- You will see little effect where there are a few large-scale areas where enhancement is required.
- Choose a lower quality to increase processing speed.
- Choose a higher quality to improve the accuracy of the result. Changing to 'High' will reduce processing speed.
- Approximate Relative Processing Time (Medium =100%): Low 98%, Medium 100%, High 220%. This may vary with environment/settings.
- This only affects the Gamma Correction part of the processing.
- Default is Medium. Range is: Low, Medium, High.

Highlights Detail Boost

- Accentuates detail in the highlights without affecting the brightness of larger contiguous areas.
- Controls the strength of detail recovery in the highlights.
- Increasing the value will bring out the detail.
- Default is 30% (Reveal), 0% (Tame, Equalize), 50% (Optimize). Range is from 0% to 100%.

Shadows Detail Boost

- Accentuates detail in the shadows without affecting the brightness of larger contiguous areas.
- Controls the strength of detail recovery in the shadows.
- Increasing the value will bring out the detail.
- Default is 30% (Reveal), 0% (Tame, Equalize), 50% (Optimize). Range is from 0% to 100%.

Signal Flow

- Specifies the signal source the algorithm works on.
- Tracked - Uses the Tracked signal which takes into account the noise evolution. Avoiding recovering 'detail' in low SNR areas. Recommended for most situations. This is only available if Tracking is on.
- Visual As-Is - Uses the stretched image without making allowances for SNR.
- Default is 'Tracked' (if Tracking is on).
- Range is: 'Visual As-Is', 'Tracked'.

Gamma Smoothen

- This parameter controls the smoothness of the transition between different locally-stretched areas.
- The default value is right for most situations.
- Increase this value if boundaries can be seen.
- Reduce this value (temporarily) to get an idea of the different areas HDR worked on.
- Default is 20.0.
- Range is from 1.0 to 40.0.

Gamma Highlight (Tame)

- Work on the larger contiguous bright areas to bring out the smaller scale detail within them.
- Redistributes dynamic range (by using local gamma correction) to move detail out of the bright areas into the mid-tones - which allows better scope for accentuating it.
- Default is 1.25 (Reveal, Equalize), 1.00 (Optimize), 3.00 (Tame). Range is from 1.00 to 5.00.

Gamma Shadow (Lift)

Works on the larger contiguous areas of shadow to bring out the smaller scale detail within them.

- Redistributes dynamic range (by using local gamma correction) to move detail out of the shadows into the mid-tones - which allows better scope for accentuating it.
- Default is 1.00 (Reveal, Optimize, Tame), 1.25 (Equalize).
- Range is from 1.00 to 5.00.

Strength

Sets how strong the effect should be overall - applies to the Reveal algorithm only.

- Default is 1.0.
- Range is from 1.0 to 3.9.

Small Detail Precision

Sets the degree of precision 'HDR' should use to identify small detail - affects speed of processing.

- Values of Max, Medium (Faster), Low (Fastest).
- Default is 'Max'.

Detail Size Range

Sets what detail size 'HDR' should concentrate on in order to bring out the most detail.

- Default is 20 pixels.
- Range is from 0 pixels to Max width.

- A small setting will separate out the details and make them stand on their own. This highlights detail but may seem harsh or flat.
- A large setting will make the details fit in with the rest of the image. This will appear soft and natural but some detail may seem obscured.

Channels

Sets whether 'HDR' should work on brightness detail, color detail or both.

- Values of 'Brightness Only', 'Color' Only' or 'Brightness and Color'.
- Default is 'Brightness Only'.

'Sharp': Multi-scale Noise-Aware Structural Detail Enhancement

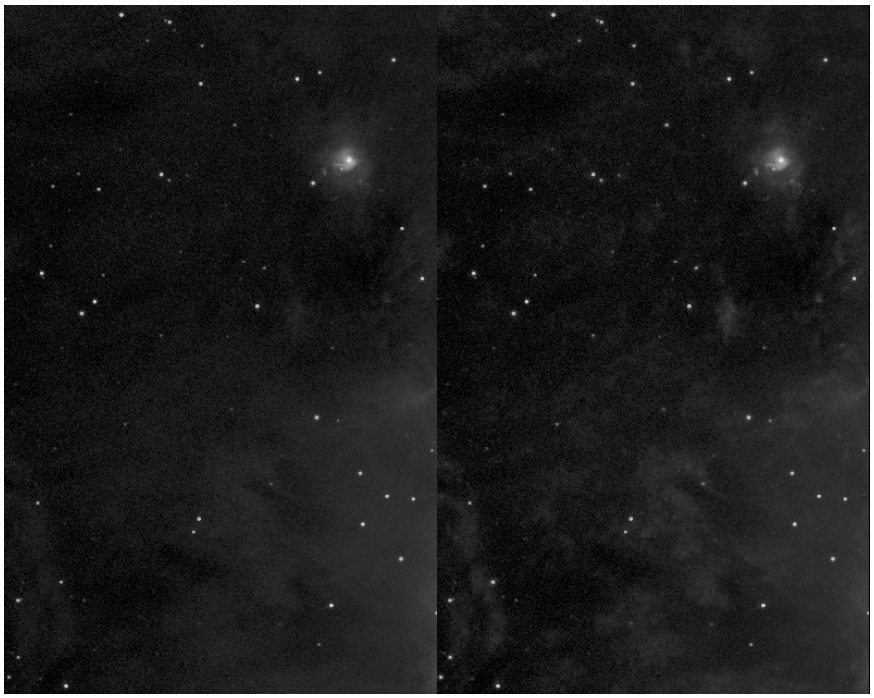

The 'Sharp' is able to dig out faint detail in the form of larger structures, entirely without exacerbating noise.

StarTools' Detail-aware Wavelet 'Sharpen'ing allows you to bring out faint structural detail in your images.

An important innovation over other, less sophisticated implementations, is that StarTools' Wavelet 'Sharpen'ing gives you precise control over how detail across different scales and SNR areas interact. This means that;

- 'Sharp' lets you control how detail is enhanced, based on the Signal-to-Noise Ratio (SNR) per-pixel in your image. This ability lets you dig out larger scale faint detail entirely without increasing perceived noise.
- 'Sharp' lets you to be the arbiter when two scales (bands) are competing to enhance detail in their band for the same pixel.

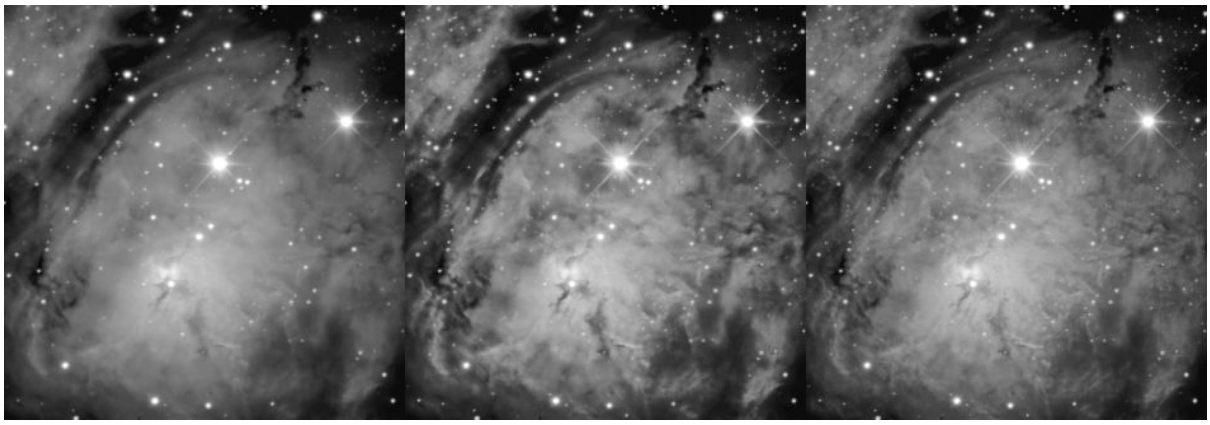

Left: source. Middle: bias towards larger scale structures. Right: bias towards smaller scale structures.

As with all modules in StarTools, the Wavelet Sharpening module will never allow you to clip your data, always yielding useful results, no matter how outrageous the values you choose, while availing of the Tracking feature's data mining. The latter makes sure that,

contrary to other implementations, only detail that has sufficient signal is emphasized, while noise grain propagation is kept to a minimum.

Usage

Using the 'Sharp' module, starts with specifying an upper limit of the size of the detail that should be accentuated via the 'Structure Size' parameter. You should only need to change this parameter if you wish very fine control over small details.

After pressing 'Next', a star mask should be created that protects stars (and their halos) from being accentuated.

An 'Amount' parameter governs the strength of the overall sharpening.

The 'Scale n' parameters allow you to control which detail sizes are getting enhanced. If you wish to keep small details from being enhanced set 'Scale 1' to 0%, similarly if you wish to keep the very largest structures from being enhanced, set 'Scale 5' to 0%.

The 'Dark/Light Enhance' parameter gives you control over whether only bright or dark (or both) detail should be introduced.

The two 'Size Bias' parameters controls the detail size that should prevail if two scales are 'fighting' over enhancing the same pixel. A higher value gives more priority to finer detail, whereas a lower value gives more priority to larger scale structures. It is this ability of the 'Sharp' modules, to dynamically switch between large and small detail enhancement that makes every combination of settings look coherent without 'overcooking' the image; the advantage is that if you try to make everything (every scale) stand out, nothing stands out. And this is precisely what the 'Sharp' module was designed for to avoid. Inherent to this approach is also the lack of ringing artifacts around sharp edges, even though the module does not employ a multi-scale median filter to try to circumvent this. This combines the benefits of the response of a pure Gaussian transform (such as precise band delineation and noise modeling) with ringing artifact-free detail enhancement.

Two version of the 'Size Bias' parameter exist; the 'High SNR Size Bias' parameter and the 'Low SNR Size Bias' parameter. The distinction lies in a further refinement of where and how detail enhancement should be applied. The 'High SNR Size Bias' parameter controls the size priority for areas with a high signal-to-noise ratio (good signal), whereas the 'Low SNR Size Bias' controls the size priority for areas with a low signal-to-noise ratio (poor signal).

When Tracking is on, the Tracking feature tends help the 'Sharp' module do a very precise job in making sure that noise is not exacerbated - you may find that the distinction is not needed for most datasets with signal of reasonable quality. However when Tracking is off, these parameters use local luminosity as a proxy for signal quality and the distinction between Low and High SNR will be much more important.

Finally, the 'Mask Fuzz' parameter increasingly smoothens the area over which the set mask goes from fully in effect, to not in effect.

Wrap Up

Description

To provide a range of tools which give the ability to sharpen the small to medium detail in the image. The 'Sharp' module uses wavelet sharpening techniques to enhance the detail in the image.

The module is scale aware - it splits the image into multiple layers of different scale elements - allowing the enhancing of detail of particular sizes only. The module is also noise-aware - allowing you to control its effect on areas depending on their SNR.

Only detail that has a high enough SNR is emphasized. Noise grain propagation is minimized.

The optional 'Intelligent Enhance' mode targets the sharpening where it will benefit most.

When to use

- After the final global stretch ('AutoDev' or Develop), 'Contrast' and 'HDR' - if used.
- Before the 'Decon' Module.
- Use only once.

What result to look for

- Look for structures becoming clearer - less blurred.
- Watch for stars bloating. If this happens, improve the star mask to counter this effect.
- Watch for noise being enhanced. If this happens at all scales then back off the Amount setting. If it is enhanced at a particular scale then back off the relevant Scale setting instead.

After Use

- Use 'Decon' module if needed.

Description of Controls

Screen 1: Structure Size

Specifies the largest (Scale 5) size the wavelet sharpening should cover:

- Large - elements of approximately 100-120 pixels downwards.
- Medium - elements of approximately 40-50 pixels downwards.
- Small - elements of approximately 10-15 pixels downwards.
- Default is Large.
- Set this to the size of the largest element you want to sharpen (e.g. spiral arms, dust lanes).

Screen 2:

Presets

- DSO - Optimizes the Settings for DSOs
- DSO Light - As DSO - but allows full local brightening to enhance contrast. Sets Dark/Light Enhance to 0%/100%
- DSO Dark - As DSO - but allows full local darkening to enhance contrast. Sets Dark/Light Enhance to 100%/0%
- Planetary - Optimizes the Settings for Planetary and Moon.

Mask

- For general instructions on using mask see Mask.
- Select what parts of the image to sharpen - The inverse star mask used by 'Decon' - which excludes the stars - is ideal here.

Auto Mask

- Create an inverse star mask. Auto Mask - Auto Generate Mask (or Auto Generate conservative Mask)

Scale

- Sets the amount of sharpening at each of 5 different scales ranges.
- Default is 100%.
- Range is 0% to 100%.
- Reduce scale from 100% at a particular scale if noise is being enhanced at that scale. If changed - this is usually only necessary at scales 1-2.
- Scales do not have absolute limits to the range - its is more like a particular scale brings detail of a certain size into focus - and that other detail is out of focus to varying degrees depending on its size.

- The following are broad guidelines:
  - The largest scale (Scale 5) is set by the Structure Size parameter.
  - The smallest size (Scale 1) is always around one pixel.
  - The intervening scale sizes increase exponentially.
- Scale Descriptions:
  - Scale 1 - Controls fine feature sharpening - This scale covers the smallest features such as single pixel detail.
  - Scale 2 - Controls the amount of medium to small feature sharpening.
  - Scale 3 - Controls the amount of medium feature sharpening.
  - Scale 4 - Controls the amount of large to medium feature sharpening.
  - Scale 5 - Controls the amount of large feature sharpening.

Amount

Adjusts the strength of sharpening across all scales

- Default is 100%.
- Range is 0% to 1,000%.
- Increase to make structures clearer. Decrease if noise is being enhanced too much at every scale.

Dark/Light Enhance

Specifies how the module can improve contrast - by specifying whether it is by darkening the surrounding area, lightening the detail, or a mixture of both.

- Increasing this value - towards 0%/100% - means the module will be able to lighten areas rather than darken them.
- Decreasing this value - towards 100%/0% - means the module will be able to darken areas rather than lighten them.
- Use low values (towards 0%/100%) to enhance faint nebulosity without darkening the surrounding area.
- Default is 50%/50% in DSO and Planetary presets
- Default is 0%/100% in DSO Light preset
- Default is 100%/0% in DSO Dark preset

High SNR Size Bias

Specifies how the module resolves any conflicts between scales in high signal-to-noise ratio (SNR) areas. For example, it may be that sharpening in a larger scale causes contrast reduction in a smaller scale.

- The larger the value, the more important is the small-scale detail compared to the larger scale.
- Normally the emphasis would be on the small-scale detail.
- Default is 85%.
- Range is 0% to 100%.

Low SNR Size Bias

Specifies how the module resolves any conflicts between scales in low signal-to-noise ratio (SNR) areas.

- The larger the value, the more important is the small-scale detail compared to the larger scale.
- For DSO's, in low SNR areas is more likely that only large scale detail can be recovered - so this should be low (0%).
- For Planetary images, the small-scale detail should be protected so use a high value (around 85%).
- Default is 0% for DSO presets, 85% for Planetary preset.
- Range is 0% to 100%.

Mask Fuzz:

If a mask is used, Mask Fuzz controls the blending of the transition between masked and non-masked parts of the image.

- Active if a mask is used to selectively sharpen the image.
- Using this control will ensure smooth transitions between sharpened and unsharpened parts of the image.
- Default is 4.0 pixels.
- Range 1.0 to 40.0 pixels.

The 'Sharp' module uses Wavelet (aka Laplacian) Sharpening techniques to enhance the detail in the image. See How Unsharp Masking and Laplacian Sharpening Work in Background Information section

Scale and regions affected

- The scale is based on the number of pixels - so if you use the same settings on an un-binned image and the same image 50% binned they will have different scales affected.
- For each scale increment the size increases exponentially, based on Structure Size. For Small the next scale size is 2x the current one, For Medium it is 2.75x and for Large it is 3.5x.
- If you want to see what size a particular scale affects - set that Scale to 100%, the Amount to 1000%, and all other scales to 0%. Do Before/After to see the scale affected.

SVDecon: Detail Recovery through Spatially Variant Distortion Correction

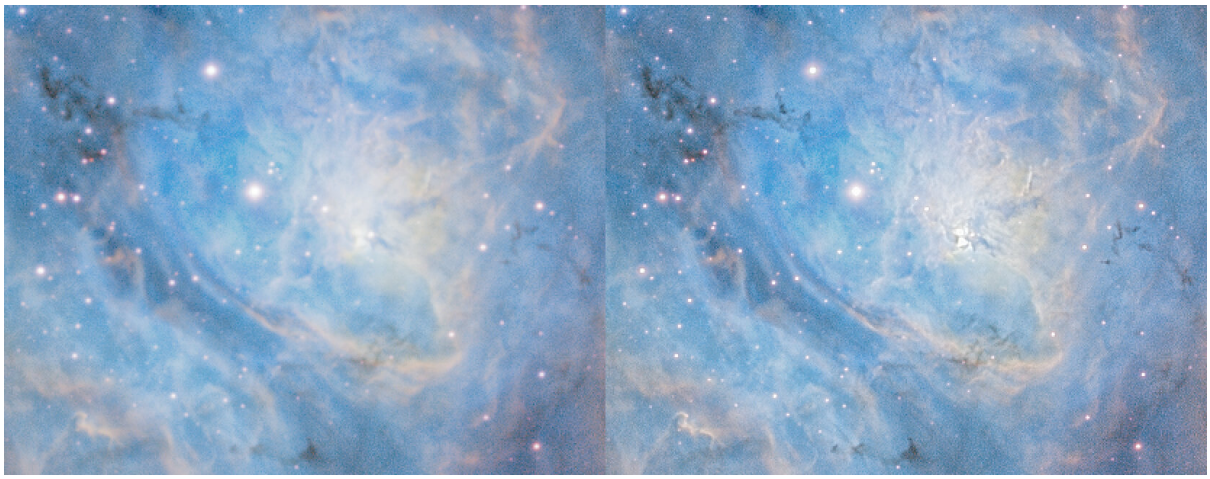

Unique amongst its peers, the GPU-accelerated SVDecon module in StarTools, is robust in the face of severe noise, singularities (such as over-exposing star cores), extreme non-linear processing and local detail enhancement, and variable Point Spread Functions. Left: original at 200% zoom. Right: deconvolved with Spatially Variant PSF Deconvolution at 200% zoom.

StarTools is the first and only software for astrophotography to implement fully generalised Spatially Variant PSF deconvolution (aka "anisotropic" or "adaptive kernel" deconvolution). The fully GPU accelerated solution is robust in the face of even severe noise, meaning it can deployed to restore detail in almost real-time in almost every dataset.

Even the best optical systems will suffer from minute differences in Point Spread Functions (aka "blur functions") across the image. Therefore, a generalised deconvolution solution that can take these changing distortions into account, has been one of the holy grails of astronomical image processing.

Innovations at a glance

The SVDecon module incorporates a series of unique innovations that sets it apart from all other legacy implementations, as found in other software;

- It corrects for different distortions at different locations in the data and not just one distortion for the entire dataset
- It preferably operates on highly processed and stretched data (provided StarTools' signal evolution Tracking is engaged)
- It performs intra-iteration resampling of PSFs
- It is almost always able to provide meaningful improvements, even when dealing with marginal datasets
- It is robust in the presence of natural singularities in the dataset (e.g. over-exposed star cores)
- Depending on your system, previews complete in near-real-time
- Any development of noise grain is tracked and marked for removal/mitigation during final noise reduction
- Smart caching allows faster tweaking of some parameters (such as de-ringing) without needing re-doing full deconvolution

Usage

It is important to understand two things about deconvolution:

- Deconvolution is 'an ill-posed problem', due to the presence of noise in every dataset. This means that there is no one perfect solution, but rather a range of approximations to the 'perfect' solution.
- Deconvolution should not be confused with sharpening; deconvolution should be seen as a means to restore a compromised (distorted by atmospheric turbulence and diffraction by the optics) dataset. It is not meant as an acuity enhancing process.

Deconvolution with a spatially variant Point Spread Function, adds to the complexity of basic deconvolution by requiring a model that accurately describes how the Point Spread Function changes across the image.

Understanding the former two important points will make clear why the various parameters exist in this module, and what is being achieved by the module.

Even this heavily distorted dataset with changing trailing direction from corner-to-corner can be improved.

Modes of operation

The SVDecon module can operate in several implicit modes, depending on how many star samples - if any - are provided;

When no star samples are provided, the SVDecon module will operate in a similar way to the pre-1.7 deconvolution modules.: a selection of synthetic models are available to model one specific atmospheric or optical distortion that is true for the entire image.

When one star sample is provided, the SVDecon module will operate in a way similar to the 1.7 module (though somewhat more effectively); a single sample provides the atmospheric distortion model for the entire image, while an optional synthetic optics model provides further refinement.

When multiple star samples are provided, the SVDecon module will operate in the most advanced way. Multiple samples provide a distortion model that varies per location in the image. An optional optical synthetic model may be used for further refinement, though is usually best turned off.

The latter mode of operation is usually the preferred and recommended way of using the module, and takes full advantage of the spatially variant PSF modelling and correction capabilities.

The module automatically grays out parameters that are not being used, and may also change (zero-out or disable) some parameters in line with the different modes as they are accessed.

When the subject is lunar or planetary in nature, no star samples are typically available. The "Planetary/Lunar" preset button configures the module for optimal use in these situations.

Finally, details of the mode being used, are reflected in the message below the image window.

Apodization Mask

The SVDecon module requires a mask that marks the boundaries of stellar profiles. Pixels that fall inside the masked areas (designated "green" in the mask editor), are used during model construction. Pixels that fall outside the masked area are disregarded during local model construction.

It is highly recommended to to include as much of a star's stellar profile in the mask as possible. Failure to do so may lead to increased ringing artifacts around deconvolved stars. Sometimes a simple manual "Grow" operation in the mask editor suffices to include the more of the stellar profiles.

Compared to most other deconvolution implementations, the SVDecon module is robust in the face of singularities (for example over-exposing star cores). In fact, it is able to coalesce such singularities further. As such, the mask is no longer primarily used for designating singularities in the image, like it was in versions of StarTools before 1.8.

The mask does however double as a rough guide for the de-ringing algorithm, indicating areas where of ringing may develop. Clearing the mask (all pixels off/not green in the mask editor) is generally recommended for non-stellar objects, including lunar, planetary or solar data. As a courtesy, this clearing is performed automatically when selecting the Planetary/Lunar preset.

Point Spread Function (PSF)

Left: original, middle: deconvolved image with appropriate settings, right: deconvolved image with ringing artifacts due to an inappropriate (too high) choice for the Radius parameter.

A Deconvolution algorithm's task, is to reverse the blur caused by the atmosphere and optics. Stars, for example, are so far away that they should really render as single-pixel point lights. However in most images, stellar profiles of non-overexposing stars show the

point light spread out across neighbouring pixels, yielding a brighter core surrounded by light tapering off. Further diffraction may be caused by spider vanes and/or other obstructions in the Optical Tube Array, for example yielding diffraction spikes. Even the mere act of imaging through a circular opening (which is obviously unavoidable) causes diffraction and thus "blurring" of the incoming light.

The point light's energy is scattered around its actual location, yielding the blur. The way a point light is blurred like this, is also called a Point Spread Function (PSF). Of course, all light in your image is spread according to a Point Spread Function (PSF), not just the stars. Deconvolution is all about modeling this PSF, then finding and applying its reverse to the best of our abilities.

Introducing Spatial Variance

Traditional deconvolution, as found in all other applications, assumes the Point Spread Function is the same across the image, in order to reduce computational and analytical complexity. However, in real-world applications the Point Spread Function will vary for (X, Y) location in a dataset. These differences may be large or small, however always noticeable and present; no optical system is perfect. Ergo, in a real-world scenario, a Point Spread Function that perfectly describes the distortion in one area of the dataset, is typically incorrect for another area.

Traditionally, the "solution" to this problem has been to find a single, best-compromise PSF that works "well enough" for the entire image. This is necessarily coupled with reducing the amount of deconvolution possible before artifacts start to appear (due to the PSF not being accurate for all areas in the dataset).

Being able to use a unique PSF for every (X, Y) location in the image solves aforementioned problems, allowing for superior recovery of detail without being limited by artifacts as quickly.

Synthetic vs sampled PSFs

The SVDecon module, makes a distinction between two types of Point Spread Functions; synthetic and sampled Point Spread Functions. Depending on the implicit mode the module operates in, synthetic, sampled, or both synthetic and sampled PSFs are used.

When no samples are provided (for example on first launch of the SVDecon module), the module will fall back on a purely synthetic model for the PSF. As mentioned before, this mode uses the single PSF for the entire image. As such the module is not operating in its spatially variant mode, but rather behaves like a traditional, single-PSF model, deconvolution algorithm as found in all other software.

A number of parameters can be controlled separately for the synthetic and sampled Point Spread Function deconvolution stages.

Synthetic PSFs

Synthetic PSF models

Atmospheric or lens-related blur is more easily modelled, as its behavior and effects on long exposure photography has been well studied over the decades. 5 subtly different models are available for selection via the 'Primary Point Spread Function' parameter:

- 'Gaussian' uses a Gaussian distribution to model atmospheric blurring.
- 'Circle of Confusion' models the way light rays from a lens are unable to come to a perfect focus when imaging a point source (aka the 'Circle of Confusion'). This distribution is suitable for images taken outside of Earth's atmosphere or images where Earth's atmosphere did otherwise not distort the image.
- 'Moffat Beta=4.765 (Trujillo)' uses a Moffat distribution with a Beta factor of 4.765. Trujillo et al (2001) propose in their paper that this value is the best fit for prevailing Atmospheric turbulence theory.

- 'Moffat Beta=3.0 (Saglia, FALT)' uses Moffat distribution with a Beta factor of 3.0, which is a rough average of the values tested by Saglia et al (1993). The value of ~3.0 also corresponds with the findings Bendinelli et al (1988) and was implemented as the default in the FALT software at ESO, as a result of studying the Mayall II cluster.
- 'Moffat Beta=2.5 (IRAF)' uses a Moffat distribution with a Beta factor of 2.5, as implemented in the IRAF software suite by the United States National Optical Astronomy Observatory.

Only the 'Circle of Confusion' model is available for further refinement when samples are available. This allows the user to further refine the sample-corrected dataset if desired, assuming any remaining error is the result of 'Circle of Confusion' issues with all other issues corrected for as much as possible.

Even this noisy and heavily drizzled Hubble dataset can be corrected by the StarTools' 'Decon' module at its native, drizzled resolution. Left: not deconvolved, middle: deconvolved, right: deconvolved and noise grain equalized.

The PSF radius input for the chosen synthetic model, is controlled by the 'Synthetic PSF Radius' parameter. This parameter corresponds to the approximate the area over which the light was spread; reversing a larger 'blur' (for example in a narrow field dataset) will require a larger radius than a smaller 'blur' (for example in a wide field dataset).

The 'Synthetic Iterations' parameter specifies the amount of iterations the deconvolution algorithm will go through, reversing the type of synthetic 'blur' specified by the 'Synthetic PSF Model'. Increasing this parameter will make the effect more pronounced, yielding better results up until a point where noise gradually starts to increase. Find the best trade-off in terms of noise increase (if any) and recovered detail, bearing in mind that StarTools signal evolution Tracking will meticulously track noise propagation and can snuff out a large portion of it during the Denoise stage when you switch Tracking off. A higher number of iterations will make rendering times take longer - you may wish to use a smaller preview in this case.

Sampled PSFs

Sampled PSF models

Ideally, rather than relying on a single synthetic PSF, multiple Point Spread Functions are provided by means of carefully selected samples; isolated stars on an even background that do not over expose, nor are too dim. Ideally, these samples are provided for all areas across the image, so that the module can analyse and model how the PSF changes for each pixel in the image.

Recommended workflow

As opposed to all other implementations of deconvolution in other software, the usage of the SVDecon module is generally recommended towards the end of your luminance processing workflow. That is, ideally, you will have already carried out the bulk of your stretching and detail enhancement before launching the SVDecon module. The reason for this, is that the SVDecon module makes extensive use of knowledge that indicates how you processed your data prior to invoking it, and how detail evolved and changed during your processing. This knowledge specifically feeds into the way noise and artifacts are detected and suppressed during the regularisation stage for each iteration.

For most datasets, superior results are achieved by using the module in Spatially Variant mode, e.g by providing multiple star samples. In cases where providing star samples is too difficult or time consuming, the default synthetic model will still very good results however.

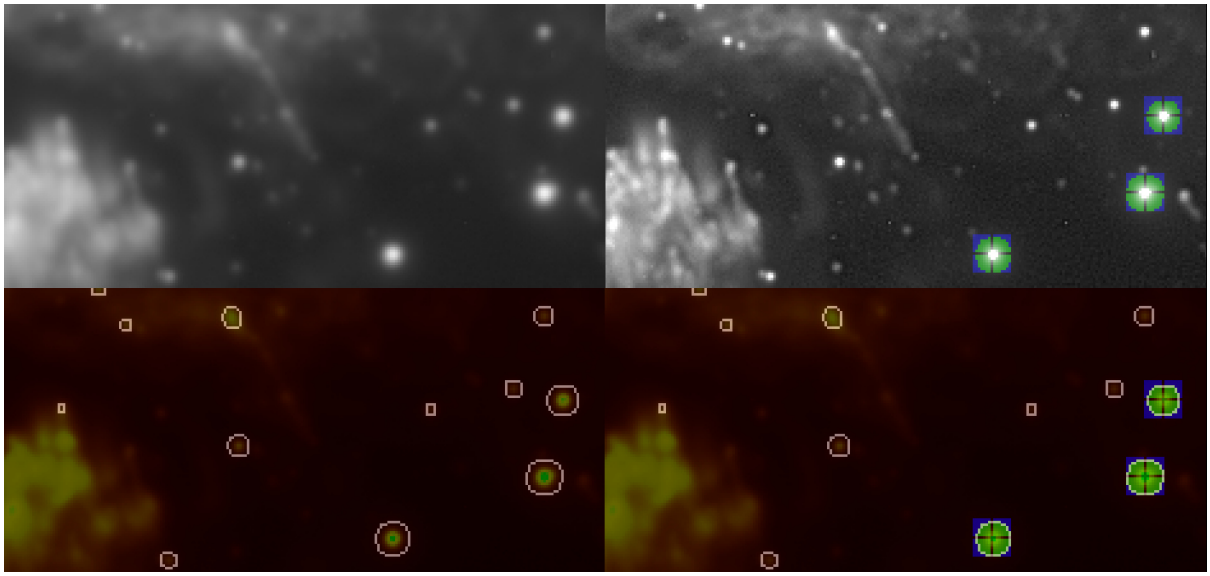

The SVDecon module is ideally operated by selecting samples of good quality. Top left; original image. Top right; the resulting deconvolved image by selecting the samples indicated. Bottom left; "Sample Quality" view without samples selected. Bottom right; "Sample Quality" view with three good quality samples selected. The blue bounding boxes should ideally fit the entire green "blobs" (signifiying the apodization mask for each sample).

Selecting samples for Spatially Variant deconvolution

To provide the module with PSF samples, a special view is available to help the user identify good quality star samples. This view is accessed by clicking the 'Sampling' button in the top right corner.

A convenient rendering of the image is shown, in which:

- Candidate stars are delineated by an outline.
- Red pixels show low quality areas
- Yellow pixels show borderline usable areas.
- Green pixels show high quality areas.

Ideally, you should endeavour to find stars that have a green inner core without any red pixels at their centre. If you cannot find such stars and you need samples in a specific area

you may choose samples that have a yellow core instead. As a rule of thumb, providing samples in all areas of the image takes precedence over the quality of the samples.

You should avoid:

- Stars that sit on top of nebulosity or other detail.
- Objects that are not stars (for example distant galaxies)
- Stars that are close to other stars
- Stars that appear markedly different in shape compared to other stars nearby
- Stars whose outline appear non-oval or concave or markedly different to the outlines of other stars nearby

Star samples can be made visible on the regular view (e.g. the view with the before/after deconvolved result) by holding the left mouse button. Star samples will also be visible outside any preview area, this also doubles as a reminder that any selected PSF Resampling algorithm will not resample those stars (see 'PSF resampling mode'). You may also quickly de-select stars via the regular before/after view by clicking on a star that has a sample over it that you wish to remove.

The Sampled Area

The immediate area of a sampled star is indicated by a blue square. This area is the 'Sampled Area'. A sampled area should contain one star sample only; e.g. you should particularly avoid parts of other stars in the blue square surrounding a sample. This size of the blue square is determined by the 'Sampled Area' parameter. The 'Sampled Area' parameter should be set in such a way that the single samples fall well within the blue area's confines and are not 'cut-off' by the blue square's boundaries.

Star sample outlines and apodization mask

The star sample outlines are constructed using the apodization mask that is generated. You may touch up this mask to avoid low-quality stars being included in the blue square 'Sampled Area', if that helps to better sample a nearby higher quality star.

Number of samples and location of samples

Ideally samples are specified in all areas of the image in equal numbers. The module will work with anywhere from one to hundreds of samples. The amount of samples you should provide is largely dependent on how severe the distortions are in the image, and the quality of the samples (see the aforementioned 'You should avoid' items).

When clicking a sample, the indicated centre of a sample will not necessarily be the pixel You clicked nor brightest pixel. Instead, the indicated centre is the "luminance centroid". It is the weighted (by brightness) mean of all pixels in the sample. This is so that, for example, samples of stars that are deformed or heavily defocused (where their centre is less bright than their surroundings) are still captured correctly.

Heavily distorted PSFs

For images with heavily distorted PSFs that are highly variant (for example due to field rotation, tracking error, field curvature, coma, camera mounting issue, or some other acquisition issue that has severly deformed stars in an anisotropic way), the 'Spatial Error' parameter may need to be increased, with the 'Sampled Iterations' increased in tandem. The 'Spatial Error' parameter relaxes locality constraints on the recovered detail, and increasing this parameter, allows the algorithm to reconstruct point lights from pixels that are much less co-located than would normally be the case. Deconvolution is not a 100% cure for such issues, and its corrective effect is limited by what the data can bear without artifacts (due to noise) becoming a limiting factor. Under such challenging conditions, improvement should be regarded in the context of improved detail, rather than perfectly point or circle-like stellar profiles.

Detail should snap into focus, stars should coalesce into point lights and halos around over-exposing stars should be diminished.

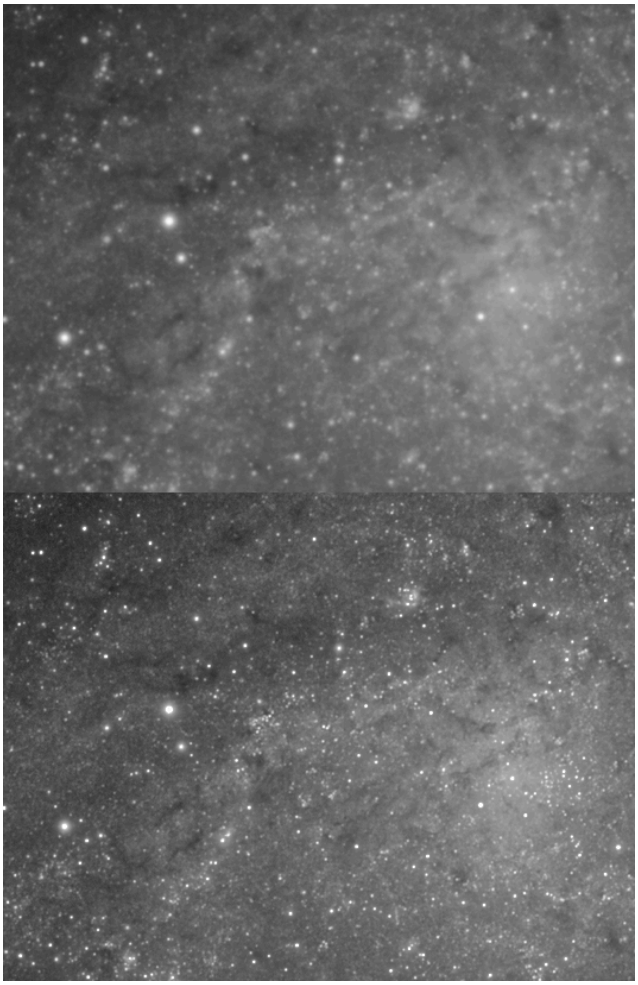

While stars may definitely become more pin-point and/or 'rounder', particularly areas that are (or are close to) over-exposing, such as very bright stars, may not contain enough data for reconstruction due to clipping or non-linearity issues. Binning the resulting image slightly may somewhat help hide issues in the stellar profiles. Alternatively, the Repair module may help correcting these stars.

PSF Resampling mode

The SVDecon module is innovative in many ways, and one such innovation is its ability to re-sample the stars as they are being deconvolved. This feedback tends to reduce the development of ringing artifacts and can improve results further.

Three 'PSF Resampling' modes are available:

- None; no resampling and model reconstruction occurs during deconvolution - the samples are used as-is.
- Intra-Iteration; all samples are resampled at their original locations for each iteration
- Intra-Iteration + Centroid Tracking; all samples are resampled after their locations have first been re-determined.

Dynamic Range Extension

The 'Dynamic Range Extension' parameter provides any reconstructed highlights with 'room' to show their detail, rather than clipping it against the white point of the input image. Use this parameter if signficant new detail is created that requires more dynamic range to be fully appreciated. Lunar datasets can often benefit from an extended dynamic range allocation.

Planetary, solar and lunar datasets

A preset for lunar, planetary, solar use quickly configures the module for lunar, planetary and solar purposes: it clears the apodization mask (no star sampling possible/needed) and dials in a much higher amount of iterations. It also dials in a large synthetic PSF radius more suitable to reverse atmospheric turbulence-induced blur for high magnification datasets. You will likely want to increase the amount of iterations further, as well as adjust the PSF radius to better model the specific seeing conditions.

Just using synthetic PSF modelling, the SVDecon module is just as effective on lunar, planetary and solar datasets, as it is on deep space datasets.

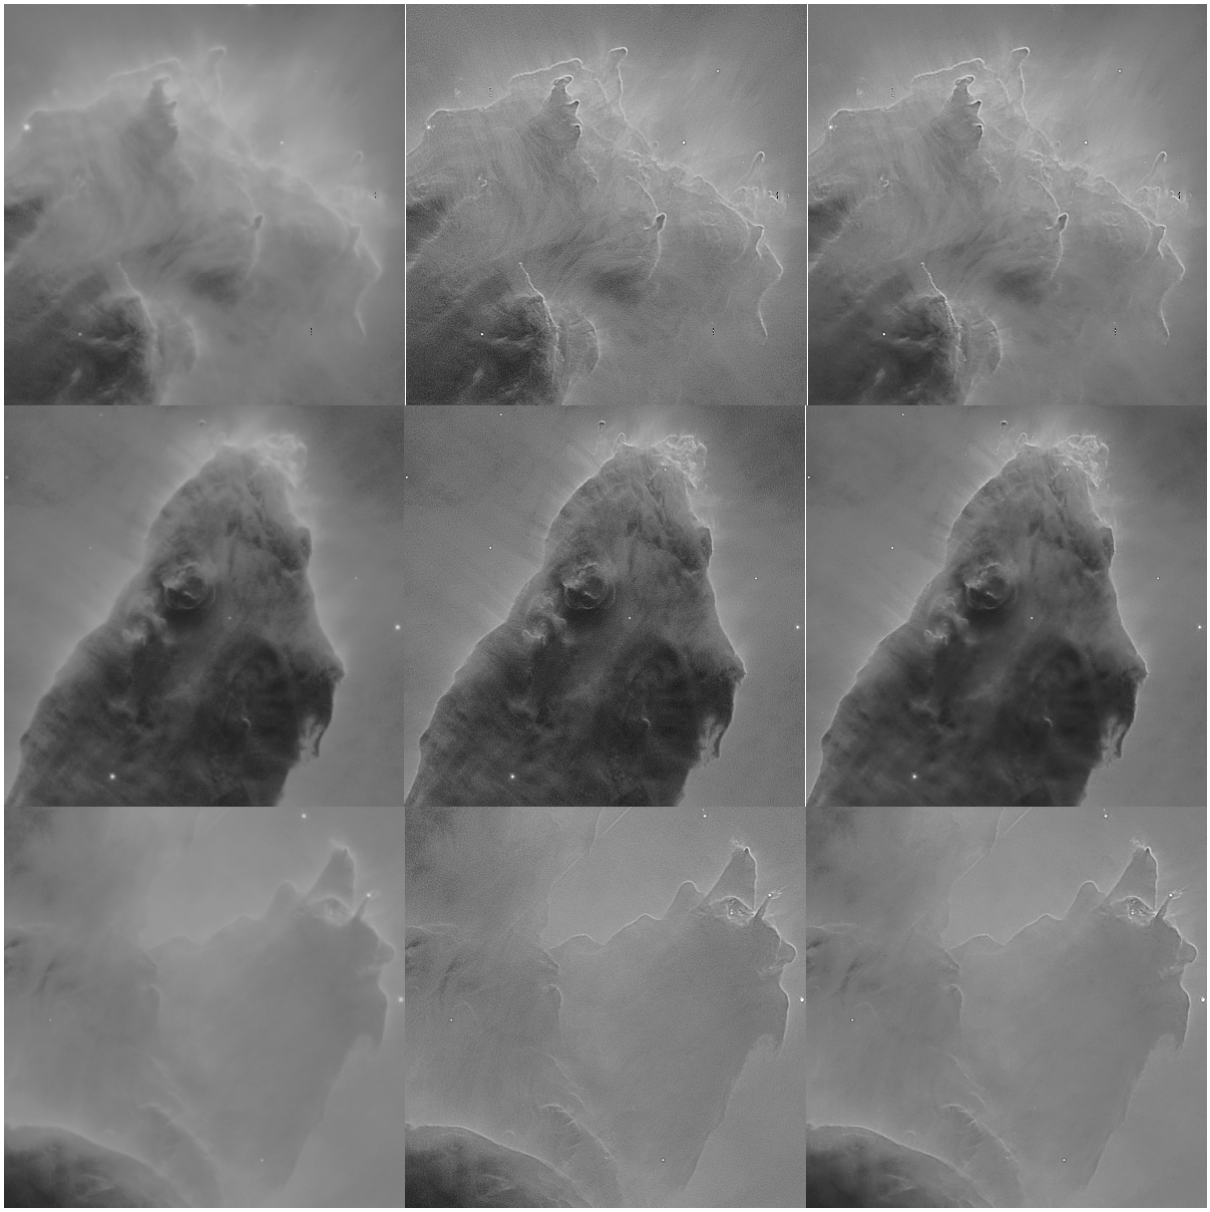

SVDecon can greatly improve the clarity of image data from even space telescopes and probes such as the Hubble Space Telescope. Right column: original data, Center column: SV deconvolved data, Left column: default noise mitigation during final signal Tracking Denoise stage, demonstrating precise, autonomous deconvolution-induced grain and artifact tracking. Evaluating the result

A considerable amount of research and development has gone into optimisation of the algorithm; an important part of image processing is getting accurate feedback as soon as possible on decisions made, samples set, and parameters tweaked.

As a result, it is possible to evaluate the result of including and excluding samples in near-real-time; you do not need to wait minutes for the algorithm to complete. This is particularly the case when a smaller preview area is selected.

Please note, however, that the 'PSF Resampling' feature is only carried out on any samples that exist in the preview area. As a result, when a 'PSF Resampling' mode is selected, previews may differ somewhat from the full image. To achieve a preview for an area when a 'PSF Resampling' mode is selected, try to include as many samples in the preview area as possible when defining the preview area's bounding box.

With the forementioned caveat with regards to resampling in mind however, any samples that fall outside the preview are still used for construction of the local PSF models for pixels inside the preview. In other words, the results in the preview should be near-identical to deconvolution of the full image, unless a specific 'PSF Resampling' mode is used.

Noise and artifact propagation

While it is best to avoid overly aggressive settings that exacerbate noise grain (for example by specifying a too large number of iterations), a significant portion of such grain will be still be very effectively addressed during the final noise reduction stage; StarTools' Tracking engine will have pin-pointed the noise grain and its severity and should be able to significantly reduce its prevalence during final noise reduction (e.g. when switching Tracking off).

Ringing artifacts and/or singularity-related artifacts are harder to address and their development are best avoided in the first place by choosing appropriate settings. As a last resort, the 'Deringing' and 'Deringing Fuzz' parameters can be used to help mitigate their prevalence.

Recovering PSF samples from the log

Any samples you set, are stored in the StarTools.log file and can be restored using the 'LoadPSFs' button.

In the StarTools.log file, you should find entries like these:

PSF samples used (8 PSF sample locations, BASE64 encoded)

VFMAAAgAOAQMA/oDEQHaAoEAIwNeAOQAUwDUAY8AbAI5AdMBMQGkAFAB

If you wish to restore the samples used, put the BASE64 string (starting with VFM... in the example) in a text file. Simply load the the file using the 'LoadPSFs' button.

Wrap Up

Description

- The SVDecon module tries to reverse the effects that atmospheric turbulence and, optionally, the optical train, has on the data. It allows recovery of detail in seeing-limited data sets that were affected by atmospheric turbulence and diffraction. It can model the variation in the effects across the image.
- The SVDecon module allows the modification of the deconvolution algorithm across the image. By deconvolving selected sample stars across the image a mapping of how the atmospheric and optical distortion - and hence deconvolution algorithm - varies across the image.
- There are three modes of operation:
    - When no star samples are selected - uses one synthetic atmospheric distortion model for the whole image - As pre v1.7 StarTools did. A number of synthetic models are available to choose from.
    - When one star sample is selected - uses that sample to provide the atmospheric distortion model for the whole image - with an option to model the optical distortion by choosing a synthetic optics model - Similar to StarTools v1.7
    - When multiple star samples are selected - uses the samples to define how the distortion model varies across the image. There is an option to model the optical distortion by choosing a synthetic optics model.

When to use

- The 'SVDecon' module should be used just after the final global stretch and local dynamic stretch optimization ('Contrast', 'HDR', 'Sharp').
- SVDecon will be able to achieve better results the closer you get to a final image - since it has better information from tracking.

- 'SVDecon' will create its own Apodisation mask to support the process of star sampling.
- If there is no oversampling, or if there is a lot of noise, the benefit of this module will be limited.
- Use only once.

What result to look for

- Elements should appear more focussed as the blurring effect of atmospheric turbulence is compensated for.
- Edges should look more distinct and without the exaggerated coalescing caused by increasing the radius a lot.
- If there are ringing artifacts around the stars it indicates the Radius parameter is too high.
- Noise may be removed by 'Denoise' to a large part later on. If noise is becoming prevalent, use 'Deringing' and 'Deringing Fuzz' to mitigate.
- An increase in blurring in the foreground in the 'After' image indicates the Regularization setting is too high.

Ways of getting better results

- The 'SVDecon' module works best when there is little noise - 'Bin' your oversampled data to improve the SNR if needed - but consider leaving some degree of oversampling to allow deconvolution to bring out finer detail. See the 'Bin' module notes for a discussion of the issues relating to 'Bin' vs. deconvolution.
- Data that is not oversampled is not a candidate for deconvolution if the aim is to reverse seeing-related issues.
- Select star samples evenly across the image
- Should only be medium to low quality star samples be available in one region, increasing the amount in that region will help.
- Use only one of the synthetic models for lunar, planetary and solar images.

After Use

- Optionally use the 'Noise' preset of 'Flux' and/or the 'SuperStructure' module.
- Use the 'Color' module.

Description of Controls:

Mask

Description of Controls:

Mask:

- Mask usage is changed compared to previous versions and other modules!
- In this case the mask is used to enable selection of sample stars.
- If no mask is set you are asked if you want to create the apodization mask. An apodization mask is needed if star samples are to be taken.
- The apodization mask identifies the light sources that will be processed.
- The apodization mask is also used to identify the potential samples in the sample view.
- You can touch up the mask in order to remove unwanted low quality stars in the blue sampled area of a sample star.
- The apodization mask identifies the boundaries of the pixels associated with a star. Only the pixels inside the boundaries are used by the module.
- It is important to include as much of the stellar profile as possible. If too few pixels are chosen there may be 'ringing' around deconvolved stars

All

- If you selected a preview area - this will apply the deconvolution to the whole image.

Sampled PSF Area

- When using sampling - specifies the area around the stars centre that is sampled.
- The area should contain only one sampled star.
- All the pixels that belong to the sample star should just fit into the area - with not much background.
- The sampled area is shown in blue - grow or shrink the area to ensure the samples just fit around the star(s).
- Default is 15x15 pixel area, Range is 7x7 to 25x25 in 2 pixel increments.

PSF Resampling

- Controls if and how the Point Spread Functions are re-sampled.
- Resampling tends to reduce the development of ringing artifacts and can also improve results.
- There are 3 modes:
  - None - No resampling or reconstruction - samples are used as they are.
  - Intra-Iteration - All samples are resampled at their original location
  - Intra-Iteration + Centroid Tracking Linear - All samples are resampled after their locations have been recalculated
  - Default is 'None'

Plnt/Lnr

- Enables Planetary/Lunar Mode.
- clears apodization mask and disables sampling mode
- sets 'Synthetic Iterations' to 50
- uses higher 'Synthetic PSF Radius'
- disables deringing

Apod Mask

- opens the menu for apodization mask creation

Sampled Iterations

Sets the number of iterations the sample based deconvolution algorithm goes through.

- Default is 10
- Range is 1 to 199
- Increase this value incrementally. In deep sky mode, 10 iterations (default) will provide a mild effect, 20 a strong deconvolution.
- With PSF Resampling on, result will be more precise and 20% higher values will still work without looking over-processed
- The algorithm does not converge to an optimum solution like V1.7. High values will create an ever higher processed look.

Sampled Area

Sets the size of the star sampling area, indicated by a blue square.

- Default is 15x15
- Options are: 7x7; 9x9; 11x11; 13x13; 15x15; 17x17; 19x19; 21x21; 23x23; 25x25

Spatial Error

Controls the locality constraints of the algorithm.

- Higher values allow less co-located pixels to be used for point light reconstruction.
- Increase this value for strongly distoreted point lights (anastigmatism, coma)
- Reduce this value should bright artifacts show up when using samples in a busy star field.

- Default is 1.00
- Range is 0.00 to 2.00

Dyn. Range Extension

Allocates more dynamic range to reconstructed highlights to prevent white point clipping.

- Default is 1.00
- Range is 1.00 to 3.00

Linearity Cutoff

Specifies the point in the dynamic range above which the sensor becomes non-linear.

- Default is 85%
- Range is 0% to 100%

Synthetic Iterations

Sets the number of iterations the synthetic deconvolution algorithm goes through.

- Default is 0 (Deep Sky), 50 (Planetary/Lunar)
- Range is 0 to 500
- Increase this value incrementally. In Planetary/Lunar mode, 50 iterations (default) will provide a mild effect.
- Synthetic PSF may help to improve the sample based deconvolution if not enough good sample stars can be selected. Use a synthetic Iteration or two to iron out strong undulation in sample based PSF

Synthetic PSF Radius

The size of the blur that 'Decon' will try and remove.

- This value can be increased until ringing starts to occur on small non-overexposed stars - then back off a little.
- Related to the seeing - Seeing-induced blur is normally 3-4.5 arc-seconds. The camera/lens combination gives a resolution between 1 and 5 arcsec/pixel depending on the equipment combination.
- Adjust the Radius until just before the smaller stars start showing signs of ringing.
- If a Secondary PSF is defined, its radius will be derived from this radius.
- Default is 1.5 pixels.
- Range is 1.0 to 20.0 pixels.

Synthetic PSF Model

Defines which model of atmospheric turbulence blurring is used in the reversal process:

- 'Gaussian' uses a Gaussian distribution to model atmospheric blurring. This model is fast to calculate and was the default model in StarTools prior to version 1.6.
- 'Circle of Confusion (Optics Only)' models basic focusing assuming no atmosphere (e.g. in space).
- 'Moffat Beta=4.765 (Trujillo)' uses a Moffat distribution with Beta factor 4.765. Recommended by Trujillo et al. (2001).
- 'Moffat Beta=3.0 (Saglia, FALT)' uses Moffat distribution with a Beta factor of 3.0, as implemented in the FALT software at ESO.
- 'Moffat Beta=2.5 (IRAF)' uses a Moffat distribution with a Beta factor of 2.5, as implemented in the IRAF software by the USNOAO.
- Default is Circle of Confusion (Optics Only)
- Allows Decon' module to calculate a total PSF for both the atmosphere using sample stars as a reference and the optical train using synthetic PSF.
- With sampling disabled, it may create the athmospheric model similar to primary PSF in V1.7, e.g. in Planetary/Lunar mode.

Deringing

Defines the aggressiveness of the de-ringing filter:

- Higher values will undo more artifacts - but may also undo some of the darker detail enhancements from 'Decon'.
- Reducing risks introducing more artifacts - which may cause worse effects in other modules.
- Lunar, Planetary and Solar images need a less aggressive deringing strategy and so this can be set to Off (0%)
- Default is 50% (DSO), 0% (Plnt/Lnr)
- Range is Off (0) to 100%.

Deringing Fuzz

Controls the smoothness of the transition between the de-ringed areas and the surrounding parts of the image.

- Default is 20.0 pixels.
- Range is 1.0 to 40.0 pixels.
- Default usually works fine. Experiment to find most natural look.

Image Type

There are different Deconvolution modes for Deep Space and Lunar/Planetary targets.

- Deep Space - Deconvolution makes no special provision for extra dynamic range.
- Lunar/Planetary - frees up dynamic range for any deconvolved highlights.
- Default is 'Deep Space'.

'Color': Advanced 'Color' Correction and Manipulation

Top: traditional processing, Bottom: StarTools color constancy showing true color of the core, regardless of brightness. (image acquisition by Jim Misti)

Thanks to StarTools' Tracking feature the 'Color' module provides you with unparalleled flexibility when it comes to color presentation in your image.

The 'Color' module fully capitalizes on the signal processing engine's unique ability to process chrominance and detail separately, yet simultaneously. This unique capability is responsible for a number of innovative features.

Firstly, whereas other software without Tracking data mining, destroys color and color saturation in bright parts of the image as the data gets stretched, StarTools allows you to retain color and saturation throughout the image with its 'Color Constancy' feature. This ability allows you to display all colors in the scene as if it were evenly illuminated, meaning that even very bright cores of galaxies and nebulas retain the same color throughout, irrespective of their local brightness, or indeed acquisition methods and parameters.

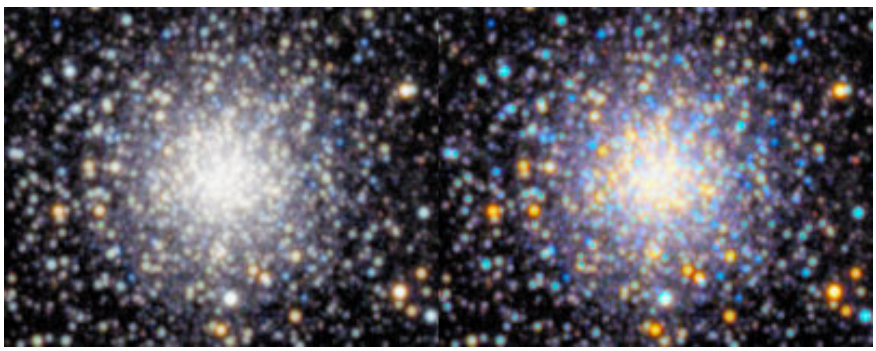

Left: traditional processing, Right: StarTools color constancy showing star temperatures evenly until well into the core.

This ability is important in scientific representation of your data, as it allows the viewer to compare similar objects or areas like-for-like, since color in outer space very often correlates with chemical signatures or temperature.

The same is true for star temperatures across the image, even in bright, dense star clusters. This mode allows the viewer of your image to objectively compare different parts and objects in the image without suffering from reduced saturation in bright areas. It allows the viewer to explore the universe that you present in full color, adding another dimension of detail, irrespective of the exposure time and subsequent stretching of the data.

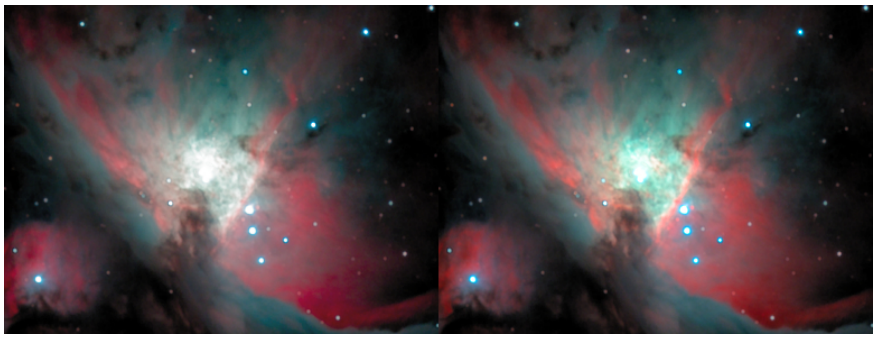

'Color' constancy (right) demonstrates how features with similar chemical/physical properties show identical colors, regardless of brightness.

For example, StarTools enables you to keep M42's color constant throughout, even in its bright core. No fiddling with different exposure times, masked stretching or saturation curves needed. You are able to show M31's true colors instead of a milky white, or resolve star temperatures to well within a globular cluster's bright core. All that said, if you're a fan of the traditional 'handicapped' way of color processing in other software, then StarTools can emulate this type of processing as well.

The 'Color' module's abilities don't stop there, however. It is also capable of emulating a range of complex LRGB color compositing methods that have been invented over the years. And it does it at the click of a button! Even if you acquired data with an OSC or DSLR, you will still be able to use these compositing methods; the 'Color' module will generate synthetic luminance from your RGB on the fly and re-composite the image in your desired compositing style.

The 'Color' module allows for various ways to calibrate the image, including by star field, sampling G2V star, galaxy sampling and - unique to StarTools - the MaxRGB calibration view. The latter allows for objective color calibration, even on poorly calibrated screens.

Because luminance (detail) and chrominance is processed separately in parallel, the module is capable of remapping channels for the purpose of color (aka "tone mapping") on the fly, without impacting detail. The result is the unique ability to flip between popular color renditions for, for example, narrowband data with a single click, whether you are processing SHO/HST datasets or duo/tri/quad-band datasets. Similarly, DSLR users benefit from the ability to use the manufacturer's preferred color matrix, yet without the cross-channel noise contamination that would otherwise impact luminance (detail).

Usage

The 'Color' module is very powerful - offering capabilities surpassing most other software - yet it is simple to use.

The primary goal that the 'Color' module was designed to accomplish, is achieving a good color balance that accurately describes the color ratios that were recorded. In accomplishing that goal, the 'Color' module goes further than other software by offering a way to negate the adverse effects of non-linear dynamic range manipulations on the data

(thanks to Tracking data mining). In simple terms, this means that coloring can be reproduced (and compared!) in a consistent manner regardless of how bright or dim a part of the scene is shown.

A second unique feature of StarTools, is its ability to process luminance (detail) and chrominance (color) separately, yet simultaneously. This means that any decisions you make affecting your detail does not affect the coloring of said detail, and vice-versa. This ability further allows you to remap color channels (aka "tone mapping") for narrowband data, without having to start over with your detail processing. This lets you try out many different popular color schemes at the click of a button.

Upon launch, the color module blinks the mask three times in the familiar way. If a full mask is not set, the 'Color' modules allows you to set it now, as color balancing is typically applied to the full image (requiring a full mask).

In addition to blinking the mask, the 'Color' module also analyses the image and sets the Red Bias Reduce, Green Bias Reduce and Blue Bias Reduce factors to a value which it deems the most appropriate for your image. This behavior is identical to manually clicking the 'Sample' button where the whole image is sampled.

Launching the 'Color' Module

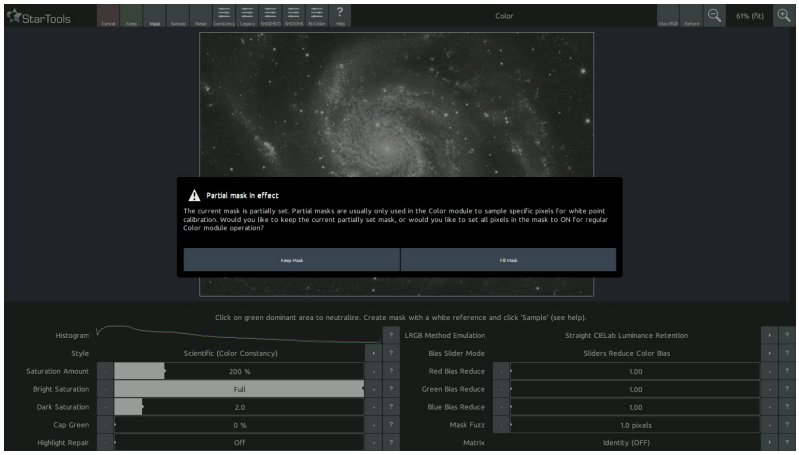

If a full mask is not set, the 'Color' modules allows you to set it now, as color balancing is typically applied to the full image (requiring a full mask).

StarTools tends to come up with a reasonable color balance by default, but may sometimes need some help if a dataset contains aberrant color information in the highlights. In this case, aberrant color information in the star cores caused the balance to be a tad core tad too green.

In cases where the image contains aberrant color information in the highlights, for example due to chromatic aberration or slight channel misalignment/discrepancies, then this initial color balance may be incorrect and may need further correct. The aberrant color information in the highlights itself, can be repaired using the 'Highlight 'Repair" parameter.

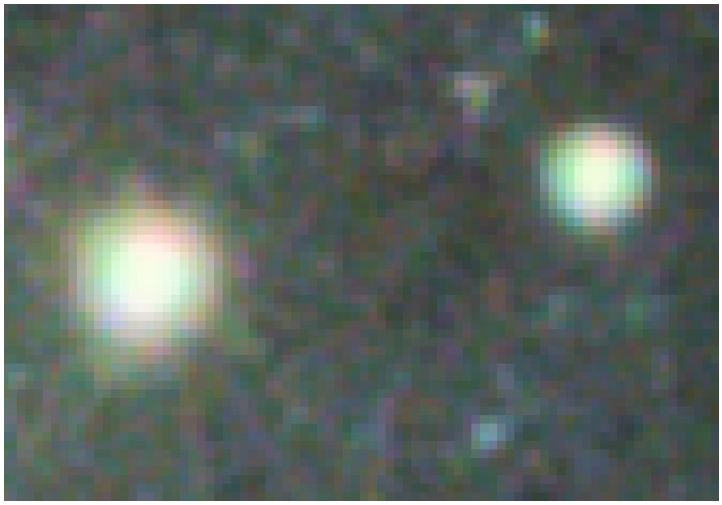

Aberrant color information due to misalignments/fringing or chromatic aberration, may throw off the initial color balance.

Setting a color balance

The Red, Green and Blue Bias controls.

The 'Red, Green and Blue Increase/Reduce' parameters are the most important settings in the 'Color' module. They directly determine the color balance in your image. Their operation is intuitive; too much red in your image? Pump up the 'Red Bias Reduce' value. Too little red in your image? Reduce the 'Red Bias Reduce' value.

If you'd rather operate on these values in terms of Bias Increase, then simply switch the 'Bias Slider Mode' setting to 'Sliders Increase 'Color' Bias'. The values are now represented in terms of relative increases, rather than decreases. Switching between these two modes you can see that, for example, a RedBias Reduce of 8.00 is the same as a Green and Blue Bias Increase of 8.00. It makes intuitive sense when you think about it - a relative decrease of red makes blue and green more prevalent and vice versa.

In addition to blinking the mask, the 'Color' module also analyses the image and sets the Red Bias Reduce, Green Bias Reduce and Blue Bias Reduce factors to a value which it deems the most appropriate for your image. This behavior is identical to manually clicking the 'Sample' button.

'Color' balancing techniques

Now that we know how to change the color balance, how do we know what to actually set it to?

There are a great number of tools and techniques that can be applied in StarTools that let you home in on a good color balance. Before delving into them, It is highly recommended to switch the 'Style' parameter to 'Scientific ('Color' Constancy)' during color balancing, even if that is not the preferred style of rendering the color of the end result, this is because the 'Color' Constancy feature makes it much easier to color balance by eye in some instances due to its ability to show continuous, constant color throughout the image. Once a satisfactory color balance is achieved, of course, feel free to switch to any alternative style of color rendering.

White point reference by mask sampling

Upon launch, or upon clicking the Sample button, the 'Color' module samples whatever mask is set (note also that the set mask also ensures the 'Color' module only applies any changes to the masked-in pixels!) and sets the Red, Green and Blue bias settings accordingly.

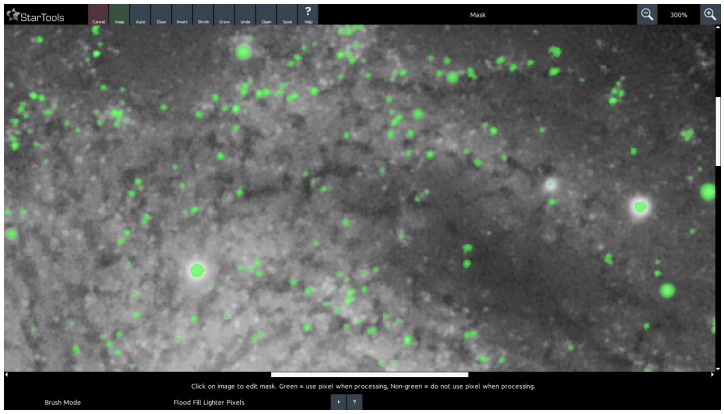

We can calibrate against a big enough population of non-associated foreground stars, by putting them in a mask, clicking 'Sample' in the 'Color' module and applying the found bias values to the whole image again.

We can use this same behavior to sample larger parts of the image that we know should be white. This method mostly exploits the fact that stars come in all sorts of sizes and temperatures (and thus colors!) and that this distribution is completely random in a wide enough field. Indeed, the Milky Way is named as such because the average color of all its stars is perceived as a milky white. Therefore if we sample a large enough population, we should find the average star to be somewhere in the middle.

A reasonably good color balance achieved by putting all stars in a mask using the Auto feature and sampling them.

We can accomplish that in two ways; we either sample all stars (but only stars!) in a wide enough field, or we sample a whole galaxy that happens to be in the image (note that the galaxy must be of a certain type to be a good candidate and be reasonably close - preferably a barred spiral galaxy much like our own Milkyway).

Whichever you choose, we need to create a mask, so we launch the Mask editor. Here we can use the Auto feature to select a suitable selection of stars, or we can us the Flood Fill Brighter or Lasso tool to select a galaxy. Once selected, return to the 'Color' module and click Sample. StarTools will now determine the correct Red, Green and Blue bias to match the white reference pixels in the mask so that they come out neutral.

To apply the new color balance to the whole image, launch the Mask editor once more and click Clear, then click Invert to select the whole image. Upon return to the 'Color' module, the whole image will now be balanced by the Red, Green and Blue bias values we determined earlier with just the white reference pixels selected.

White balancing in MaxRGB mode

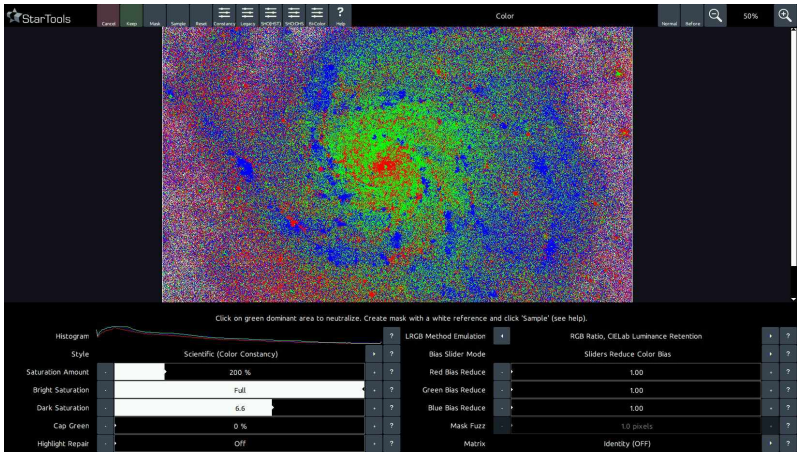

Major green channel dominance in the core points to color imbalance in that area.

StarTools comes with a unique color balancing aid called MaxRGB. This mode of color balancing is exceptionally useful if trying to color balance by eye, but the user suffers from color blindness or uses a screen that is not color calibrated very well.

The MaxRGB aid allows you to view which channel is dominant per-pixel. If a pixel is mostly red, that pixel is shown red, if a pixel is mostly green, that pixel is shown green, and if a pixel is mostly blue, that pixel is shown blue.

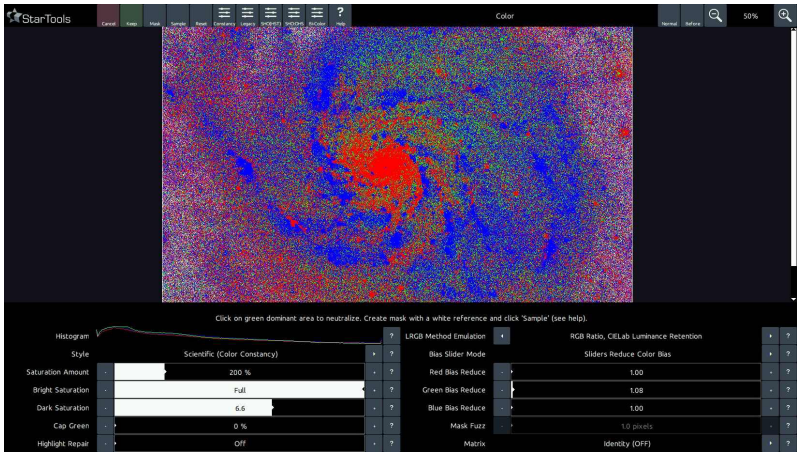

Reducing the green bias has removed green dominance in the core, leaving only spurious/random green dominance due to noise.

Switching from MaxRGB mode to Normal mode confirms the image still looks good.

Therefore, if we see large areas of green, we know that we have too much green in our image and we should adjust the bias accordingly. Similarly if we have too much red or blue in our image, the MaxRGB mode will show many more red than blue pixels in areas that should show an even amount (for example the background). Again we then know we should adjust red or green accordingly.

Clicking on an area to neutralize green

A convenient way to eliminate green dominance is to simply click on an area. The 'Color' module will adjust the 'Green Bias Reduce' or 'Green Bias Increase' in response so that any green dominance in that area is neutralized.

White balancing by known features and processes

M101 exhibiting a nice yellow core, bluer outer regions, red/brown dust lanes and purple HII knots, while the foreground stars show a good distribution of color temperatures from red to orange, yellow, white to blue.

StarTools' 'Color' Constancy feature makes it much easier to see colors and spot processes, interactions, emissions and chemical composition in objects. In fact, the 'Color' Constancy feature makes coloring comparable between different exposure lengths and different gear. This allows for the user to start spotting colors repeating in different features of comparable objects. Such features are, for example, the yellow cores of galaxies (due to the relative over representation of older stars as a result of gas depletion), the bluer outer rims of galaxies (due to the relative over representation of bright blue young stars as a result of the abundance of gas) and the pink/purplish HII area 'blobs' in their discs. Red/brown (white light filtered by dust) dust lanes complement a typical galaxy's rendering.

Similarly, HII areas in our own galaxy (e.g. most nebulae), while in StarTools 'Color' Constancy Style mode, display the exact same color signature found in the galaxies; a pink/purple as a result of predominantly deep red Hydrogen-alpha emissions mixed with much weaker blue/green emissions of Hydrogen-beta and Oxygen-III emissions and (more dominantly) reflected blue star light from bright young blue giants who are often born in these areas, and shape the gas around them.

Dusty areas where the bright blue giants have 'boiled away' the Hydrogen through radiation pressure (for example the Pleiades) reflect the blue star light of any surviving stars, becoming distinctly blue reflection nebulae. Sometimes gradients can be spotted where (gas-rich) purple gives away to (gas-poor) blue (for example the Rosette core) as this process is caught in the act.

Diffraction spikes, while artifacts, also can be of great help when calibrating colors; the 'rainbow' patterns (though skewed by the dominant color of the star whose light is being diffracted) should show a nice continuum of coloring.

Finally, star temperatures, in a wide enough field, should be evenly distributed; the amount of red, orange, yellow, white and blue stars should be roughly equal. If any of these colors are missing or are over-represented we know the color balance is off.

Color balancing of data that was filtered by a light pollution filter

Color balancing of data that was filtered by a light pollution filter is fundamentally impossible; narrow (or wider) bands of the spectrum are missing and no amount of color balancing is going to bring them back and achieve proper coloring. A typical filtered data set will show a distinct lack in yellow and some green when properly color balanced. It's by no means the end of the world - it's just something to be mindful of.

Correct coloring may be achieved however by shooting deep luminance data with light pollution filter in place, while shooting color data without filter in place, after which both are processed separately and finally combined. Color data is much more forgiving in terms of quality of signal and noise; the human eye is much more sensitive to noise in the luminance data that it is in the color data. By making clever use of that fact and performing some trivial light pollution removal in 'Wipe', the best of both worlds can be achieved.

A visual spectrum color balance will not be possible with datasets shot through a light pollution filter, however pleasing results showing important coloring (for example emissions and reflection nebulosity) quite accurately, can still be achieved.

OSC (One-Shot-'Color') instruments

Many more modern OSC cameras have a spectrum response that increases in sensitivity across all channels beyond the visual spectrum red cut-off (the human eye can detect red wavelengths up until around 700nm). This is a feature that allows these cameras pick up detail beyond the visual spectrum (for example for use with narrowband filters or for recording infrared detail).

However, imaging with these instruments without a suitable IR/UV filter (also known as a "luminance filter") in place, will cause these extra-visual spectrum wavelengths to accumulate in the visual spectrum channels. This can significantly impact the "correct" (in terms of visual spectrum) coloring of your image. Just as a light pollution filter makes it fundamentally impossible to white-balance back the missing signal, so too does imaging with extended spectrum response make it impossible to white-balance the superfluous signal away.

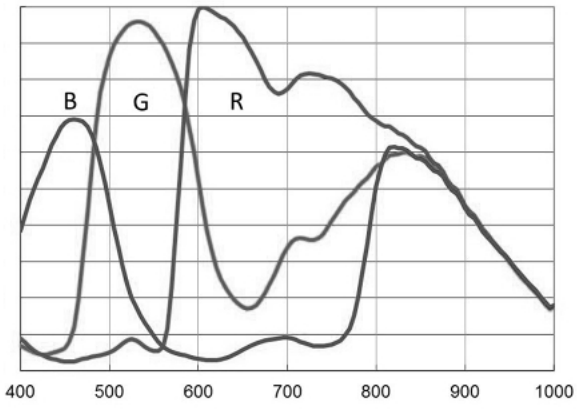

This example spectral response graph of a ZWO ASI290MC camera shows a marked "bump" in the green and blue response beyond the red visual spectrum cut-off (approximately 700nm).

Hallmarks of datasets that have been acquired with such instruments, without a suitable IR/UV filter in place, is a distinct yellow cast that is hard (impossible) to get rid of, due to a strong green response coming back in combined with extended red channel tail.

The solution is to image with a suitable IR/UV filter in place that cuts-off the extended spectrum response before those channels increase in sensitivity again. The needed IR/UV filter will vary per OSC. Consult the respective manufacturers' spectral graphs to find the correct match for your OSC.

Tweaking your colors

A more 'handicapped' way of showing colors is also available, emulating the way other software distorts and destroys hues and saturation along with stretching the luminance data.

Once you have achieved a color balance you are happy with, the StarTools 'Color' module offers a great number of ways to change the presentation of your colors.

Style

The parameter with the biggest impact is the 'Style' parameter. StarTools is renowned for its 'Color' Constancy feature, rendering colors in objects regardless of how the luminance data was stretched, the reasoning being that colors in outer space don't magically change depending on how we stretch our image. Other software sadly lets the user stretch the color information along with the luminance information, warping, distorting and destroying hue and saturation in the process. The 'Scientific ('Color' Constancy)' setting for Style undoes these distortions using Tracking information, arriving at the colors as recorded.

To emulate the way other software renders colors, two other settings are available for the Style parameter. These settings are 'Artistic, Detail Aware' and 'Artistic, Not Detail Aware'. The former still uses some Tracking information to better recover colors in areas whose dynamic range was optimized locally, while the latter does not compensate for any distortions whatsoever.

LRGB Method Emulation

The LRGB Method Emulation allows you to emulate a number of color compositing methods that have been invented over the years. Even if you acquired data with an OSC or DSLR, you will still be able to use these compositing methods; the 'Color' module will generate synthetic luminance from your RGB on the fly and re-composite the image in your desired compositing style.

The difference in coloring can be subtle or more pronounced. Much depends on the data and the method chosen.

- 'Straight CIELab Luminance Retention' manipulates all colors in a psychovisually optimal way in CIELab space, introducing color without affecting apparent brightness.

- 'RGB Ratio, CIELab Luminance Retention' uses a method first proposed by Till Credner of the Max-Planck-Institut and subsequently rediscovered by Paul Kanevsky, using RGB ratios multiplied by luminance in order to better preserve star color. Luminance retention in CIELab color space is applied afterwards.
- '50/50 'Layer'ing, CIELab Luminance Retention' uses a method proposed by Robert Gendler, where luminance is layered on top of the color information with a 50% opacity. Luminance retention in CIELab color space is applied afterwards. The inherent loss of 50% in saturation is compensated for, for your convenience, in order to allow for easier comparison with other methods.
- 'RGB Ratio' uses a method first proposed by Till Credner of the Max-Planck-Institut and subsequently rediscovered by Paul Kanevsky, using RGB ratios multiplied by luminance in order to better preserve star color. No further luminance retention is attempted.
- '50/50 'Layer'ing, CIELab Luminance Retention' uses a method proposed by Robert Gendler, where luminance is layered on top of the color information with a 50% opacity. No further luminance retention is attempted. The inherent loss of 50% in saturation is compensated for, for your convenience, in order to allow for easier comparison with other methods.
- Note that the LRGB Emulation Method feature is only available when Tracking is engaged.

Saturation

Increasing saturation makes colors more vivid, while increasing the Dark Saturation response parameter introduces more color in the shadows.

The 'Saturation' parameter allows colors to be rendered more, or less vividly, whereby Bright Saturation and Dark Saturation control how much color and saturation is introduced in the highlights and shadows respectively. It is important to note that introducing color in the shadows may exacerbate color noise, though Tracking will make sure any such noise exacerbations are recorded and dealt with during the final denoising stage.

Cap Green

The 'Cap Green' parameter, finally, removes spurious green pixels if needed, reasoning that green dominant colors in outer space are rare and must therefore be caused by noise. Use of this feature should be considered a last resort if color balancing does not yield adequate results and the green noise is severe. The final denoising stage should, thanks to Tracking data mining, pin pointed the green channel noise already and should be able to adequately mitigate it.

Matrix correction and remapping

The 'Color' module comes with a vast number of camera color correction matrices for various manufacturers (Canon, Nikon, Sony, Olympus, Pentax and more), as well as a vast number of channel blend remappings for narrowband dataset (e.g. HST/SHO or bi-color duo-band/quad-band filter data).

Uniquely, thanks to the signal evolution Tracking engine, this color calibration is preferably performed towards the end of your processing workflow. This allows you to switch color rendering at the very last moment at the click of a button without having to re-composite and re-process, while also allowing you to use cleaner, non-whitebalanced, non-matrix corrected data for your luminance component, aiding signal fidelity.

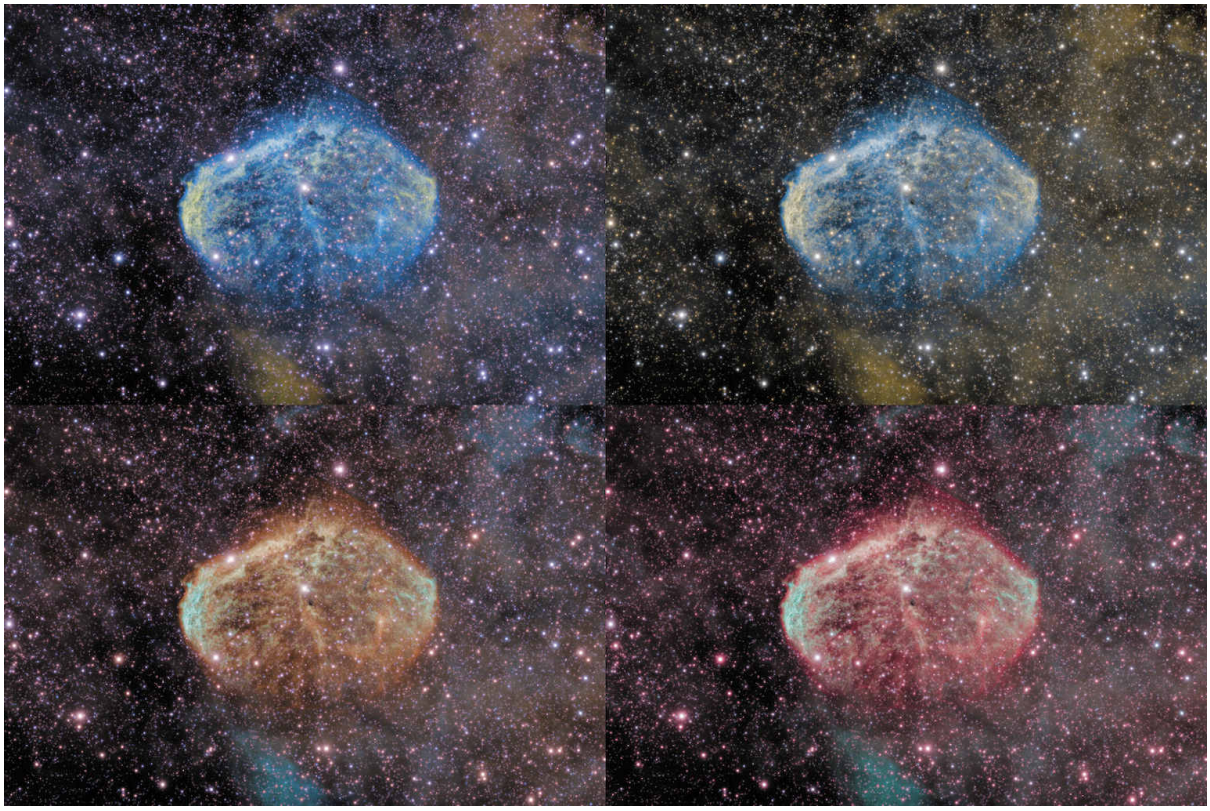

Four example renderings of the same S-II + H-alpha + O-III narrowband data. Top: two examples of a 1-click SHO (HST palette) rendering. Bottom: two examples of a OHS rendering. No recompositing is needed, and detail is kept fully intact - only the coloring changed.

Camera Matrix correction is performed towards the end of your processing workflow on your chrominance data only, rather than in the RAW converter during stacking. This helps improve luminance (detail) signal, by not contaminating it with cross-channel camera-space RGB and XYZ-space manipulations.

The matrix or channel blend/mapping is selected using the 'Matrix' parameter. Please note that the available options under this parameter are dependent on the type of dataset you imported. Please use the Compose module to import any narrowband data separately.

Presets

As in most modules in StarTools, a number of presets are available to quickly dial in useful starting points.

- 'Constancy' sets the default 'Color' Constancy mode and is the recommended mode for diagnostics and color balancing in.

- 'Legacy' switches to a color rendition for visual spectrum datasets that is closest to what legacy software (e.g. software without signal evolution Tracking) would produce. This will mimic the way such software (incorrectly) desaturates highlights and causes hue shifts.
- 'SHO(HST)' dials in settings that are a good starting point for datasets that were imported as S-II, H-alpha and O-III for red, green and blue respectively (also known as the 'SHO, 'SOH:RGB', 'HST' or 'Hubble' palette). This standard way of importing datasets and mapping the 3 bands to the 3 channels in this way (via the Compose module), allows for further channel blends and remapping via the 'Matrix' parameter. Please note the specific blend's parameters/factors under the 'Matrix' parameter. This preset also greatly reduces the green bias to minimize green, while attempting to bring out the popular golden hues.
- 'SHO:OHS' is similar to the 'SHO(HST)' preset, except that it further remaps a SHO-imported dataset to a channel blend that is predominantly mapped as OHS:RGB instead. Renditions typically yield a pleasing "glowing ice-on-fire" effect.
- 'Bi-'Color" assumes a dataset was imported as HOO:RGB, that is Ha-alpha imported as red, and O-III (sometimes also incorporating H-beta) imported as green and also blue. This yields the popular red/cyan bi-color renderings that are so effective at showing dual emission dominance. This preset is also particularly useful and popular for people who use a duo-band filter (aka as a tri-band or quad-band filter) with an OSC or DSLR.

Wrap Up

Description

- To achieve a good color balance that accurately describes the color ratios that were recorded.
- Processing can modify the color balance - in particular stretching the image to bring out detail can have a bad effect. This module helps correct those changes.

Important

- To get best results this module should be used on a color calibrated monitor. If you can't do that the 'Max RGB' button helps by showing the dominant channel (R,G or B) for each pixel. This can help you adjust the color if you have an uncalibrated monitor.

When to use

- Towards the end of the workflow (after all the stretching) just before turning off Tracking and doing the final 'Denoise'.
- You can re-use this module if you use one of the 'LRGB Method Emulation' modes that uses CIELab Luminance Retention.
- Re-use if you want to adjust the color or saturation in different areas by using a mask.

Sampling Methods

Parts of the image can be sampled to establish a good color balance. The following ways are available:

See also the topic Auto 'Color' Balance and 'Color' Balancing Techniques.

- Sampling overrides the current settings.
- Use Sampling when automatic color balance is difficult due to, for example, noisy data.
- Sampling doesn't work when a filter (e.g. a narrow band or light pollution filter) has been used as the stars are all missing part of their spectrum - so no good reference white objects, or broad range of spectral types, exist. See 'Special Techniques' for an alternative approach.

- This color balance technique is based on the assumption that the sampled object(s) is, on average, a good reference white so we can establish the correct relative balance of the R,G and B channels from this.
- White light contains all the visible spectrum so we look for objects which contain the same broad range:
  - 'White' Galaxies and some Star fields contain a broad range of spectral class stars and so are considered good reference white.
  - Spectral class G2V stars, like our sun, are often considered suitable white references but some people argue that they don't have a unique role as reference source. See the discussion in the Background Notes.
- Globular Clusters often don't have a good mix - they mostly have very old stars - mainly yellow with some orange and reds - so should be excluded from a star field calibration.
- For a discussion of these color balancing techniques see the following references:
- 'Getting the colors right in astrophotos'.
- 'PixInsight color calibration methodology'.
- See also this Starizona article on 'True 'Color' Imaging'.

Single sample

See also 'Setting white reference by clicking pixel'.

- Click on an area in the image that should be white.

Sample from Mask

See also 'Setting white reference by mask sampling'.

- Define the elements that make up your sample in the mask (e.g. a galaxy, G2V star, or star field with good mix of star temperatures).
- Click the 'Sample' button - this uses the mask to define the sample from which the white balance is determined and the Red, Green and Blue bias settings are established.

Star Field Calibration - assumes the star field has a good mix of star temperatures

See also Starfield Color Calibration.

- Mask - Auto - 'Fat Stars' preset - Do - (optional grow 1 pixel) - Keep.
- Click Sample (uses the star samples and sets RGB bias control settings based on that).
- Mask - Clear - Invert - Keep.
- The module remembers the RGB settings but now applies them to the whole image (based on the new mask).

What result to look for

The color distribution you are looking for depends to a large degree on your preferences.
(See also the topic M31 in Moonlight).
If you look at these Thumbnails of Images of M8 you can see the range of colors of the processed M8 images on the internet. Many of them have a red bias and this can be a side-effect of the non-linear stretching causing a color skew that has not been compensated for.

- If using the 'Scientific ('Color' Constancy)' approach you should look for the following:
  - Good distribution of star colors - Foreground stars should show a good distribution of color temperatures from red through orange, yellow and white to blue.
  - If a light pollution filter is used the star colors may just be orange and blue with little in between - yellow is often missing.

    - Check for green - This should be rare unless there is an OIII emission region (e.g. M42 core or Tarantula Nebula).
    - The H-alpha should look red, H-beta should look cyan.
    - HII areas (H-alpha + H-beta) should look purplish/pink.
    - Galaxy cores tend to look yellow (older stars) and their outer rims tend to look bluer (younger stars & star formation).
    - Dust tends to let through lower wavelength light (if any) - mainly browns and reds.
- If using the 'Artistic, Detail Aware' or 'Artistic, Not Detail Aware' styles you should look for the following:
    - Bright areas will be paler, less colorful than above.
- If using narrow-band techniques - such as the Hubble palette or similar.
    - The automatic color balancing that occurs on loading the module will probably give odd results because there is so much of the spectrum missing.
    - Set the 3 Bias Sliders back to 1.00 and manually adjust until you get the result you want. This is false color so there is no 'right' result.
    - Ensure the relative strength of each channel highlights the detail you want.
- Light Pollution filter - tends to have a dip where the yellows should be - so galaxy cores lose their yellow.

Ways of getting better results

- Try to correct any color problems (such as those caused by light pollution and gradients) using the 'Wipe' module first.
- Help the 'Wipe' module by using Flats.
- Temporarily increase the Saturation Amount control (to say 300%) while working with color to help when gauging color balance.
- If using a light pollution filter visually getting the right balance will be very difficult as there are parts of the spectrum missing. It will often show a lack of yellow and some green when properly color balanced - There is a way around this - see the Special Techniques section below.

After Use

If needed, use 'Shrink', 'Filter', 'Entropy' or 'SuperStructure' modules. You might also stop Tracking and to use the 'Denoise' module.

Description of Controls

Top Buttons

- Reset - Sets the controls to neutral settings. Saturation to 100%, Bright and Dark Saturation to 1.0, R,G and B Bias to 1.0, All other settings to defaults.
- Sample - Uses the current mask setting as a sample set from which the white balance is determined and the Red, Green and Blue bias settings are established.
- Mask - For general instructions on using mask see Mask:
    - The mask can be used to select areas to sample as a color reference - see Sampling Methods section above. Click 'Sample' when done.
    - The mask can also be used to selectively adjust the color of areas in the image. To do this you need to clear the mask before starting the module and set it when prompted at the start.
- Constancy - Optimized preset settings to support the 'Scientific ('Color' Constancy)' style.
- Legacy - Optimized preset settings to support the 'Artistic, Not Detail Aware' style.
- Hubble - Optimized preset settings for narrow band data sets using the Hubble palette (v1.5).

- SHO(HST) - Optimized preset settings for narrow band data sets using the Hubble palette (v1.6).
- SHO:OHS - Optimized preset settings for narrow band data sets using the SHO (v1.6).
- Duo-band - Optimized preset settings for narrow band HO data sets (v1.6).
- Cancel - Exits the module - discarding any changes.
- Keep - Exits the module - saving any changes.
- Max RGB - For each pixel, shows which channel - R, G or B - is dominant. See Max RGB Mode.
- If your image is too red, pixels that are supposed to be 'neutral' (such as the background) will show mostly red. If your image is too green they will show mostly green. If, however, your image is well calibrated, these neutral pixels will alter between red, green and blue.
- Before/After - Toggles the display of the image between the current and initial view.

Histogram

The histogram shows the distribution of pixel intensity for the separate R,G and B channels.

- Displays the pixels intensity distribution split into RGB channels.
- Only pixels set in the current mask are counted.

Style

See also the article Tweaking your colors.

Options:

- Scientific ('Color' Constancy) - keeps the color regardless of brightness by separating luminance and color processing.
- Artistic, Detail Aware - Emulates much other software where bright areas can look washed out. Tries to compensate for local brightness manipulations during processing (e.g when using 'HDR'), In these areas it will compensate for these changes and show more color.
- Artistic, Not Detail Aware - As above but does not try to compensate for local brightness manipulations during processing.
- The Style control is only available when Tracking is engaged.

Saturation Amount

Specifies the amount of color saturation relative to the original image.

- Default is 200%
- Range is 0-1,000%.
- Reducing to 0% turns the image monochrome.

Bright Saturation

Specifies the color saturation in the lighter areas.

- Default is Full (10)
- Range is 1.00-10.00 (Full)
- Reduce this value where there are color artifacts noticeable in the highlights. For example where there are color fringes around bright star cores with one side blue and the opposite side red.

Dark Saturation

Specifies the color saturation in the darker areas.

- Default is 2.00
- Range is 1-10
- Reduce this value if there is a lot of color noise in the dark background.

Cap Green

- Green is produced by OIII emission regions (e.g. M42 core or Tarantula Nebula) which are rare.
- Stretching color data with luminance causes a skew in the color balance.
- Very few objects in space are predominantly green when imaged in RGB. So if we find a green pixel, and we are sure the color balance is right, we can assume any pixels that are green are made that way by noise so we convert them to something more natural like yellow or brown.
- Default is '0%'.
- Range is 0% to 100%.
- Use as a final change if necessary.

Highlight 'Repair'

To remove color defects around highlights due to debayering or channel alignment issues.

- Specifies the area size around highlights where color data should be repaired.
- Default is 'Off'
- Range is Off-10 pixels

LRGB Method Emulation

With Tracking on, StarTools has from the start separated L and RGB (if the data was linear when imported).

Now is the time to combine them - there are a number of different approaches to combining them to choose from here:

- Straight CIELab Luminance Retention - adjusts all colors in a psychovisually optimal way in CIELab space, introducing color without affecting apparent brightness.
- RGB Ratio, CIELab Luminance Retention - Applies the RGB Ratio technique, with luminance retention in CIELab color space being applied afterwards.
- 50/50 Layering, CIELab Luminance Retention - Applies the '50/50 Layering' technique, with luminance retention in CIELab color space being applied afterwards.
- RGB Ratio - This uses RGB ratios multiplied by luminance in order to better preserve star color saturation.
- 50/50 Layering - Here the luminance is layered on top of the color information with a 50% opacity.
- The default is 'Straight CIELab Luminance Retention'.

Bias Slider Mode

Sets whether the Bias sliders increase or reduce the channel influence.

- Default is 'Sliders Reduce Color Bias'.

Bias Sliders

These sliders set the color balance. They can be adjusted manually or by one of the color sampling techniques described above.

- Red Bias Reduce (or Red Bias Increase):
  - Default is 1.00 (no reduction)
  - Range is 1.00-20.00.
- Green Bias Reduce (or Green Bias Increase):
  - Default 1.00 (no reduction)
  - Range is 1.00-20.00.
- Blue Bias Reduce (or Blue Bias Increase):
  - Default is 1.00 (no reduction)
  - Range is 1.00-20.00.

Mask Fuzz

- If a mask is used, Mask Fuzz controls the blending of the transition between masked and non-masked parts of the image.
- Default is 1.0 pixels.
- Range is 1.0 to 41.0 pixels.

Highlight 'Repair'

Specifies the size of the area around highlights (e.g. stars) where inaccurate color data may exist and will be repaired.

- 'Color' inaccuracy may be due to channel alignment or debayering issues.
- Increasing this value increases the area around the highlight which will be repaired.
- Default is Off (0 pixels)
- Range is 0-10 pixels.

Matrix

Contains either:

- a list of DSLR cameras - if the dataset is imported as 'Linear, from OSC/DSLR and not white-balanced'.
  - This allows selection of the right color correction matrix for your camera.
  - Options include cameras from Canon, Nikon, Olympus, Pentax, Samsung & Sony
- a list of common channel blends to choose from - if the dataset is imported as 'Linear'.
  - This allows the selection of your preferred channel blend remapping.
  - Options include combinations based on: SHO, SHO:OHS, SHO:HOS, Duo-band HO, HO(H+O), RGB:RGB combinations, False 'Color': Solar
- Default is 'Identity (Off)'.

'Shrink': Star Appearance Manipulation

Top left: input image. Top right: Tighten preset. Bottom left: Dim preset. Bottom right: Un-glow preset with all else turned off; a subtle contrast increase can be seen around bright stars.

The 'Shrink' module offers comprehensive stellar profile modification by shrinking, tightening and re-coloring stars.

Usage

A good star mask is essential for good results. Even though the 'Shrink' module is much more gentle on structural detail, ideally, only stars are treated and not any structural detail.

The 'AutoMask' button launches a popup with access to two quick ways of creating a star mask. This same popup is shown upon first launch of the module. The generated masks tend to catch all major stars with very few false positives. If you also wish to include fainter, small stars in the mask, then more sophisticated techniques are recommended to avoid including other detail.

Finally, if your object is mostly obscured by a busy star field, for example in a widefield, then also consider using the Super Structure module to enhance the super structures in your image and push back the busy star field. Combining both the 'Shrink' module's output and the Super Structure module's output can greatly transform a busy looking image in positive ways.

Parameters

Two 'Mode' settings are available:

- 'Tighten' has the effect of tightening a stars around their central core.
- 'Dim' has the effect of dimming stars luminosity.

The 'Shrink' module uses an iterative process; the strength of the Tighten or Dim effect is controlled by the number of 'Iterations', as well as the 'Regularization' parameter that dampens the effect. The stringing and pitting artifacts commonly produced by less sophisticated techniques, is thereby avoided.

Top: original image. Bottom: 'Color Taming' parameter used. Note the stars now appear less conspicuous as their colors blend in with the rest of the object.

The 'Color' Taming' parameter forces stars to progressively adopt the coloring of their surroundings, like 'chameleons'.

The 'Halo Extend' parameter effectively grows the given mask temporarily, thereby including more of each star's surroundings.

If the image has been deconvolved or sharpened and the stars may be subject to subtle ringing artifacts, then the 'De-ringing' parameter will take this into account when shrinking the stellar profiles, as to not exacerbate the ringing.

The 'Un-glow' feature attempts to reduce the halos around bright, over-exposing stars. 'Un-glow Strength' throttles the strength of the effect. The 'Un-glow Kernel' specifies the width of the halos.

The 'Shrink' module's iterative algorithm avoids these sort of stringing and pitting artifacts typically produced by unsophisticated morphological transformations

Creating a suitable star mask

A good star mask is essential for good results. Though the 'Shrink' module is much more gentle on structural detail than the basic unsophisticated morphological transformations (such as minimum filters) found in other software, ideally, only stars are treated and not any nebulosity, gaseous filaments or other structural detail.

The 'AutoMask' button launches a pop-up with access to two quick ways of creating a star mask. This same popup is shown upon first launch of the module. The generated masks tend to catch all major stars with very few false positives. If you also wish to include fainter, small stars in the mask, then more sophisticated techniques are recommended to avoid including other detail. Besides touching up the mask by hand, it is also possible to combine the results of an aggressive auto-generated star mask (catching all faint stars), with a less aggressive auto-generated star mask (catching fewer faint stars, but also leaving structural detail alone):

1. Clear the mask, and select the part of the image you wish to protect with the Flood Fill Lighter or Lasso tool, then click Invert.

2. In the Auto mask generator, set the parameters you need to generate your mask (here we chose the 'Stars' preset and set the 'Source' parameter to 'Stretched' to avoid any noise mitigation measures that may otherwise filter out faint stars for selection). Be sure to set 'Old Mask' to 'Add New Where Old Is Set'.

3. After clicking 'Do'. The auto-generator will generate the desired mask, however excluding the area we specified earlier.

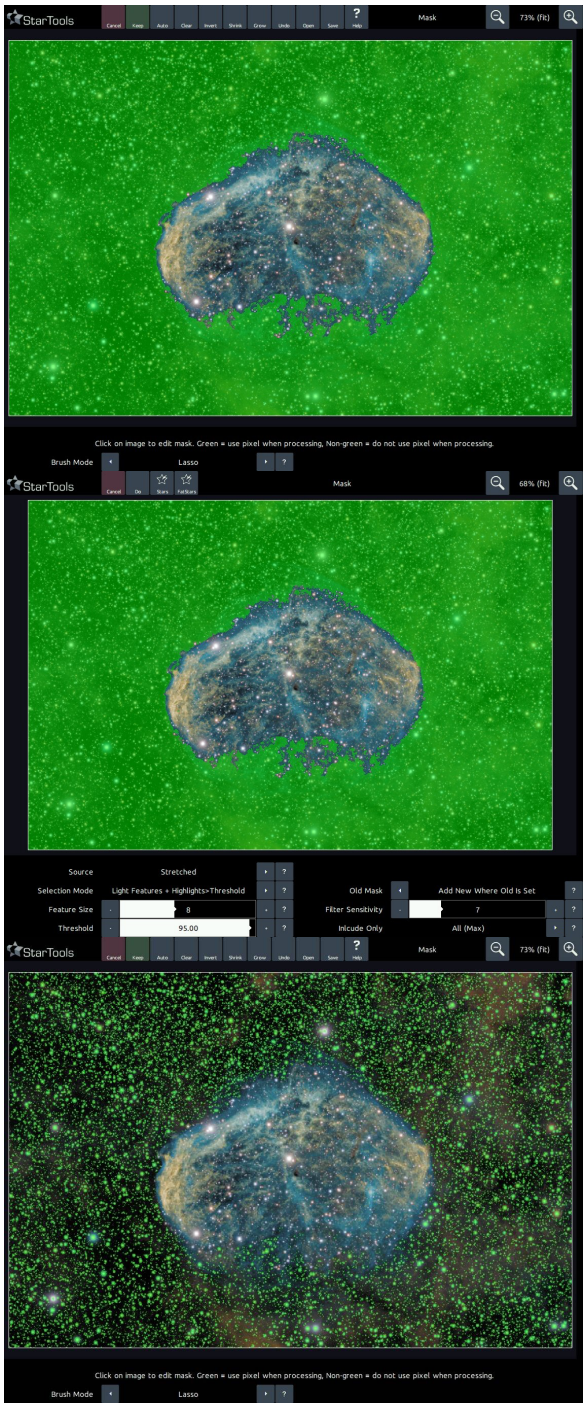

4. Launch the Auto mask generator once more. Click the 'Stars' preset again. This time set 'Old Mask' to 'Add New To Old' to add the newly generated mask to the mask we already have. This will fill in the area we excluded earlier with the less aggressive mask as well.

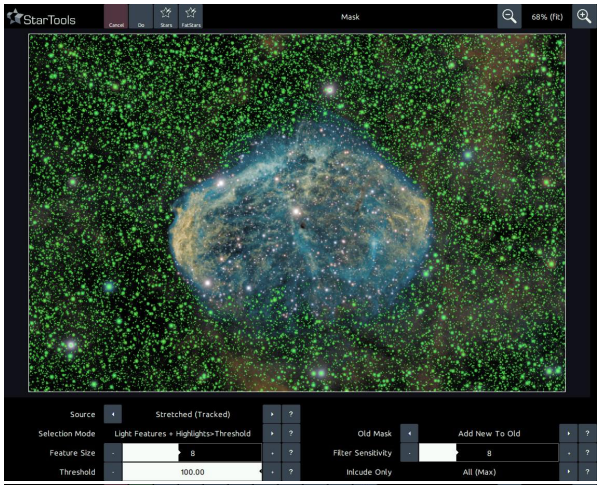

5. We now have a mask that is less aggressive in the area we specified earlier and more aggressive elsewhere

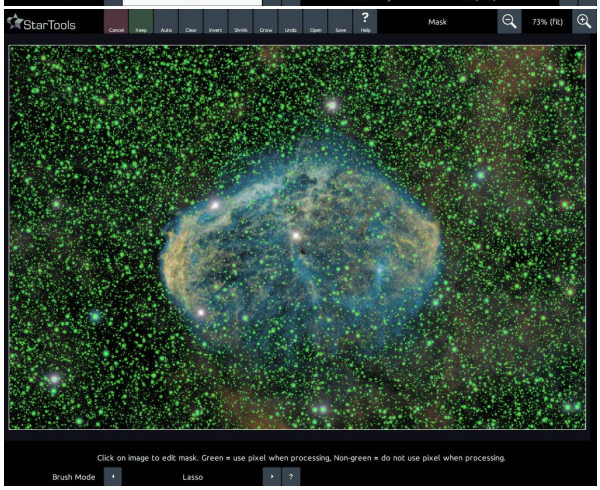

Wrap Up

Description

- To modify the way stars look in the image - by shrinking bloated stars, recover the color and manipulate the core and halo.
- This module can be used to control bloated stars (with the Dim or Tighten algorithm) that the 'Decon' module had trouble with - possibly because they were overexposed.
- It changes the star sizes.
- It complements the 'Repair' module which is used for severely deformed stars.

When to use

- After the Color module but before Tracking is turned off.
- Be careful when shrinking stars on widefield images. There is often not much background for the algorithm to draw from.
- Can be used more than once.

What result to look for

If the Mode is set to 'Tighten'

- Stars are more concentrated and have improved color.

After Use

Consider using other modules that require Tracking to be on - like 'Filter', 'Entropy' and 'SuperStructure'.

Description of Controls:

Mask & Auto Mask

For general instructions on using masks see Mask

- The use of a star mask is highly recommended.
- The Auto-generated star mask should catch all major stars but leave the smaller ones.
- You may want to exclude areas from being affected by Shrink - or to be less affected by using a less aggressive mask in places.
- To select other stars as well you can run the mask module more than once by setting 'Old Mask' to 'Add New to Old' - see 'Creating a suitable mask' here

Presets

- Tighten
  - Reduces star size, pulling halos into the core.
  - Sets Mode to Tighten, Iterations to 10
- Dim
  - reducess the luminosity of stars
  - Sets Mode to Dim, Iterations to 10, Halo extend to 2 pixels
- Classic
  - Mimics the legacy module's 'Shrink' algorithm
  - Sets Mode to Dim, Iterations to 1, Deringing to Off, Color taming to 0 pixels, Regularization to 0.00.
- Un-glow
  - Reduces Star Halos
  - Sets Un-glow Strength to 35% and Un-glow Kernel to 30.0 pixels

Iterations

Sets the strength of the effect.

- Default is 10 ('Tighten' preset), 7('Dim' preset), 1 ('Classic', 'Unglow' presets) .
- Range is 0 to 30.

Mode

Defines the Mode the 'Shrink' module operates in:

- Dim - Dims the stars luminosity
- Tighten - Reduces the star halos, drawing them in to the star. Does not affect the star core.
- Default is 'Tighten' ('Tighten' preset), 'Dim' ('Dim', 'Classic', 'Unglow' presets)

De-ringing

Mitigates ringing artifacts caused by sharpening or deconvolution

- Default is 10 pixel ('Tighten', 'Dim' presets), 'Off' ('Classic', 'Unglow' presets)
- Range is 0 to 10 pixels.

Color Taming

This setting controls how the star progressively takes on the color of the surroundings. Improves embedding of stars in the image and mitigates rugged edges.

- Default is 2 pixel ('Tighten' preset), 40 ('Dim' preset), 0 pixel ('Classic', 'Unglow' presets).
- Range is 0 to 40 pixels.

Regularization

This controls how much each iteration of the process is dampened by. This is used to avoid pitting or stringing artifacts.

- Reduce this setting if any pitting or stringing artifacts appear around stars.
- Default is 0.85 ('Tighten', 'Dim' presets), 0.00 ('Classic', 'Unglow' presets)
- Range is 0.00 to 1.00.

Halo Extend

Has the effect of temporarily growing the selected mask - so including more of the stars surroundings. Use this if rings occur in or around the stars when using a star mask.

- Default is 1 pixel ('Tighten', 'Classic', 'Unglow' presets). 2 pixel ('Dim' preset)
- Range is 0 to 10 pixels.

Un-glow Strength

Un-glow attempts to reduce the halos around bright stars. This parameter controls the strength of that effect.

- Default is 'Off' ('Tighten', 'Dim', 'Classic' presets), 35% ('Unglow' preset)
- Range is 0% to 100%.

Un-glow Kernel

This defines the width of the halos that the Un-glow feature should detect and control.

- Default is 'Off' ('Tighten' preset), 30.0 pixel ('Dim', 'Classic', 'Unglow' presets)
- Range is 'Off' (0.0) to 30.0 pixels.

'Filter': Feature Manipulation by Color

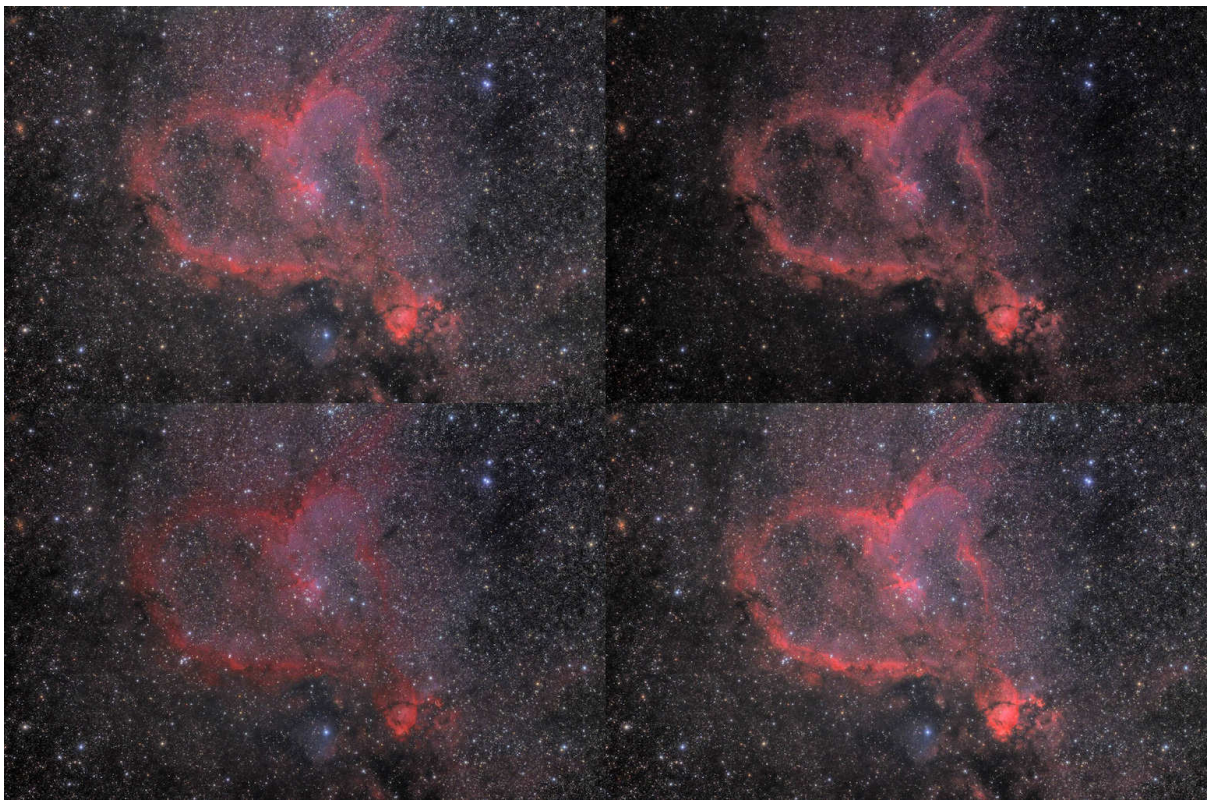

A simple demonstration of the 'Filter' module. All images generated by selecting a filter mode and clicking twice on a part of the red H-alpha region. Top Left: Source. Top Right: Pass. Bottom Left: Reject. Bottom Right: Nudge.

The 'Filter' module allows for the modification of features in the image by their color by simply clicking on them. It's as close to a post-capture color filter wheel as you can get.

'Filter' can be used to bring out detail of a specific color (such as faint Ha, Hb, OIII or S2 details), remove artifacts (such as halos, chromatic aberration) or isolate specific features. It functions as an interactive color filter.

The 'Filter' module is the result of the observation that many interesting features and objects in outer space have distinct colors, owing to their chemical make up and associated emission lines. Thanks to the 'Color' Constancy feature in the 'Color' module, colors still tend to correlate well to the original emission lines and features, despite any wideband RGB filtering and compositing. The 'Filter' module was written to capitalize on this observation and allow for intuitive detail enhancement by simply clicking different parts of the image with a specific color.

Usage

A 'Filter Mode' parameter selects the mode of the filter. Available modes are:

- 'Conservative Nudge'; this mode boosts the selected signal linearly, but only if the boost would not yield any overexposure
- 'Nudge (Screen)'; this mode boosts the selected signal by using a Screen overlay operation, boosting the signal non-linearly.
- 'Pass'; only lets through the selected signal and attenuates all other signal.
- 'Reject'; blocks the selected signal, leaving all other signal intact.
- 'Fringe Killer'; Draws color from neighboring pixels that are not masked and gives these colors to masked pixels. Note that this mode requires a mask to be set.
- 'Saturate Visual H-alpha'; saturates red coloring. In this mode, the user must click on the coloring that is to be preserved while the H-alpha is boosted.
- 'Saturate Visual H-beta/O-III'; saturates cyan coloring. In this mode, the user must click on the coloring that is to be preserved while the H-beta/O-III is boosted

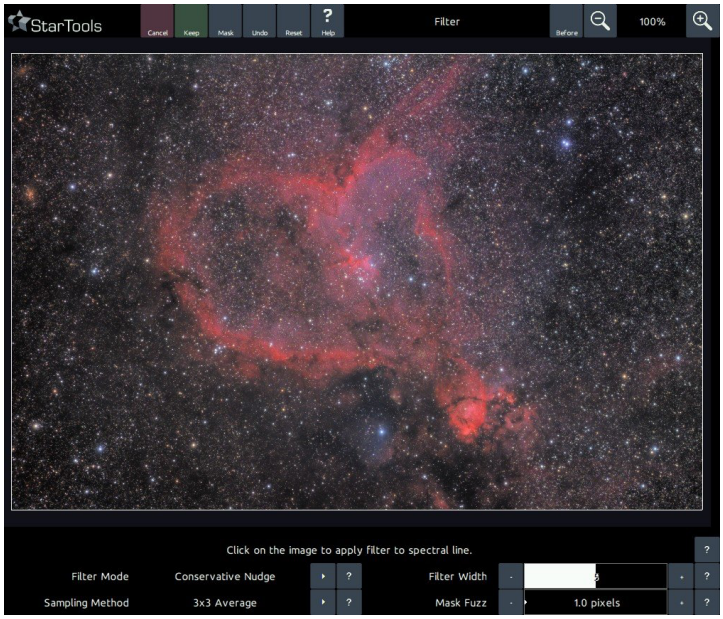

Operating the module is as easy as picking a 'Filter Mode' and clicking on areas you wish to affect with your chosen filter.

The 'Filter Width' parameter specify the responsiveness of neighboring colors in the spectrum. A small 'Filter Width' will see the module only modify areas with a very precise match in color to the area selected, while a larger 'Filter Width' will see the module progressively modify areas that deviate in color from the selected area as well.

The 'Sampling Method' mode selects how a click on the image samples the image. The '3x3 Average' mode samples a 3x3 area around the clicked pixel and uses the resulting 9-pixel average as the input color. The 'Single Pixel' mode, samples only the precise pixel that was clicked.

Finally, a 'Mask Fuzz' parameter allows for the result to progressively masked in cases where a mask is set.

Mitigating chromatic aberration

The 'Filter' module's 'Fringe Killer' mode is an easy and very effective way to remove unsightly blue and purple halos caused by chromatic aberration.

The 'Filter' module's 'Fringe Killer' mode is an easy and very effective way to remove unsightly blue and purple halos caused by chromatic aberration.
Simply put the offending stars, including their halos in a mask (one can be automatically generated from within the 'Filter' module, by clicking Mask, Auto, Stars or FatStars, Do, Keep). Next click a few times on different parts of the purple or blue halos and they will slowly disappear with each click.

Wrap Up

Description

- To modify features in the image on the basis of their color.
- The 'Filter' module can be used to bring out detail of a specific color (such as faint Ha, Hb, OIII or S2 details), remove artifacts (such as halos, chromatic aberration) or isolate specific features.
- It functions as an interactive color filter.

When to use

- This is best used after the 'Color' module and before the 'Denoise' module (Tracking off).

What result to look for

- The selected color range should be enhanced or reduced as specified by the 'Filter' Mode setting.
- Only the areas where the mask is set (green) should be changed.

After Use

- If used after the 'Color' module then the usual next step is either 'Entropy' module or to turn off Tracking and use the 'Denoise' module.

Description of Controls

Mask

For general instructions on using masks see Mask.

- Select the elements that you want to be modified. The action will be applied to areas of the selected color range within the masked areas.

'Filter' Mode

Controls what a filter will do with the selected spectral line.
Values are:

- Conservative Nudge - make the selected color range more pronounced as long as it doesn't cause over-exposure (clipping).
- Nudge - make the selected color range more pronounced - even if it means over-exposure.

- Pass - keep the selected color range and reduce all other parts of the spectrum.
- Reject - reduce the selected spectral band and keep all other parts of the spectrum.
- Fringe Killer - tries to remove halos and fringes around stars that have the selected color range.
- Saturate H-Alpha - tries to make the Hydrogen Alpha (red) spectral line more prominent.
- Saturate H-Beta - tries to make the Hydrogen Beta (cyan) spectral line more prominent.
- Default is 'Conservative Nudge'.

Sampling Method

The central color to filter is selected by clicking a pixel in the image. This setting controls how the color range will be identified from this:

- 3 x 3 Average - averages the color of a 3x3 pixel block around the selected pixel.
- Single Pixel - uses the color of the selected pixel only.
- Default is '3x3 Average'.

'Filter' Width

Controls how much of neighboring parts of the color spectrum that the selected action will be applied to.

- Values can range from 0 to 10 with 10 being the widest color range.
- The width does not have a sharp cutoff - there is a peak and the intensity of the effect fades away as you get further from the peak. The fading is more gradual at higher 'Filter' Width values.
- Default is 5.
- Range is 0 to 10.

Mask Fuzz

- If a mask is used, Mask Fuzz controls the blending of the transition between masked and non-masked parts of the image.
- Using this control will allow smooth transitions between filtered and unfiltered parts of the image.
- Default is 1.0 pixels.
- Range is 1.0 to 31.0 pixels.

'Entropy': Inter-channel 'Entropy'-driven Detail Enhancement

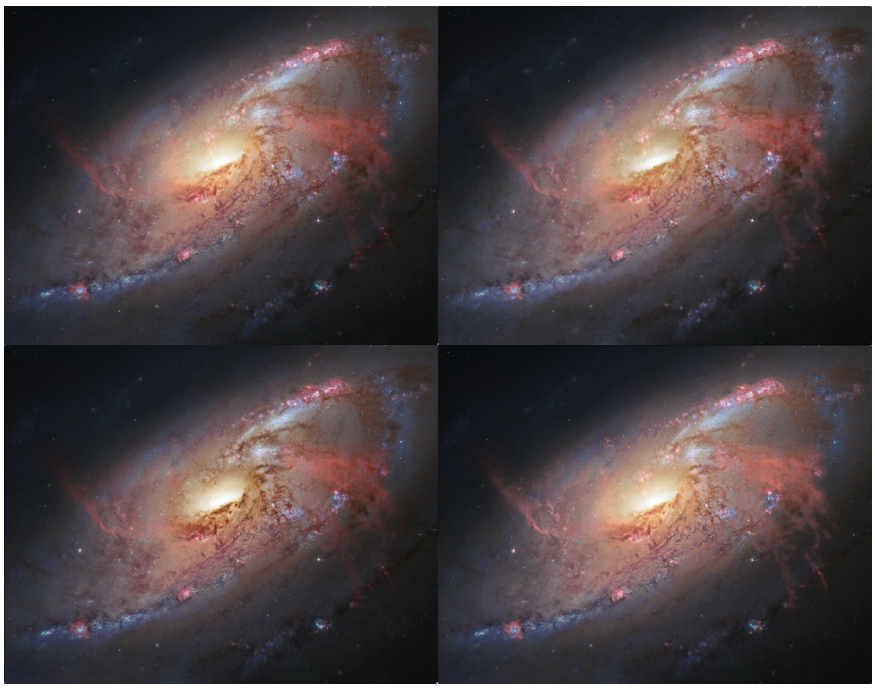

Top left: original, top right: all channel optimization, bottom left: blue channel optimization, bottom right: red + green optimization. Original image courtesy of NASA, ESA, the Hubble Heritage Team (STScI/AURA), and R. Gendler (for the Hubble Heritage Team). Acknowledgment: J. GaBany

The 'Entropy' module is a novel module that enhances detail in your image, using latent detail cues in the color information of your dataset.

The 'Entropy' module exploits the same basic premise as the 'Filter' module; that is, the observation that many interesting features and objects in outer space have distinct colors, owing to their chemical make up and associated emission lines. This correlation become 100% when considering a narrowband composite, where each channel truly is made up of data from distinct parts of the spectrum.

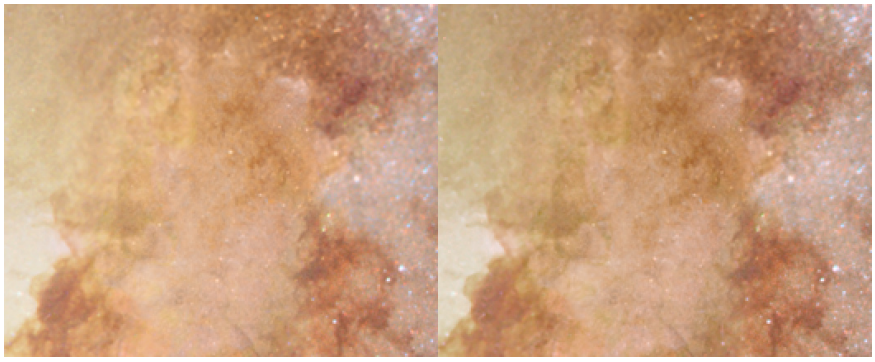

200% zoom detail from NASA 106 image. Left: original, right: 'Entropy' module processed image. Very subtle difference in clarity and contrast can be spotted.

The 'Entropy' module works by evaluating entropy (a measure of 'busyness' or 'randomness') as a proxy for detail. It does so on a local level in each color channel for each pixel. Once this measure has been established for each pixel, the individual channel's contribution to luminance for each pixel is re-weighted in CIELab space to better reflect the contribution of visible detail in that channel.

The result is that the luminance contribution of a channel with less detail in a particular area is attenuated. Conversely, the luminance contribution of a channel with more detail in a particular area is boosted. Overall, this has the effect of accentuating latent structures and detail in a very natural manner. Operating entirely in CIELab space means that, psychovisually, there is no change in color, only brightness.

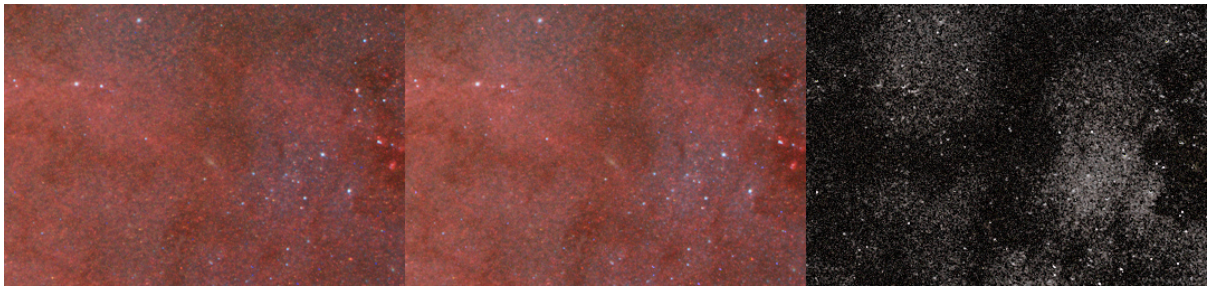

Subtle large-scale structure enhancement. Left: original, middle: 'Entropy' module processed, right: difference map. Of course, the strength of the effect is wholly decided by the user.

The above attributes make the 'Entropy' module an extremely powerful tool for narrowband composites in particular.

The 'Entropy' module is effective both on already processed images, as well as Tracked datasets. The module is available as of StarTools 1.5.

Usage

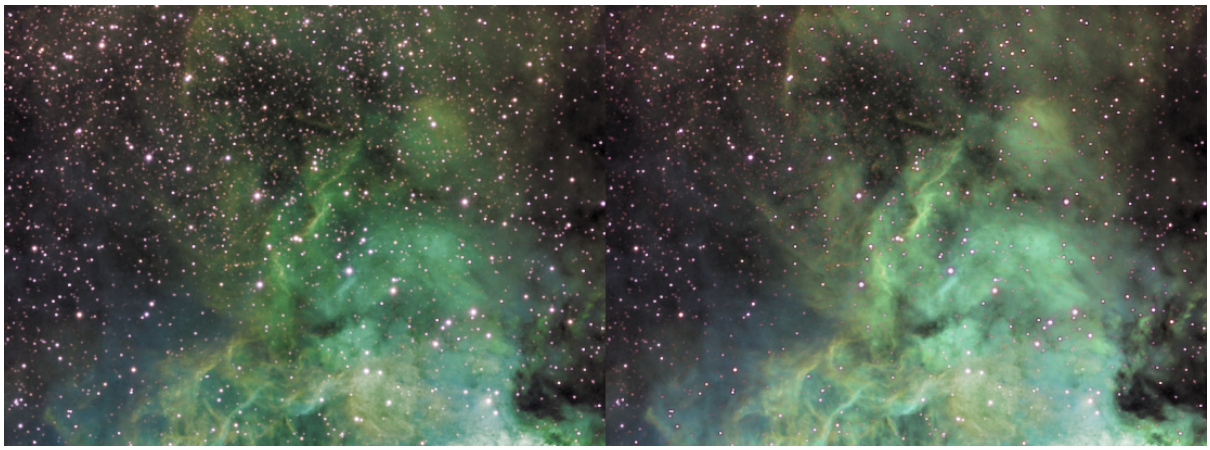

The 'Entropy' module makes for a fantastic narrowband manipulation tool; here it was used to effortlessly boost the prevalence of O-III emissions in a SHO-mapped image.

The 'Entropy' module is very flexible in its image presentation. To start using the 'Entropy' module, an entropy map needs to be generated by clicking the 'Do' button. This map's resolution/accuracy can be chosen by using the 'Resolution' parameter. The 'Medium' resolution is sufficient in most cases.

For the entropy module to be able to identify detail, the dataset should ideally be of an image-filling object or scene.

After obtaining a suitable entropy map, the other parameters can be tweaked in real-time;

- The 'Strength' parameter governs the overall strength of the boost or attenuation of luminance. Overdriving the 'Strength' parameter too much may make channel transitions too visible. In this case you may wish to pull back, or increase the 'Midtone Pull 'Filter" size to achieve a smoother blend.
- The 'Dark/Light Enhance' parameter enables you to choose the balance between darkening and brightening of areas in the image. To only brighten the image (for example if you wish to bring out faint H-alpha, but nothing else), set this

parameter to 0%/100%. To only darken the image (for example to better show a bright DSO core) bring the balance closer to 100%/0%.

- The 'Channel Selection' parameter allows you to only target certain channels. For example, if you wish to enhance S-II more visible in a Hubble-palette image, set this parameter to red (to which S-II should is mapped). S-II will now be boosted, and H-alpha and O-III will be pushed back where needed to aid S-II's contrast. If you wish to avoid the other channels being pushed back, simply set the 'Dark/Light Enhance' to 0/100%.
- The 'Midtone Pull Filter' and 'Midtone Pull Strength' parameters, assist in keeping any changes in the brightness of your image confined to the area where they are most effective and visible; the midtones. This feature can be turned off by setting 'Midtone Pull Strength' to 0%. When on, the filter selectively accepts or rejects changes to pixels, based on whether they are close to half unity (e.g. neutral gray) or not. This feature works analogous to creating a 'HDR' composite from different exposure times. The transition boundaries between accepted and rejected pixels are smoothened out by increasing the 'Midtone Pull Filter' parameter.

Wrap up

Description

- To enhance the local detail in the image - using information in the color within your data. Works with both narrowband and broadband data.
- It uses the idea that interesting features have distinct colors - derived from the emission lines of what they are made of, or the light they reflect.
- This idea can be well exploited when using narrowband data where each channel represents a different part of the spectrum.
- The module identifies local detail in each band and locally adjusts the weighting of the bands luminance contribution to highlight the detail.
- By default the 'Entropy' module works on all bands and darkens and brightens to bring out the detail evenly.
- With Tracking on, noise will be tracked and taken care of when 'Denoise' is done when Tracking is switched off.
- With Tracking off, noise reduction is applied within the module itself.

When to use

- When you want to improve the detail in the image - and highlight the changes in color.
- When the object fills much of the image.
- Very effective with narrowband data - also works with broadband data from OSC and DSLR cameras.
- Use after using the 'Color' module. The 'Entropy' module will be greyed out until the 'Color' module has been applied.
- With Tracking on, prior color balancing does not affect the result.
- Works on both unstretched and stretched data (e.g. .JPG data)

What result to look for

- Increased detail in the image where the color changes.

After Use

You may want to use retry the 'Color' Module to adjust the color of the recovered detail.
Then move on to Wavelet 'Denoise'.

Description Of Controls

Presets

Define a preset areas of the spectrum you want to use to enhance the detail.
Many of the presets assume SHO (e.g. Hubble Palette) RGB mapping.

- Default(All) - Uses all of the color data to enhance the detail
- SHO SII - When using SHO narrowband data - Enhances the detail identified using the SII (red) data.
- SHO Ha - When using SHO narrowband data - Enhances the detail identified using the Ha (green) data.
- SHO OIII - When using SHO narrowband data - Enhances the OIII detail identified using the OIII (blue) data.
- Visual Ha - When using broadband data - Enhances the detail identified using the Ha (red) data.

Resolution

Specifies the 'Entropy' mapping resolution of the analysis. Higher resolution takes longer to process.

- Low - If you are trying to highlight larger, more general, variations in the image.
- Medium - If you are trying to highlight medium sized detail.
- High - If you are trying to highlight small details in the image.
- Default is Medium.

Channel Selection

Selects the channels which are used in enhancing the detail. Hubble palette (HSO) is assumed for narrow band, or full visual spectrum.

- All - Enhances the detail using all the color data.
- Red (SHO SII, Vis.HA) - Enhances the detail identified using red data (e.g. SHO SII or Visible Ha)
- Red+Blue (SHO SII+OIII,Vis.HII) - Enhances the detail identified using the red and blue data (e.g. SHO SII & OIII or Visible HII)
- Red+Green (SHO SII+Ha) - Enhances the detail identified using the red and green data (e.g. SII & Ha data)
- Green (SHO Ha) - Enhances the detail identified using the green data (e.g. SHO Ha data)
- Green+Blue (Vis.O-III, SHO Ha+OIII) - Enhances the detail identified using the green and blue data (e.g. SHO Ha and OIII or visible OIII)
- Blue (Vis.Reflection, SHO OIII) - Enhances the detail identified using the blue data (e.g. SHO OIII or visible reflection nebulae)
- Default is 'All' - Uses all channels to identify the detail to enhance.

Strength

- The overall strength of the increase or decrease of luminance.
- Increasing the Strength parameter increases the effect.
- Be careful not to introduce artifacts with the higher values under certain circumstances.
- Default is 100%.
- Range is 0% to 1,000%.

Dark/Light Enhance

Selects the balance between brightening the selected channels and darkening areas in order to enhance the detail locally.

- Minimum 0%/100% - Just brightens to highlight the detail. Use this to just enhance faint nebulosity without darkening the surrounding area.
- Maximum 100%/0% - Just darkens to highlight the detail.
- Default is 50%/50% - Both darkens and brightens evenly to highlight the detail.

Midtone Pull Strength

Sets the amount of priority given to the midtones. How much luminance should be pulled towards the midtones.

- 0% Turns off the midtone filter.
- Default is 50%.
- Range is 0% to 100%.

Midtone Pull 'Filter'

Defines the kernel size of the filter that blends the enhanced and non-enhanced detail to give a smooth transition.

- Default is 20.0 Pixels.
- Range is 1.0 pixel to 51.0 pixels

Super Structure:
Global Light Diffraction Remodeling of Large Scale Structures

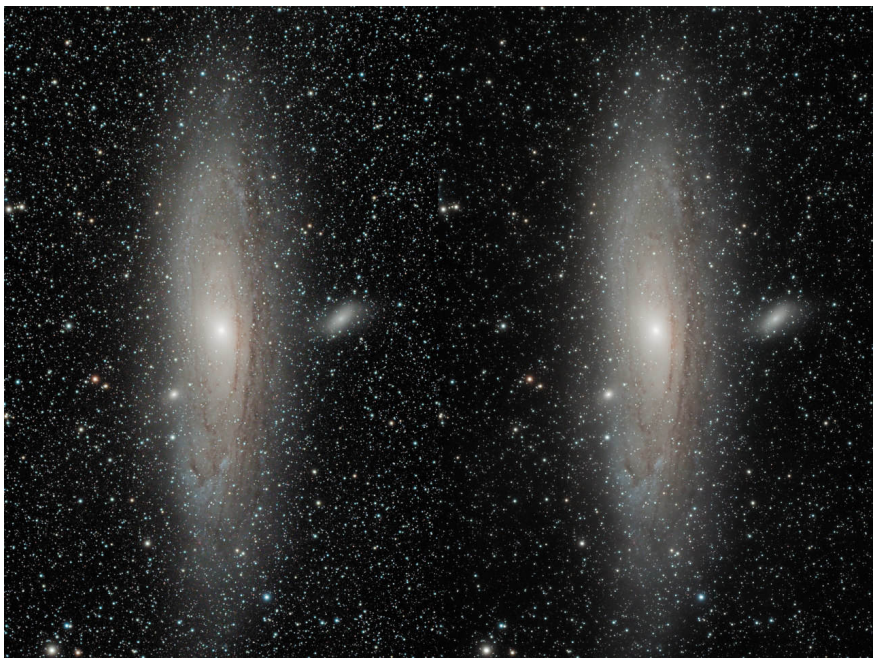

The Isolate preset of Super Structure at work, 'pushing back' a busy star field and refocusing attention on the nebulosity.

The Super Structure module brings back 'life' into an image by remodeling uniform light diffraction, helping larger scale structures such as nebulae and galaxies stand out and (re)take center stage.

Throughout the various processing stages, light diffraction (a subtle 'glow' of very bright objects due to lens or mirror diffraction) tends to be distorted and suppressed through the various ways dynamic range is manipulated. This can sometimes leave an image 'flat' and 'lifeless'. The Super Structure module attempts to restore the effects of uniform light diffraction by an optical system, throughout a processed image. It does so by means of modeling an Airy disk pattern and re-calculating what the image would look like if it were diffracted by this pattern. The resulting model is then used to modulate or enhance the source image in various ways. The resulting output image tends to have a re-established natural sense of depth and ambiance, with better visible super structures.

For example, the Super Structure module's 'Isolate' preset, when applied to the whole image, is particularly adept at pushing back busy star fields and noisy backgrounds, refocusing the viewer's attention to the larger scale structures. As such it is a very powerful, yet easy to use tool to radically change the feel of an image.

The Super Structure module may additionally be used locally by means of a mask. In this case the Super Structure module can be used to isolate objects in an image and lift them from an otherwise noisy background. By having the Super Structure module augment an object's super-structure, faint objects that were otherwise unsalvageable can be made to stand out from the background. Please note that, depending on the nature of the used selective mask, the super structures introduced by using the Super Structure module in this particular way with a selective mask, should be regarded as an educated guess rather than documentary detail, and technically falls outside of the realm of photography.

Usage

As with most modules in StarTools, the Super Structure module comes with a number of presets:

- 'DimSmall' pushes back anything that is not a super structure while retaining energy allocated to super structures. Overall image brightness is compensated for.
- 'Brighten' brightens detected super structures
- 'Isolate' is similar to the 'DimSmall' preset, however does not compensate for lost energy (image brightness).
- 'AiryOnly' shows the AiryDisk model only for fine tuning or use in other ways.
- 'Saturate' saturates the colors of detected super structures.

Going beyond the presets, very detailed adjustments can be made, starting with the 'Glow Threshold' parameter. This parameter is determines how bright a pixel needs to be before it is considered for diffraction by the Airy disk diffraction model.

To view just the model that Super Structure is using to enhance the image, the 'Output Glow Only' parameter can be set to 'Yes'. Optionally this output can be used to manipulate the image later using the 'Layer' module, or in a separate application.

The 'Strength' parameter governs the overall strength of the effect.

The 'Inherit Brightness, Color' parameter determines whether brightness or color information is inherited (and thus unchanged) from the source image.

The 'Saturation' parameter controls the color saturation of the output model (viewable by setting 'Output Glow Only' to 'Yes'), before it is applied to the source image to generate the final output. This parameter can be quite effective for enhancing the color of nebulosity.

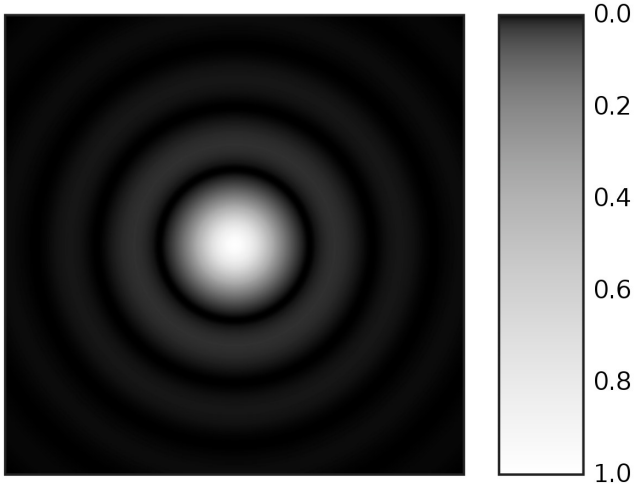

A computer-generated image of an Airy disk. The grayscale intensities have been adjusted to enhance the brightness of the outer rings of the Airy pattern. Source: Wikipedia.

The 'Detail Preservation' parameter selects the detail preservation algorithm the Super Structure module should use to merge the model with the source image to produce the output image:

- Off - does not attempt to preserve any detail.
- Min Distance to 1/2 Unity - uses the pixel that is closest to half unity (e.g. perfect gray).
- Max 'Contrast' - uses whatever pixel maximizes contrast with its neighboring pixels.
- Linear Brightness Mask - uses a brightness mask that progressively masks-out brighter values until it uses the original values instead.
- Linear Brightness Mask Darken - uses a brightness mask that progressively masks out brighter values. Only pixels that are darker than the original image are kept.

The 'Detail Preservation Radius' sets a filter radius that is used for smoothly blending processed and non-processed pixels, according to the algorithm specified by the 'Detail Preservation' parameter.

The 'Compositing Algorithm' parameter defines how the calculated diffraction model is to be generally combined with the original image:

- Screen - works like projecting two images on the same screen.
- Power of Inverse - Power of Inversed Pixels (PIP) function.
- Multiply, Gamma Correct - multiplies foreground and background and then takes the square root.
- Multiply, 2x Gamma Correct - similar to 'Multiply, Gamma Correct' but doubles the Gamma Correction.

The 'Airy Disk Sampling' parameter controls the accuracy of the point spread function (PSF) that describes the diffraction model (an Airy disk).

- Default is 128 x 128 pixels. Range is 128 x 128, 256 x 256, 512 x 512 pixels.
- Increasing this value will give a more accurate simulation but will take longer.

The 'Airy Disk Radius' parameter sets the radius of the Airy disk point spread function (PSF) that is used to diffract the light. Just like in nature, you may spot some (very) subtle rings around the stars after processing. The way this looks can be adjusted using this setting.

Finally, as with most modules in StarTools that employ masks, a 'Mask Fuzz' parameter is available to smoothly blend the transition between masked and non-masked pixels.

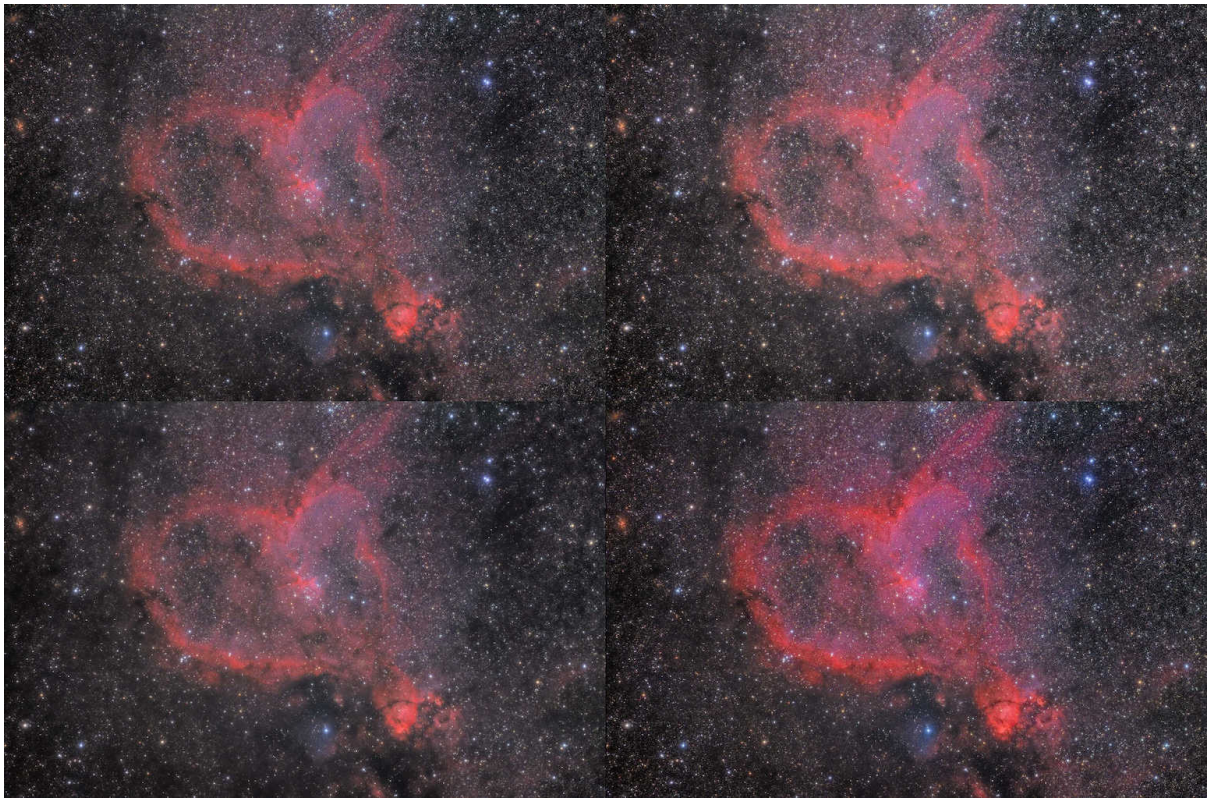

3 examples of the Super Structure module presets manipulating the visibility of large scale structures. Top left: original. Top right: 'Brighten' preset. Bottom left: 'DimSmall' preset. Bottom right: 'Saturate' preset.

Wrap Up

Description

- To process large scale structures in the image (e.g. galaxies & large scale nebulosity) independent of the rest of the image to allow effects such as de-emphasizing busy star fields, emphasizing large structures or enhancing their color.
- To bring back life into an image by remodeling uniform light diffraction.
- Removes the flatness that heavy processing can cause bringing back some of the 3d effects.
- To lift objects from very noisy data or dense star fields.

When to use

- Towards the end of the workflow - usually just before using the 'Color' module or the final 'Denoise'.

What result to look for

- 'Isolate' Preset - objects should be lifted from the noise and background stars. Watch for halos around bright stars.
- 'DimSmall' Preset – objects should be lifted from the noise and background stars with overall brightness remaining. Watch for halos around bright stars.
- 'Brighten' - Large scale objects should be brighter
- 'Saturate' - Large scale objects should have stronger Saturation

After Use

- Often followed by turning Tracking Off & the 'Denoise' module.

Description of Controls

Mask

For general instructions on using mask see Mask.

- A mask can be used to selectively 'add life'.
- Use of the mask will compromise the 'documentary' value of the final image. If this is of concern only use a mask derived purely from the image (e.g. Auto Star Mask with optional Invert)
- The mask can be used in conjunction with the 'Isolate' or 'DimSmall' presets to select the object(s) to be lifted.

Presets

- 'Isolate' – Makes large scale structures stand out and pushes back stars.
- 'DimSmall' – makes large scale structures stand out and pushes back stars while compensating for overall brightness loss.
- 'Brighten' - brightness of large scale structures will be increased
- 'Saturate' - saturation of large scale structures will be increased
- 'Airy Only' - displays the airy disk modeling only

Gamma

Specifies gamma correction to the model before compositing.

- Range is 0.00 to 10.00.
- Default is 0.50 ('DimSmall', 'Brighten', 'Isolate', 'Saturate' presets), 1.00 (AiryOnly).

Compositing Algorithm

This setting defines how the calculated super structure is combined with the original image:

- None - output super structure only
- Screen - works like projecting two images on the same screen.
- Power of Inverse - Combination based on Power of Inversed Pixels (PIP) function.

- Multiply, Gamma Correct - multiplies foreground and background and then takes the square root.
- Multiply, 2x Gamma Correct - as above but doubles the Gamma Correction.
- Default is 'Multiply, Gamma Correct' ('DimSmall', 'Isolate', presets), 'Power of Inverse' ('Brighten' preset), 'Screen' ('Saturate' preset), 'None' ('AiryOnly' preset)

Brightness, 'Color'

Sets whether brightness or color or both will be processed.

- Process Both
- Only 'Color'
- Only Brightness
- Default is 'Process Both' ('DimSmall', 'Brighten', 'AiryOnly' presets), 'Only Brightness' ('Isolate' preset) 'Only 'Color'' ('Saturate' preset).

Strength

Controls the overall strength of the effect.

- Range is 0% to 500%.
- Default is 100% ('DimSmall', 'Isolate', 'Saturate' 'AiryOnly' presets), 75% ('Brighten' preset)

Saturation

Sets the color saturation level of the extracted super structure before it is combined with the original.

- Default is 50% ('DimSmall' preset), 100% ('Brighten', 'Isolate', 'AiryOnly' presets), 200% ('Saturate' preset).
- Range is 0% to 500%.
- Change this to adjust the amount of 'glow'.

Brightness Retention

Specifies the algorithm compensating brightness discrepancies between input and output.

- Off - does not compensate
- Local Median - corrects locally re-applying the median for the pixel's area
- Global Mode Align, Darken Only - performs a gamma correction aligning histogram peaks of input and output. Then only the darker pixel of both is chosen for output.
- Default is 'Local Median' (DimSmall preset , 'Off' ('Brighten', 'Isolate', 'Saturate', 'AiryOnly' presets)

Detail Preservation

Sets the way that 'SuperStructure' preserves the detail in the parts of the image to be brightened.

- Off - does not preserve any detail.
- Min Distance to 1/2 Unity - uses the pixel that is closest to half unity.
- Linear Brightness Mask - uses a brightness mask that progressively masks out brighter values and uses the original values instead.
- Linear Brightness Mask Darken - uses a brightness mask that progressively masks out brighter values. Only pixels that are darker than the original image are kept.
- Default is 'Linear Brightness Mask Darken' ('DimSmall', preset), 'Linear Brightness Mask' ('Isolate', 'Saturate' presets), 'Min Distance to 1/2 Unity' ('Brighten' preset), 'Off' 'AiryOnly' preset.

Airy Disk Radius

Sets the radius of the Airy disk point spread function (PSF) that is used to diffract the light.

- Default is 50% ('DimSmall', 'Brighten', 'Isolate', 'Saturate', 'AiryOnly' presets)
- Range is 0% to 100%

- There are subtle rings around the stars after processing. The way this looks can be adjusted using this setting.

Detail Preservation Radius

Sets a filter radius that is used for smoothly blending processed and non-processed pixels, according to the algorithm specified. This setting is only active for 'Detail Preservation' set to 'Min Distance to 1/2 Unity' or 'Linear Brightness Mask Darken'

- Default is 40.0 pixels ('Brighten', 'DimSmall', presets).
- Range is 1.0 pixel to 50.0 pixels.

Mask Fuzz

When using a mask, this sets the smoothness of the transition between masked and non-masked elements by applying a gaussian blur to the mask.

- Default is 1.0 pixel ('DimSmall', 'Brighten', 'Isolate', 'Saturate', 'AiryOnly' presets).
- Range is 1.0 to 40.0 pixels.
- Increase to make the transition smoother - values of 10 pixels are common.

NBAccent: Adding Narrowband Accents to Visual Spectrum Datasets

The NBAccent module adds narrowband accents to visual spectrum datasets. Visibility of HII areas in this image of M33 is greatly enhanced, while visual spectrum coloring is largely maintained. Top: before NBAccent module, Bottom: after NBAccent module.

Adding narrowband accents to visual spectrum datasets has traditionally been a daunting, difficult and laborious process, involving multiple workflows. The NBAccent module is a powerful module that starts its work as soon as your load your data in the Compose module. Crucially it adds only a single, easy step to an otherwise standard workflow, while yielding superior results in terms of color fidelity/preservation.
By making narrowband accents an integral part of the complete workflow and signal path, results are replicable, predictable and fully tracked by StarTools' unique signal evolution Tracking engine, yielding perfect noise reduction every time.

Usage

Activating the NBAccent module functionality, starts with importing a suitable narrowband dataset via the Compose module. The Compose module will extract the relevant channels from the dataset you give it, as directed by its 'NB Accents Type' parameter.

The narrowband dataset is processed in parallel during your workflow; the Bin, Crop, Mirror, Rotate and - most notably - Wipe modules all operate on the narrowband accent dataset in parallel as you process the main luminance (and optionally chrominance) signal.

Understanding the module's purpose and use case

There are many different ways and techniques of incorporating narrowband data into your workflow. Which method is suitable or desirable, depends on the object, the availability of datasets/bands, and the quality of those available dataset.

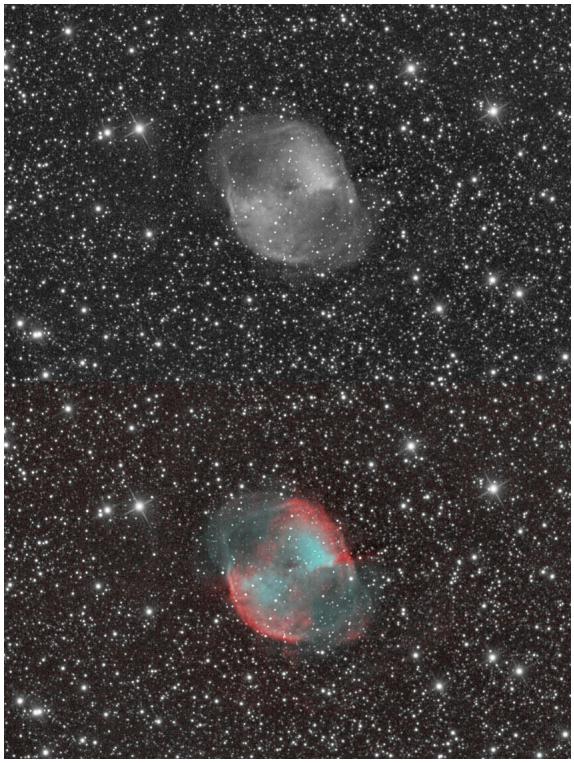

Through the power of StarTools' three-way signal separation (luminance, chrominance, and narrowband accents), the module can also be deployed for more esoteric uses. In this example it is used to endow a luminance ("L") dataset with narrowband accents acquired through a popular duoband filter. No chrominance (color) dataset was used here. Notice that, as a result, stars remain perfectly white.

The NBAccent module was specifically designed for the most difficult compositing use case: that of using narrowband as means to accentuate detail in a visual spectrum 'master' dataset. In other words, in this use case, the narrowband is used to support, enhance and accentuate small(er) aspects of the final image, rather than as a basis for the initial signal luminance/detail or chrominance/coloring itself. This is a subtle, but tremendously important and consequential distinction.

As such, the narrowband accent dataset is processed entirely independent of the luminance and chrominance signal of the 'master' dataset; it sole purpose is to accentuate detail from the 'master' (luminance/chrominance) dataset through careful - but deliberate - local brightness and/or color manipulation.

If you wish to use the narrowband signal as luminance or chrominance itself, rather than for accentuating luminance or chrominance, then the NBAccent module will not apply, and you should use the Compose module to load your narrowband as luminance and/or chrominance instead.

Given the module's use case, it is best invoked late in the processing flow, after the Color module.

Examples of use cases for the NBAccent module are:

- accentuating HII areas in galaxies (by passing it a Hydrogen-alpha dataset) such as M31, M33
- accentuating or adding large scale background nebulosity to already rich visual spectrum widefield renditions of HII areas such as NGC 7635, M16
- accentuating or better resolving intricate features in objects such as planetary nebula

Ideal datasets for augmenting visual spectrum (mono or colour) datasets are Ha datasets, O-III datasets, Ha+O-III datasets or datasets from the popular duo/tri/quadband filters for OSCs and DSLRs such as the Optolong L-Extreme, the STC Duo , the ZWO Duo-Band and other similar filters with narrow spectrum responses.

Stage 1: Signal stretch and contribution calibration

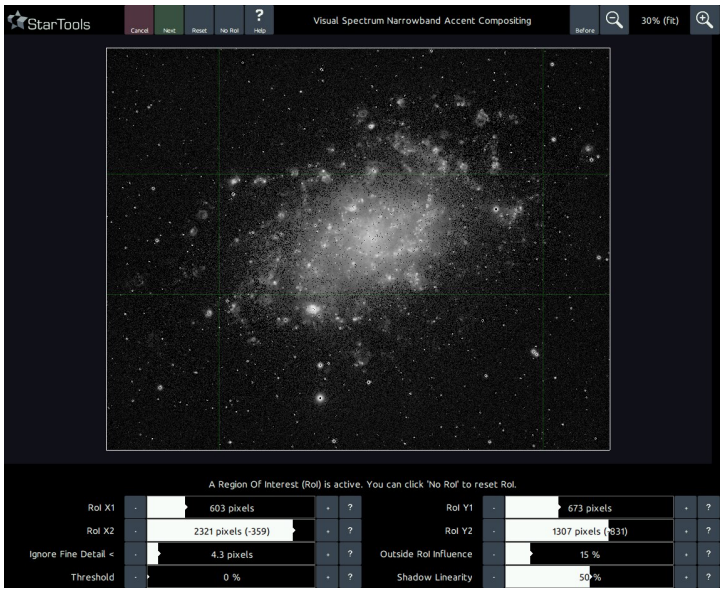

In this setup stage, pixels that will be affected in the visual spectrum image will show narrowband signal, while pixels that will not be affected are clipped to black. This will allow you to gauge how the image will be transformed by the parameters you choose here.

The first screen allows you to fine control which areas will receive narrowband enhancement. The procedure and, hence, interface is closely related to the AutoDev module. Familiarizing with AutoDev is key to achieving good results with StarTools, and being able to use it effectively is a prerequisite to being able to use the NBAccent module.

One notable difference compared to AutoDev, is the way the stretched narrowband data is presented; areas that will not be considered for the final composite, will be clipped to black. Areas that will be considered in the final composite, will appear sretched as normal. The other difference from the AutoDev module, is the removal of the 'Detector Gamma' parameter and its replacement by the 'Threshold' parameter; this parameter allows for intentional clipping of the narrowband image, for example to avoid any background imperfections being added to the final composite. It is important to note that this parameter should be used as a last resort only (for example if the narrowband accent data is of exceedingly poor quality) as it is a very crude tool that will inevitably destroy faint signal.

It is important to understand that the signal as show during this first stage, is merely signal that is up for consideration by the second stage. Its inclusion is still contingent on other parameters and filters in the second stage. In other words, during this first stage, you should merely ensure that, whichever signal is visible, is actual useful narrowband signal, and not the result of background imperfections or other artificial sources.

For your convenience, the NBAccent module will, by default, use the same Region of Interest that was specified during AutoDev.

Stage 2: Accentuating your image with narrowband accents

The second stage is all about using the signal from the first stage in a manner you find aesthetically pleasing.

Straight up, there are two presets that are useful in two of the NBAccent's major use cases;

- 'Nebula'; to accentuate detail associated with Milkyway nebulosity
- 'Galaxy'; to accentuate smaller detail in other galaxies

These presets dial in the most useful settings for these two usecases.

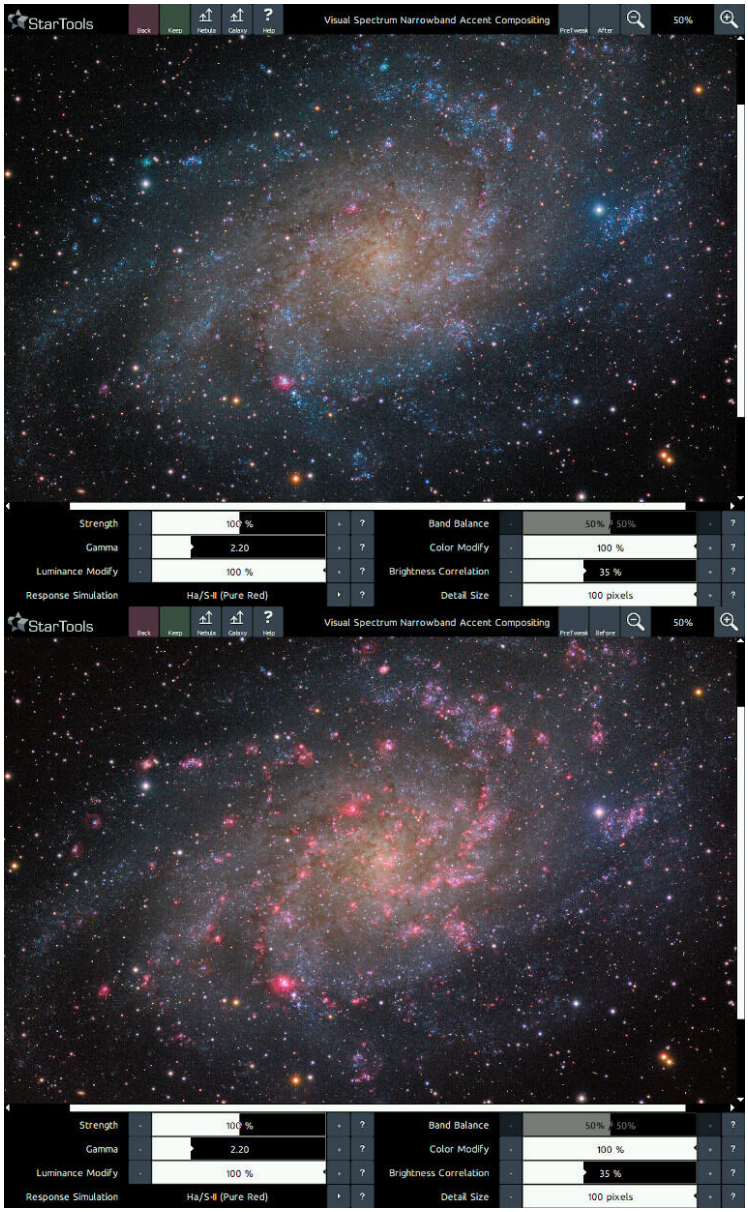

Isolated detail from the first stage is now modulated, color mapped and added to the visual spectrum image. Top: original image ("Before" switch toggled), bottom: image with added Hydrogen-alpha accents mapped to pure red.

The 'Response Simulation' parameter is responsible for the visual spectrum coloring equivalent that is synthesised from the narrowband data. The NBAccent module was designed to synthesise plausible visual spectrum coloring for a wide range of scenarios and filters:

- Ha/S-II (Pure Red): uses the narrowband data's red channel to add pure, deep red accents to the image. While pure red is rather rare in visual spectrum images (due to these emissions almost never existing by themselves and instead being accompanied by other emissions that are much bluer), it can nevertheless be useful to make these areas stand out very well.
- HII/Balmer Series (Red/Purple): uses the narrowband data's red channel to add the familiar red/purple colour of HII areas to the image. This mode makes the assumption that the other visual spectrum emissions from the Balmer series (almost all blue) are also present where the H-alpha line was detected. This mode tends to yield renditions that matches closely with the colouring of HII areas in actual visual spectrum data.

- Hb/O-III (Cyan): uses the narrowband data's green and blue channels to add pure cyan accents, corresponding to the colour of areas of strong Hb/O-III emissions as powered by nearby O or B-class blue giant stars.
- O-III (Teal); uses the narrowband data's green and blue channels to add teal green accents, corresponding to the colour of areas of strong O-III emissions
- Ha/S-II (Pure Red) + Hb/O-III (Cyan): uses pure deep red accents for data from the red channel, while using cyan accents for data from the blue and green channels. This mode is particularly useful for narrowband data acquired through the popular duo/tri/quadband filters.
- Ha/S-II (Pure Red) + O-III (Teal): uses pure deep red accents for data from the red channel, while using teal green accents for data from the blue and green channels. This mode is particularly useful for narrowband data acquired through the popular duo/tri/quadband filters.
- HII/Balmer Series (Red/Purple) + Hb/O-III (Cyan): synthesises the full Balmer series (red/purple) from the red channel, while using cyan accents for data from the blue and green channels. This mode is particularly useful for narrowband data acquired through the popular duo/tri/quadband filters.
- HII/Balmer Series (Red/Purple) + O-III (Teal): synthesises the full Balmer series (red/purple) from the red channel, while using green accents for data from the blue and green channels. This mode is particularly useful for narrowband data acquired through the popular duo/tri/quadband filters.

The 'Luminance Modify' and 'Color Modify' parameters, precisely control how much the module is allowed to modify of the visual spectrum image's luminance/detail and colour respectively. For example, by setting 'Luminance Modify' to 0%, and leaving 'Color Modify' at 100%, only the colouring will be modified, but the narrowband accent data will not (perceptually) influence the brightness of any pixels of the final image. Conversely, by setting 'Color Modify' to 0% and 'Luminance Modify' to 100%, the narrowband accent data will significantly brighten the image in areas of strong narrowband emissions, however the colouring will remain (perceptually) the same as the visual spectrum input image.

Wrap up

Description

- This module, and the changes to other modules, greatly simplifies the adding of narrowband accents to an image. After the initial loading of the data this work can now be done in one module.
- The narrowband data is processed in parallel to the other data 'behind the scenes' by the Bin, Crop and Wipe modules.
- Since the purpose of the narrowband data is in this case to accentuate the visual data the narrowband data is processed independently from, but parallel to the main dataset.

When to use

- When you want the narrowband data to accentuate certain elements - and not to contribute to the luminance or chrominance of the main data.
- Use the narrowband data either to supplement LRGB, or to accent it - don't use it for both.
- For example - don't include Ha as luminance if you want to use it for accents.
- If you want the narrowband data to contribute to the luminance or chrominance of the main data use the Compose module instead.
- Use it when you want to:
  - Accentuate areas of H-II in galaxies.

    - Accentuate large scale background nebulosity.
    - Accentuate features in planetary nebulae and other objects.
- Good candidate datasets for accentuating are:
    - Ha, O-III or Ha + O-III datasets
    - Datasets from OSC and DSLRs taken using duo/tri/quad band filters with narrow spectral responses.

What result to look for

- The accents should be clear and well balanced with the rest of the image.
- Coloring of non-narrowband regions should not change.
- Then move on to Wavelet 'Denoise'.

Description Of Controls

Screen 1

ROI sliders

Allow to set ROI.

- The RoI is initially inherited from the AutoDev settings

Ignore Fine Detail <

Allows 'NB Accent' to ignore small features such as noise when allocating dynamic range.

- Particularly important when the image has a low signal to noise ratio (SNR).
- Default is Off.
- Range is Off then 1.1 to 50.0 pixels.
- Increase to exclude noise - usually until the image doesn't darken any more.

Threshold

Introduces intentional background clipping in the narrowband dataset.

- Increasing the value will produce an even background.
- The effect will isolate the accent objects in the narrowband dataset.
- Use as last resort only.
- Default is 0%
- Range is 0% to 100%

Outside ROI Influence

Defines how much dynamic range to reserve for outside the Region of Interest (ROI).

- Defaults to 15%.
- Range is 0% to 100%.
- Reduce if you have a high dynamic range subject and you want to allocate more dynamic range to the subject.
- Avoid 'NB Accent' clipping the data by never reducing this value to 0%.
- Increase if there is some detail outside the ROI which you want to show up.

Shadow Linearity

Sets how much linearity is applied to the shadows in the narrowband dataset - before non-linear stretching takes over.

- Values above 50% allocate more dynamic range to the shadows and background.
- Values below 50% allocate more dynamic range to the highlights and foreground.
- Default is 50%.
- Range is from 0% to 100%.

Screen 2

Presets

Provide a baseline which fits your image subject.

- 'Nebula' - Uses HII Balmer series coloring (purple) for HII regions, no 'Detail Size' restriction and no 'Brightness Correlation'

- Galaxies - Uses Ha/SII (red) coloring for HII regions, 100 pixel maximum 'Detail Size' and 50% 'Brightness Correlation'

Strength

- Controls the overall strength of the enhancement effect.
- Increasing the Strength parameter increases the enhancement.
- Avoid to introduce artifacts under certain circumstances using higher values.
- Default is 100%.
- Range is 0% to 1,000%.

Gamma

Allows you to apply an additional Gamma correction to the stretched dataset

- Increasing the value above 1.00 will make bright details stand out more - so more dynamic range will be allocated to bright detail.
- Decreasing the value below 1.00 will make bright details stand out less.
- Keep at 1.00 for no change.
- Default is 1.00.
- Range is 0.00 to 10.00

Luminance Modify

- Sets how much luminance may change by adding NB accents.
- Increasing this parameter increases the brightness of the accents.
- Default is 100%.
- Range is 0% to 1,000%.

Response Simulation

Selects the visual spectrum coloring of the accent signal synthesized from relevant narrowband data channels.

- Options are:
  - Ha/S-II (Pure Red): creates deep red Ha accents out of narrowband data's red channel.
  - HII/Balmer Series (Red/Purple): creates red/purple H-II accents out of narrowband data's red channel.
  - Hb/O-III (Cyan): creates cyan Hb/O-III accents out of narrowband data's green and blue channels.
  - O-III (Teal): creates teal O-III accents out of narrowband data's green and blue channels.
  - Ha/S-II (Pure Red) + Hb/O-III (Cyan): creates deep red Ha accents out of narrowband data's red channel and cyan Hb/O-III accents out of green and blue channels. Useful for datasets acquired by duoband filters.
  - Ha/S-II (Pure Red) + O-III (Teal): creates deep red Ha accents out of narrowband data's red channel and teal O-III accents out of green and blue channels. Useful for datasets acquired by duoband filters.
  - HII/Balmer Series (Red/Purple) + Hb/O-III (Cyan): creates red/purple H-II accents out of narrowband data's red channel and cyan Hb/O-III accents out of green and blue channels. Useful for datasets acquired by duoband filters.
  - HII/Balmer Series (Red/Purple) + O-III (Teal): creates red/purple H-II accents out of narrowband data's red channel and teal O-III accents out of green and blue channels. Useful for datasets acquired by duoband filters.
- Default is Ha/S-II (Pure Red) for 'Galaxy' preset, and HII/Balmer Series (Red/Purple) for 'Nebula' preset

Band Balance

Sets the balance between red vs. green/blue channels. Only for duoband modes of response simulation parameter

- 100%/0% leaves red signal only, 0%/100% leaves green/blue signal only
- Default is 50%/50%.
- Range is 100%/0% to 0%/100%.

Color Modify

- Sets how much chrominance may change by adding NB accents.
- Increasing this parameter allows stronger color change in the accent's area.
- Default is 100%.
- Range is 0% to 1,000%.

Brightness Correlation

Defines how much detail must pre-exist before it is allowed to be enhanced by narrowband data.

    - Higher values allow less narrowband data to be added
- Default is Off ('Nebula' preset), 50% ('Galaxy' preset)
- Range is Off to 100%

Detail Size

Specifies the maximum size of detail used for narrowband accents.

- IValues above 0 are helping to reduce color or brightness shifts on large scales
- Default is 0 Pixel ('Nebula' preset) or 100Pixel ('Galaxy' preset).
- Range is 0 pixel to 100 pixels

Unified De-Noise:
Detail Aware Wavelet-based Noise Reduction and Noise Grain Shaper

200% zoom showing pin-point accurate, autonomous, fully configurable noise reduction based on data mined statistics. No masks, no other subjective crutches.

The Unified De-Noise module offers detail-aware, astro-specific noise reduction. Paired with StarTools' Tracking feature, it yields pin-point accurate results that have no equal.

Whereas generic noise reduction routines and plug-ins for terrestrial photography are often optimized to detect and enhance geometric patterns and structures in the face of random noise, the Unified De-Noise module is optimized to do the opposite. That is, it is careful not to enhance structures or patterns, and instead attenuates the noise and gives the user control over its appearance. Its unified noise reduction routines are specifically designed to be "permissible" even for scientific purposes - that is, it was designed to only carefully remove energy from the image and not add it; it strictly does not sharpen, edge-enhance or add new "detail" to the image.

The Unified De-Noise module is the ultimate application of the signal evolution Tracking feature (which data mines every decision and noise evolution per-pixel during the user's processing). The results that Unified De-Noise is able to deliver autonomously because of this deep knowledge of your signal and its evolution, are absolutely unparalleled.

In addition, StarTools is currently the only software that also targets walking noise (streaks) caused by not being able to dither during acquisition (for example when conducting Electronically-Assisted Astronomy).

Usage

Denoising starts when switching Tracking off. It is therefore generally the last step, and for good reason. Being the last step, Tracking has had the longest possible time to track and analyze noise propagation.

Setup

The first stage of noise reduction involves helping StarTools establish a baseline for visual noise grain and the presence (and direction) of walking noise. To establish this baseline, increase the 'Grain size' parameter until no noise grain of any size can be seen any longer. StarTools will use this baseline as a guide as to what range of details in your image is affected by visible noise.

If walking noise is present, then temporarily set 'Grain Size' parameter to 1.0. Next, use the 'Walking Noise Angle' level setter, or click & drag an imaginary line on the image in the direction of the walking noise to set the 'Walking Noise Angle' that way. Now increase the 'Walking Noise Size' parameter until individual streaks are no longer visible in the direction you detected them in (though other imperfections may still be visible). After that, increase the 'Grain Size' parameter until other noise grain can no longer be seen.

StarTools is the only software solution for astrophotography that can target walking noise.

After clicking 'Next', the wavelet scale extraction starts, upon which, after a short while, the second interactive noise reduction stage interface is presented.

Main operation

Noise reduction and grain shaping is performed in three stages.

Stage one

The first-pass algorithm is an enhanced wavelet denoiser, meaning that it is able to attenuate features based on their size. Noise grain caused by shot noise (aka Poisson noise) - the bulk of the noise astrophotographers deal with - exists on all size levels, becoming less noticeable as the size increases. Therefore, much like the 'Sharp' module, a number of scale sizes ('Scale n' parameters) are available to tweak, allowing the denoiser to be more or less aggressive when removing features deemed noise grain at different sizes. Tweaks to these scale parameters are generally not necessary, but may be desirable if - for whatever reason - noise is not uniform and is more prevalent in a particular scale.

'Noise grain caused by shot noise (aka Poisson noise) - the bulk of the noise astrophotographers deal with - exists on all size levels, becoming less noticeable as the size increases'

Different to basic wavelet denoising implementations, the algorithm is driven by the per-pixel signal (and its noise component) evolution statistics collected during the preceding image processing. E.g. rather than using a single global setting for all pixels in the image, StarTools' implementation uses a different setting (yet centered around a user-specified global setting) for every pixel in the image.

The wavelet denoising algorithm is further enhanced by a 'Scale Correlation' feature parameter, which exploits common psychovisual techniques, whereby noise grain is generally tolerated better in areas of increased (correlated) detail.

The general strength of the noise reduction by the wavelet denoiser, is governed by the 'Brightness Detail Loss' and 'Color Detail Loss' for luminance (detail) and chrominance (color) respectively.

The noise reduction solution in StarTools is based wholly around energy removal - that is attenuation of the signal and its noise components in different bands in the frequency domain - and avoids any operations that may add energy. It does not enhance edges, does not manipulate gradients, and does not attempt to reconstruct detail. These important attributes make its use generally permissible for academic and scientific purposes; it should never suggest details or features that were never recorded in the first place.

Stage two

Any removed energy, is collected per pixel and re-distributed across the image in a second pass, giving the user intuitive control, via the 'Grain Dispersion' parameter, over a hard upper size limit beyond which grain is no longer smoothed out.

Stage three

Left: input image. Middle: denoised image without Grain Equalization. Right: denoised image with Grain Equalization. Notice the vast differences in noise grain prevalence independent of brightness, and notice subtle pre-existing noise grain has been re-introduced in an unobtrusive and visually pleasing manner.

The 'Grain Equalization' parameter lets the user reintroduce removed noise grain in a modified, uniform way, that is; appearing of equal magnitude across the image (rather than being highly dependent per-pixel signal strength, stretches and local enhancements as seen in the input image).

The 'Grain Equalization' feature an acknowledgement of the "two schools" of noise reduction prevalent in astrophotography: there are those who like smooth images with little to no noise grain visible, and there are those who find a tightly controlled, uniform measure of noise grain desirable for the purpose of creating visual interest and general aesthetics (much like noise grain is added for a "filmic" look in CGI). The noise signature of the deliberately left-in noise, is precisely shaped to be aesthetically pleasing for precisely this purpose.

Lastly, it should be noted that the 'Grain Equalization' feature only shapes and re-introduces noise in the luminance portion of the signal, but not in the chrominance (color) portion of the signal.

Wrap Up

Description

To get rid of different types of noise while preserving detail.

- Separates Brightness and Color - this allows separate control of brightness and color noise.
- Uses information gained while Tracking was on to help target the noise.

The 'Denoise' module uses the following methods to separate detail from noise:

- Tracking identifies the areas of higher noise.
- Scale Correlation techniques are used to identify detail.
- This allows a high level of control over noise - with extra control over areas of detail.

When to use

- Final 'Denoise' is usually done after the 'Color', 'Entropy' or 'SuperStructure' module. Click 'Apply noise reduction' when switching Tracking off.

- Also, try using the 'Super Structure' modules' 'Isolate' preset with no mask set just before using the 'Color' module then final 'Denoise' module - this will help to push back the noise. Watch out for halos around the stars though.

What result to look for

- Background noise should be greatly reduced or eliminated without affecting detail significantly.
- Look out for any remaining noise blotches - if found go back and check the Grain Dispersion settings.
- Look out for any reductions in the detail - if found try reducing Brightness Detail Loss or 'Color' Detail Loss.
- Use the 'Before'/'After' button to see the effect of the module.

Ways of getting better results

- Improve the Signal-to-Noise ratio (SNR) of the original image - by taking more subs. Also, make sure the subs are long enough.
- With light polluted data you will need many more subs to get equivalent results.
- If you have pattern noise try Dithering if you don't already.
- Make sure you have used the 'Bin' module to reduce the resolution (and improve the SNR) if the image is oversampled.
- Try using the 'SuperStructure' module 'Isolate' preset with no mask set just before using 'Denoise' module - to help push back the noise.
- If there is background color noise this may be de-emphasized by using the Dark Saturation control in the 'Color' module.
- If the background noise cannot be controlled successfully in 'Denoise' - it may be necessary to go back and redo Develop/'AutoDev' to control the final stretch a little to limit the noise to a level 'Denoise' can handle. To do this use the 'Restore' - 'Linear, Wiped' button.

After Use

- Save the image and finish - or apply one of the modules not available when Tracking is on, i.e.: 'Heal' 'Repair' or 'Synth' as needed.

Description of Controls

Screen 1 - Select filter type and grain size

Walking Noise Angle

If you have walking noise in your image, for example if dithering has not been used, use this setting and Walking Noise Size to identify and remove it. Sets the angle of the walking noise to be removed. Set Walking Noise Size and Angle parameters as necessary so as to eliminate any Walking noise in the image. Then set the Grain Size

- Increase the Walking Noise Size above 1.0 to enable Walking noise removal
- Set the Walking Noise Angle before you do final adjustment of the Walking Noise Size. Set the Grain Size last.
- Adjust the value to reflect the angle of the walking noise. The angle is measured from the vertical.
- The angle can be set by clicking and dragging a line on the image at the angle of the walking noise.
- Default is 0 degrees.
- Range is 0 to 359 degrees.

Walking Noise Size

Sets the size of the walking noise to be removed.

- Increase the value above 1.0 to enable walking noise removal.
- Set the Walking Noise Angle before you do final adjustment of Walking Noise Size.
- Once the walking noise has been removed you can adjust the Grain Size to remove the other noise.
- Default is 1.0 pixels (Off).
- Range is 1.0 to 30.9 pixels.

Grain Size

Specifies the maximum size of the noise grain that is visible in the image.

- Once set tells module that anything larger in scale than this is not noise.
- Specifies over how large an area it spreads the energy that was contained in pixels that are smoothed.
- Default is 2.0 pixels. Range is 1.0 to 30.9 pixels.
- Experiment until you find a value which causes the noise grain and clumps to be dispersed so that can no longer be seen at any scale. Values up to 15-30 are fairly common with noisy data.
- Do not exceed what is needed so as to preserve large scale detail as much as possible.
- Concentrate on the noise and don't worry about the detail. This is a visual representation to help find the right setting and the signal is not being affected.

Screen 2 - Customize noise reduction

Identifying and protecting detail

- Structures larger than the Grain Dispersion are not considered to be noise.
- Scale Correlation identifies how much the smaller structures are analyzed when looking for detail.

Scale Settings

Defines how hard the noise reduction is done for different sizes of noise.

- Scales do not have absolute limits to the range - its is more like a particular scale brings detail of a certain size into focus - and that other size detail is out of focus to varying degrees depending on its size.
- The following are broad guidelines:
  - The largest scale (Scale 5) is approximately 100-120 pixels.
  - The smallest size (Scale 1) is around one pixel.
  - The intervening scale sizes increase exponentially.
- Increase the scale value if noise is noticeable at that scale. Decrease it if detail is being affected.
- For Scales 1-4 the default is 95%. For Scale 5, default is 50%. Range is 0% to 100%.
- Scale 5 may need to be increased if there is large scale noise. If there is noise at this scale it has often been introduced artificially during debayering or subsequent processing and is not from natural Poisson noise sources. This type of noise can show scale correlation too - so we need to reduce the Scale Correlation (from the default of 6 to 4 - 2) to avoid the algorithm mistakenly identifying noise for signal.
- Values up to 95% aren't unusual.

Scale Descriptions

- Scale 1 - Controls the extent of noise reduction in fine detail (e.g. single pixels)
- Scale 2 - Controls the extent of noise reduction in small to medium detail.
- Scale 3 - Controls the extent of noise reduction in medium detail.
- Scale 4 - Controls the extent of noise reduction in medium to large detail.
- Scale 5 - Controls the extent of noise reduction of large noise blotches/grain.

Grain Equalization

Reintroduces grain in a way that takes into account how noise grain was affected during processing. This helps in cases where the background looks unnaturally smooth. It creates a film-like look.

- If set to 0% then no grain is reintroduced and the detail identified using Scale Correlation is smoothed completely.
- If set to 100% then all grain is retained and there is no smoothing.
- Set to 0 initially and then, at the end, after the noise has been removed, increase until the correct amount of noise grain is found.
- Default is 0%.
- Range is from 0% to 100%.

Brightness Detail Loss

This balances noise reduction with detail loss in brightness. Larger values will do more aggressive noise reduction possibly causing some detail loss. Smaller values will reduce noise reduction. 0% is no noise reduction.

- Default is 50%.
- Range is 0% to 100%.
- Reducing to 0% turns off noise reduction completely.

'Color' Detail Loss

This balances noise reduction with detail loss in color. Larger values will do more aggressive noise reduction possibly causing some color loss.

- Default is 50%.
- Range is 0% to 100%.

Scale Correlation

Usually, when there is a correlation between image elements over multiple scales it indicates important detail in an image. This is how 'Denoise' identifies detail. It can then provide the control to protect this detail from the denoising algorithm.

For every scale, the scale correlation algorithm looks at the immediate neighboring scales to see if detail in that scale exists. If detail in the neighboring scale exists, this is taken into account when determining how much noise reduction is applied. The Scale Correlation parameter specifies how much neighboring scales are evaluated.

- Defines how much 'Denoise' identifies smaller scale features as being detail that correlates with large scale detail that contains it. The scale correlation value controls how far the correlation propagates to other scale levels.
- Default value is 50%.
- The range is 0% to 100%
- Certain types of noise can have scale correlation that makes them look like detail. To avoid this mis-identification in noisy images the Scale Correlation value can be reduced so it doesn't search for correlation in the smaller elements. This can be a problem when:
  - There are too few sub-frames taken when using an OSC or DSLR or when using drizzle, or
  - Using insufficient dithering.
  - The noise has been introduced artificially during debayering or subsequent processing and not from natural Poisson noise sources.
- Larger values mean smaller elements of a larger structure are searched for to identify detail.
- Smaller values means that more of the smaller elements of detail will not be identified and so will not have the additional control over the denoise process.

Grain Dispersion

See also the description for Grain Size above.

- The Grain Dispersion influences the noise reduction of all the other controls apart from Scale Correlation - which defines how parts of the image are protected from noise reduction.
- Defaults to Grain Size value set previously.
- Range is 1.0 to 30.9 pixels.
- Structures larger than the Grain Dispersion are considered detail, not noise.
- Defines a surface area over which it can safely redistribute energy that was taken away (denoised).
- Typical values <30 pixels - there will be a value beyond which there will be little effect - don't exceed the maximum size needed.

'Flux': Automated Astronomical Feature Recognition and Manipulation

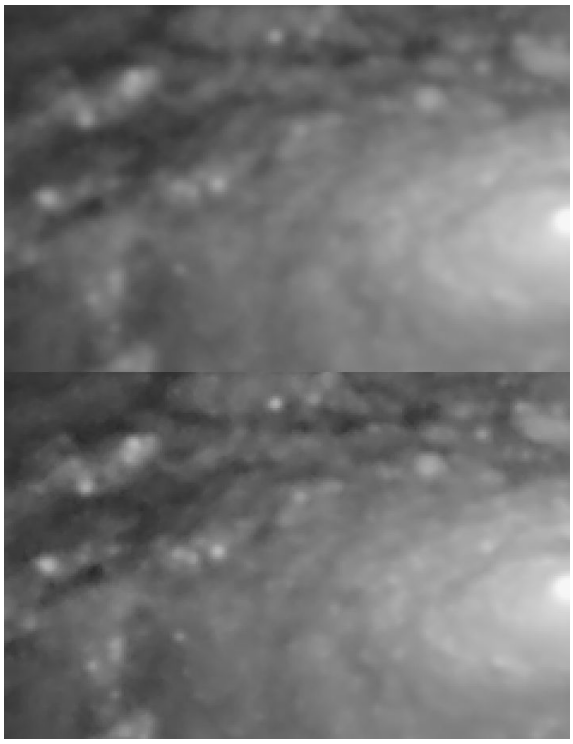

'Flux' sharpening by self similarity.

The Fractal 'Flux' module allows for fully automated analysis and subsequent processing of astronomical images of DSOs.

The one-of-a-kind algorithm pin-points features in the image by looking for natural recurring fractal patterns that make up a DSO, such as gas flows and filaments. Once the algorithm has determined where these features are, it then is able to modify or augment them.

Knowing which features probably represent real DSO detail, the Fractal 'Flux' is an effective denoiser, sharpener (even for noisy images) and detail augmenter.

Detail augmentation through flux prediction can plausibly predict missing detail in seeing-limited data, introducing detail into an image that was not actually recorded but whose presence in the DSO can be inferred from its surroundings and gas flow characteristics. The detail introduced can be regarded as an educated guess.

'Flux' Sharpening by self-similarity feature detection. Only areas that are deemed recurring detail are sharpened.

It doesn't stop there however – the Fractal 'Flux' module can use any output from any other module as input for the flux to modulate. You can use, for example, the Fractal 'Flux' module to automatically modulate between a non-deconvolved and deconvolved copy of your image – the Fractal 'Flux' module will know where to apply the deconvolved data and where to refrain from using it.

Wrap Up

Description

- To de-noise, sharpen or augment DSO detail.
- The Fractal 'Flux' module uses an algorithm which pin-points features in the image by looking for natural recurring fractal patterns that make up a DSO, such as gas flows and filaments.
- Once the algorithm has determined where these features are, it then is able to modify or augment them.

When to use

- Use after extracting the most detail (i.e. 'Contrast', 'HDR', 'Sharp', 'Decon') and before using the 'Color' module.
- The noise reduction capabilities of the 'Flux' module can be used in place of the 'Denoise' module where Tracking is not used (for example where an image is not linear to begin with and StarTools is being used just to touch up an image).
- It could be a decent alternative to the 'Denoise' module after working in the 'Layer' module with Tracking off. In particular, for fine noise due to the way the self-similarity works.
- Only use once.

What result to look for

- Any change may be subtle - Use Before/After button to see the change.
- If 'Sharpen' preset is chosen the image will look sharper in the areas where there is detail.
- If Noise preset is chosen the noise in the image will be reduced.

Description of Controls

Screen 1

Wavelet Library

This specifies the amount of 'knowledge' the fractal flux module starts off with. The smaller the amount of knowledge the more flux you might get.

- Small.
- Large.
- Default is Small.
- The smaller the amount of knowledge the more flux you might get.
- A larger wavelet library will make the module better at relating detail to other detail at larger scales (essentially making it more precise), but will also make the effect a lot less 'dramatic', especially when noise is involved.

Screen 2

'Sharpen' preset

This button optimizes the settings for sharpening of the image.

Detail preset

This button optimizes the settings for adding natural-looking detail to the image.

- This adds artificial detail but...
- It can distort stars so using a star mask to mask out stars is advisable.

Noise preset

This button optimizes the settings for removing noise from the image.

Mask

For general instructions on using masks see Mask.

- Recommended to use an inverted star mask with this module - Auto, Stars, Do, Invert, Keep.

Algorithm

This selects the filtering algorithm to use. Choose from:

- Add Detail - Used to add detail to the image.
- Modulate Unsharp Mask - the presence of flux turns on and off an Unsharp Mask sharpening algorithm. Effectively performs a localized sharpening of only ‘interesting’ features, while leaving others (including noise) alone.
- Modulate Gaussian
- Modulate Median
- Modulate MOMH - Mean of Median Half - 25% upper and lower outlier rejection, taking mean of remaining values - Used to reduce noise.
- Modulate Undo Buffer- Use when you want to combine an image with another image using the flux to modulate.
- Default is Modulate Unsharp Mask ('Sharpen' preset), Modulate MOMH ('Noise' preset), Add Detail ('Detail' Preset)
- Use the default setting unless you understand and need the effects the other settings give.

Positive 'Flux'

Controls the amount any positive flux is amplified

- Default is 500% ('Sharpen' preset), 25% ('Detail' preset), 600% ('Noise' preset)
- Range is 0% to 1000%

Negative 'Flux'

Controls the amount any negative flux is amplified:

- Default is 0% ('Sharpen', 'Detail' presets), 600% ('Noise' preset)
- Range is 0% to 1000%

'Detail Filter'

Sets the minimum allowable size of any flux concentrations.

- Increasing this value allows only the bigger clumps of flux through.
- Decreasing this value allows more and more smaller, detailed flux to pass through.
- Default is 1.0 Pixel ('Sharpen', 'Noise' presets), 2.5 pixels ('Detail' preset)
- Range is 1.0 to 40.0 pixels.

Brightness Mask Power

Sets the power that should be applied to a pixel in the luminance mask if one is set in Brightness Mask Mode.

- Default is 3.00 ('Noise' preset), N/A (for 'Sharpen', 'Detail' presets)
- Range is 0.0 to 5.0

Brightness Mask Mode

Sets the mode of the optional masking using the luminance information. Options are:

- Off
- Where fg is dark, use bg
- Where fg is light, use bg
- Where fg is light & dark, use bg
- Where fg is grey, use bg
- Default is 'Off', 'Where fg is light, use bg' ('Noise' preset)

'Filter' Amount

Sets the strength of the Unsharp Mask sharpening.

- Default is 200% ('Sharpen' preset), 100% ('Noise' preset), N/A (for 'Detail' preset)
- Range is 0% to 1000%

'Filter' Radius

Specifies the maximum size of any details to be sharpened by the Unsharp Mask algorithm (if used) or the filter width for the other algorithms.

- Default is 1.5 pixels ('Sharpen' preset), 12.5 pixels ('Noise' preset), N/A ('Detail' preset)
- Range is 1.0 to 40.0 pixels.
- The larger the radius the larger the details that are sharpened.

'Filter' Fuzz

Specifies the kernel radius of an optional Gaussian blur, to be applied to the flux before it is passed to the Unsharp Mask routine for modulation.

- Default is 1.0 pixels ('Sharpen', 'Noise' presets), N/A ('Detail' preset)
- Range is 1.0 to 40.0 pixels.
- Increasing this value effectively 'smears out' the flux over a larger area so the detail around the flux is also sharpened.
- As the flux is spread out its intensity becomes lower. You may compensate for this by increasing the Positive 'Flux' and Negative 'Flux' amplifiers, or by increasing the 'Filter' Amount (sharpening strength) parameter.

Mask Fuzz

If a mask is used, Mask Fuzz controls the blending of the transition between masked and non-masked parts of the image.

- Default is 1.0 pixels
- Range is 1.0 to 40.0 pixels

'Heal': Unwanted Feature Removal

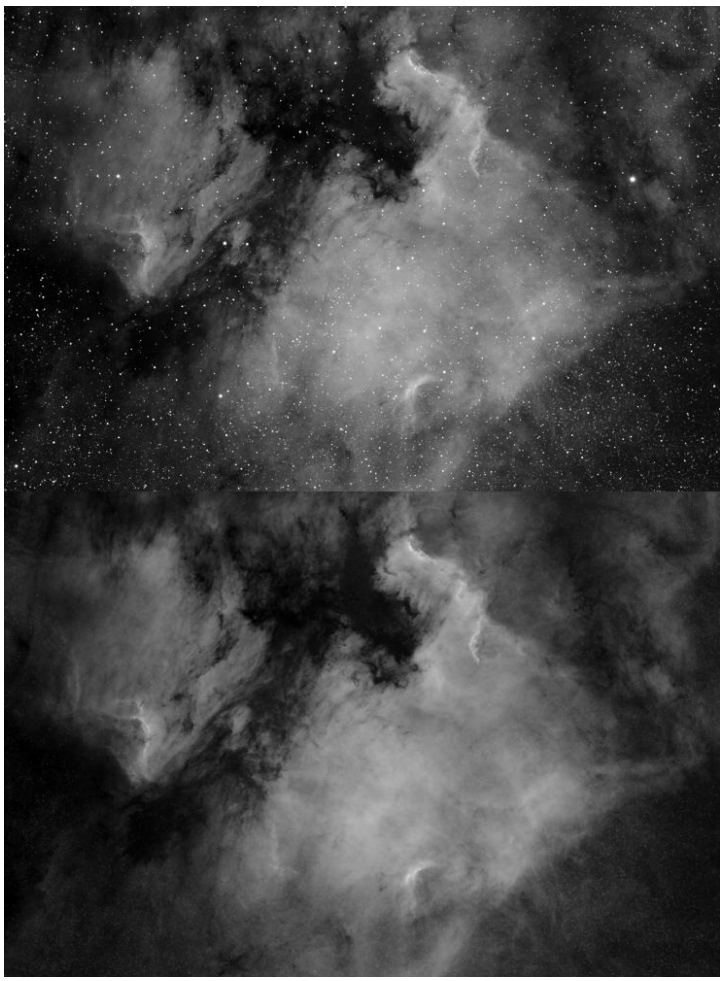

Removal of stars is an effective way to draw attention to the underlying nebulosity, or to process nebulosity separately from the stars.

The 'Heal' module was created to provide a means of substituting unwanted pixels in a neutral way.

Cases in which healing pixels may be desirable may include the removal of stars, hot pixels, dead pixels, satellite trails and even dust donuts.

The 'Heal' module incorporates an algorithm that is content aware and is able to synthesize extremely plausible substitution pixels for even the large areas. The algorithm is very similar to that found in the expensive photo editing packages, however it has been specifically optimized for astrophotography purposes.

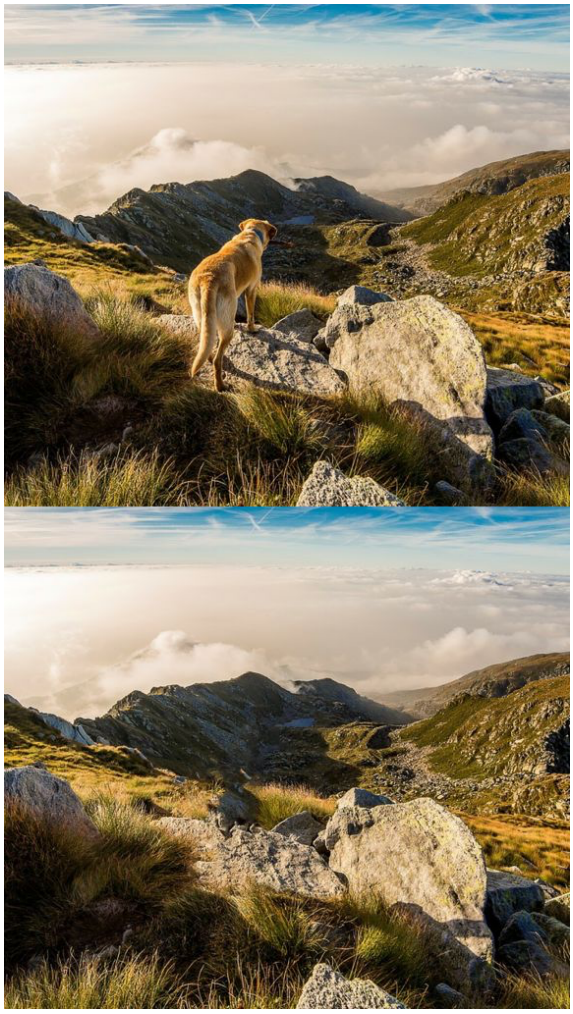

The 'Heal' module's algorithm is similar to that found in expensive photo editing packages

Usage

Getting started with the 'Heal' module in StarTools is a fairly straightforward affair; simply put any unwanted pixels in a mask and let the module do its thing. The more pixels are in the mask, the more the 'Heal' module will have to 'invent' and the longer the 'Heal' module will take to produce a result.

By using the advanced parameters, the 'Heal' module can be made useful in a number of advanced scenarios.

The 'New Must Be Darker Than' parameter lets you specify a brightness value that indicates the maximum brightness a 'new' (healed) pixel may have. This is useful if you are healing out areas that you later wish to replace with brighter objects, for example stars. By ensuring that the 'new' (healed) background is always darker than what you will be placing on top, you can simply use, for example, the Lighten mode in the 'Layer' module.

The 'Grow Mask' parameter is a quick way of temporarily growing the mask (see Grow button in the Mask editor). This is useful if your current mask did not quite get all pixels that needed removing.

The 'Quality' parameter influences how long the 'Heal' module may look for substitutes for each pixel. Higher quality settings give marginally better results but are slower.

The 'Neighborhood Area' parameter sets the size of the local area where the algorithm can look for good candidate seed pixels.

The 'Neighborhood Samples' parameter is useful if you are looking to generate more 'interesting' areas, based on other parts of the image. It can be useful for a large area being healed to avoid small repeating patterns. This feature is useful for terrestrial photography, however, this is often not needed or desirable for astrophotographical images. If you do not wish to use this feature, keep this value at 0.

The 'New Darker Than Old' parameter sets whether newly created pixels should always be darker than the old pixels. This may be useful for manipulation of the image in the 'Layer' module (for example subtracting the healed image from the original image).

Wrap Up

Description

- To remove unwanted elements in an image (e.g. satellite trails, stars, scratches, stuck pixels, etc) and replace them with pixels similar to the surrounding pixels.
- The 'Heal' module was created to provide a means of substituting unwanted pixels in a neutral way.
- The elements to remove are indicated by the mask.
- The 'Heal' algorithm is content aware and is able to synthesize extremely plausible substitution pixels for even the large areas.

Note: 'Heal' is not well suited for removing noise

When to use

- The 'Heal' Module can only be used when Tracking is Off.
- It is normally used after most of the processing has been done - after the Tracking has been turned off and 'Denoise' done.
- Can be used more than once.

What result to look for

- Ensure elements to be removed are completely gone.
- Ensure the replaced pixels look natural and are not too bright (or dim) - set 'New Must be Darker Than' value to adjust.
- Ensure the area around the removed elements don't have a halo.
- Make sure you don't have star 'doughnuts' - if so, grow the mask a bit to select more pixels and then retry.

Ways of getting better results

- If using the 'Heal' module to remove elements such as satellite trails try stacking with a different stacking algorithm - such as median or sigma stacking, or some other outlier rejection stacking algorithm.
- If possible, avoid using the 'Heal' module by removing the problem earlier on in processing - for example - by discarding bad sub-frames.

After Use

Consider using other modules that require Tracking to be off - like 'Repair' and 'Synth'.

Description of Controls

Mask

For general instructions on using mask see Mask.

- Mask areas which need to be healed.
- Set 'Grow Mask' as needed to ensure all pixels are masked.

New Must Be Darker Than

Sets the maximum value of new pixels (as a percentage of full white).

- This is useful when the 'Heal' algorithm generates elements that look like stars. By capping the maximum brightness these elements will not be mistaken for stars.
- Adjust this if the replacement pixels are too bright.
- Default value is 0% (Pure Black).
- Range is 0% to 100% (Off).

Grow Mask

Temporarily grows the mask - like pressing the 'Grow' button the set number of times.

- Default value is 0 pixels. Range is 0 to 10 pixels.
- Set Grow Mask as needed to ensure all pixels are masked.
- A value of 2-3 is normally all that is needed.

Quality

Sets the quality with which the new pixels are rendered. Higher quality settings give marginally better results but are slower.

- Low
- Medium
- High
- Ultra
- Default value is Medium.

Neighborhood Area

Sets the size of the local area where the algorithm can look for good candidate seed pixels.

- Default value is 200 pixels.
- Range 0 pixels to the image size.

Neighborhood Samples

For a large area being healed using seed samples avoid repeating patterns. Leaving it to 0 means the heal algorithm will not use any seed samples.

- Default value is 0.
- Range 0 to 500.

New Darker Than Old

Sets whether any new pixels created must be darker than the ones they replace.
This setting makes it possible to subtract the healed image from the original image.

- Yes
- No
- Default value is No.
- Set to Yes if creating a starless image or processing stars separately - as described in the Special Techniques section.

'Repair': Star Rounding and 'Repair'

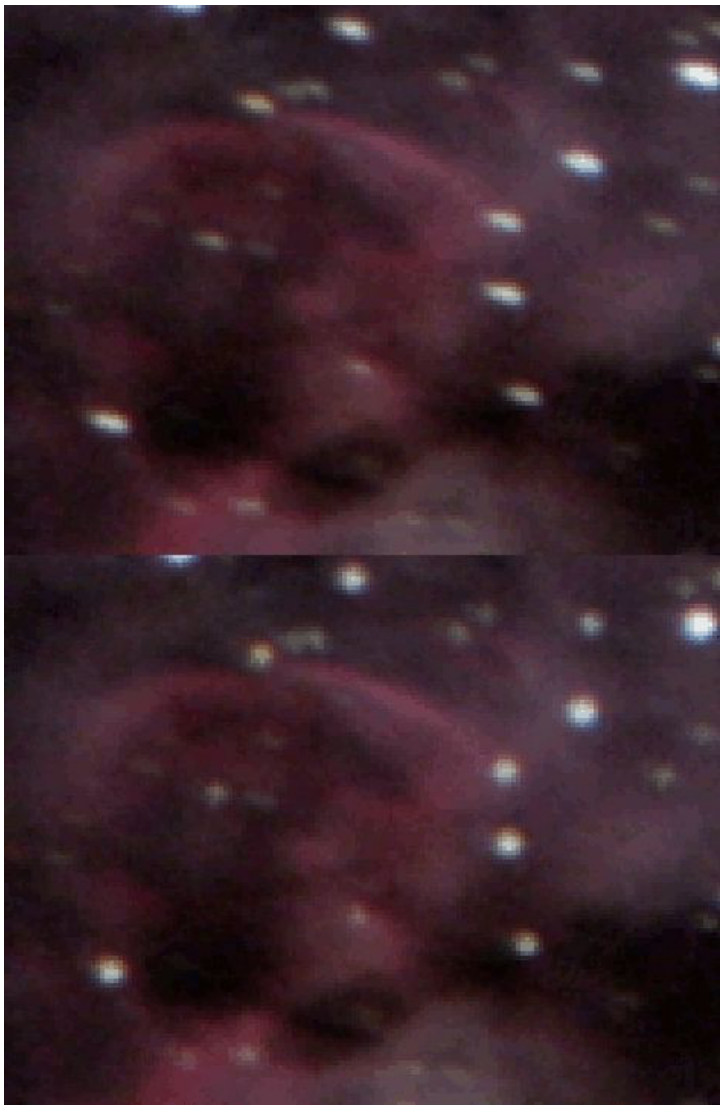

The 'Repair' module's 'Warp' algorithm uses the original pixels from the image to reverse-warp stars back into shape.

The 'Repair' module attempts to detect and automatically repair stars that have been affected by optical or guiding aberrations.

'Repair' is useful to correct the appearance of stars which have been adversely affected by guiding errors, incorrect polar alignment, coma, collimation issues or mirror defects such as astigmatism.

The 'Repair' module allows for the correction of more complex aberrations than the much less sophisticated 'offset filter & darken layer' method, whilst retaining the star's exact appearance and color.

The repair module comes with two different algorithms. The 'Warp' algorithm uses all pixels that make up a star and warps them into a circular shape. This algorithm is very effective on stars that are oval or otherwise have a convex shape. The 'Redistribution' algorithm uses all pixels that make up a star and redistributes them in such a way that the original star is reconstructed. This algorithm is very effective on stars that are concave and cannot be repaired using the 'Warp' algorithm.

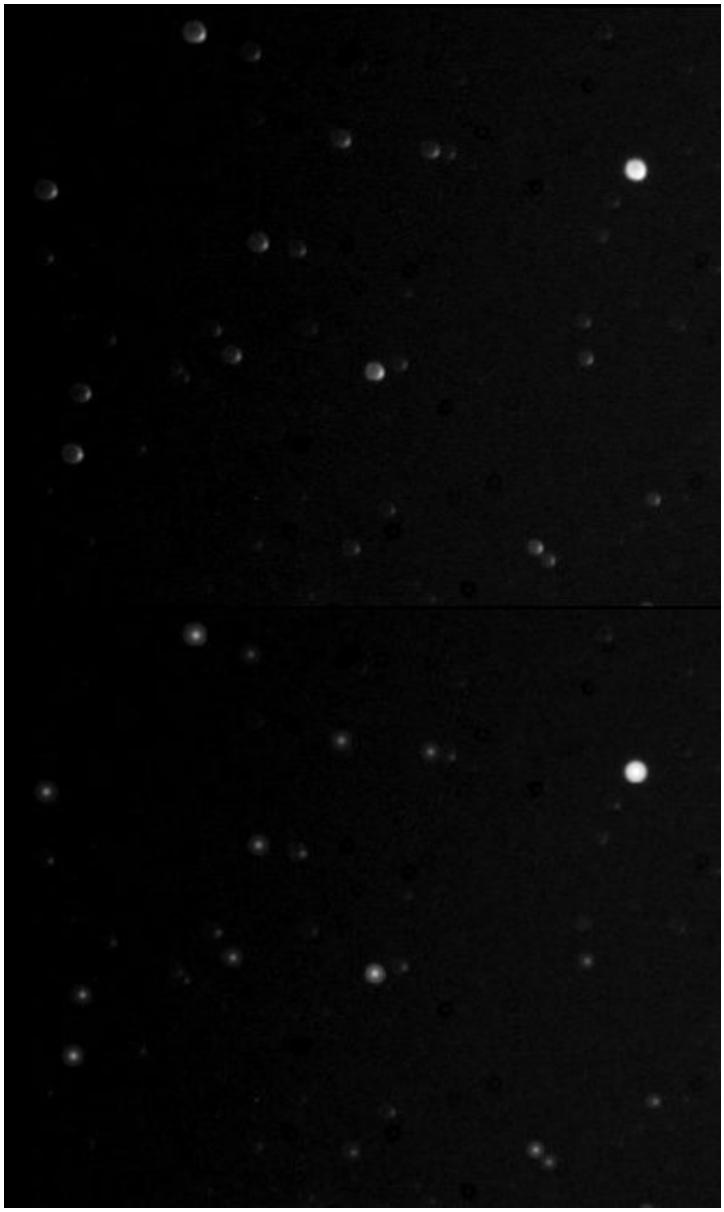

The 'Repair' module's 'Redistribute' algorithm uses the original pixels from the image and recalculates their appearance and position as if they originated from a point light source.

Wrap Up

Description

- To detect and automatically repair star defects such as those caused by blooming, guiding errors, bad polar alignment, collimation problems, or mirror defects such as astigmatism.
- The 'Repair' module is used to repair severely deformed stars.
- It complements the 'Shrink' module which is used to reduce the stars size only.

When to use

- Can only be used when Tracking is off. Towards the end of the processing workflow.
- Can be used more than once. Using different algorithms perhaps.

What result to look for

- Stars should be round and the color should be well distributed. If not, see the method for 'Improving the Stellar Profile' in the Special Techniques section.
- Check for artifacts around the corrected stars - increase Grow Mask if they exist.

- Check for places where two stars have merged into one - if found make sure that all stars in the mask are separated by at least one 'off' pixel.

After Use

Consider using other modules that require Tracking to be off - like 'Heal', 'Shrink' and 'Synth'.

Description of Controls

Mask

For general instructions on using masks see Mask

- For best results make sure that in the mask all stars are separated by one 'off' pixel - otherwise the stars may be merged.

Algorithm

Sets the algorithm the 'Repair' module uses when repairing stars:

- Warp - rounds stars by warping them into shape.
- Redistribute, Core is Avg Location - Tries to redistribute all starlight pixels. Use where stars are severely distorted by astigmatism or heavy coma where starlight is scattered, out of focus (doughnuts) or otherwise badly distorted. Tries to find the original centre of the star by averaging all the sample locations.
- Redistribute, Core is Brightest Pixel - As above but sets the centre of the star based on the brightest pixel.
- Debloom (Vertical Streaks) - Tries to recover stars that were affected by CCD blooming. The algorithm assumes that bloom is in a straight vertical direction and there is negligible rotation.

Grow Mask

Sets the amount to temporarily grow the mask by.

- Use a higher value if you can see artifacts around corrected stars.
- Default is 1 pixel.
- Range is 0 to 10 pixels.

Radial Samples

Applies to the Warp algorithm only - Sets the number of samples taken to determine the roundness of a star.

- A higher value may be needed to correct severely distorted stars but will be slower.
- Default is 32.
- Range is 16, 32 or 64.

Sub Sampling

Applies to the Warp algorithm only - Sets the number of subsamples per pixel to use while creating the Radial Samples.

- A higher value may make the stars more precisely round but will be slower.
- Default is 4x.
- Range is 2x, 4x or 8x.

'Synth': Star Resynthesis and Augmentation

Diffraction patterns are not painted on; they can be quite subtle.

The 'Synth' module generates physically correct diffraction and diffusion of point lights (such as stars) in your image, based on a virtual telescope model.

Besides correcting and enhancing the appearance of point lights (such as stars), the 'Synth' module may even be 'abused' for aesthetic purposes to endow stars with diffraction spikes where they originally had none.

It is worth noting that any other tools on the market today simply approximate the visual likeness of such star spikes and 'paint' them on. However the 'Synth' module can physically model and emulate most real optical systems and configurations to obtain a desired result.

Please note that the use of this module on your images falls outside of the realm of documentary photography.

Wrap Up

Description

To allow the user to adjust, augment or even replace stars in an image by the modeling of a telescope's diffraction of starlight.

This module can be used to:

- Make stars brighter and tighter. However the 'Shrink' module may be more appropriate if this is all you need to achieve.
- Re-model stars to show physically correct starlight energy distribution.
- Accurately model diffraction spikes.

When to use

- Can only be used when Tracking is off. Towards the end of the workflow.
- Only use once.

What result to look for

- Ensure the stars have the same sharpness as the rest of the image. If not, adjust the Blur parameter.

Ways of getting better results

- The 'Synth' module works best on images which are sharp and have low noise.

After Use

- Use other modules that require Tracking to be off - e.g. 'Heal', 'Shrink', 'Repair'.

Description of Controls:

Screen 1 - Specify the parameters of the virtual telescope

Newtonian preset

Sets default values relevant to a Newtonian reflector telescope.

- With this preset it may only be necessary to set Aperture and Focal Length. The default values of other settings may be Ok in most cases.

Refractor preset

Sets default values relevant to a refractor telescope.

- With this preset it may only be necessary to set Aperture and Focal Length. The default values of other settings may be Ok in most cases.

Aperture

Sets the scope's aperture. Apart from being a factor in the concentration of starlight and the diffraction limit, this setting also influences the relative size of all other components.

- Default is 200mm.
- Range is 50mm to 1050mm

Focal Length

Sets the scope's focal length.

- From this and the Aperture setting the FoV of the scope is derived.
- From the FoV and the angular size of the image object, the virtual stars' light concentration is estimated.
- Default is 1200mm.
- Range is 50mm to 10050mm

Sample Size

Sets the accuracy of the diffraction pattern.

- Larger sizes yield more accurate simulations, at the expense of increased computing time.
- Default is 2048 x 2048 pixels.
- Range is 512x512, 1024x1024, 2048x2048, 4096x4096 pixels.

Vanes

Sets the number of spider vanes present.

- This setting is the main influence on the number of distinctly visible diffraction spikes.
- Default is 4 [Newtonian], 0 [Refractor].
- Range is 0 to 4.

Vane Width

Sets the width of the virtual spider vanes.

- This setting influences the frequency of the occurrence of the 'rainbow patterns' in the spikes.
- Default is 1.6mm, 0.00mm [Refractor].
- Range is 0.0 to 100.0 mm.

Vane Support Width

Sets the width of the virtual spider vanes.

- This setting has a subtle influence on the central flare's diffraction pattern.
- Default is 1.6mm, 0.00mm [Refractor].
- Range is 0.0 to 100.0 mm.

Backscatter

Sets how much light is reflected back by the primary mirror and mirror Clips onto the Central Obstruction, the Vanes and the Vane Supports and subsequently emitted as diffused light.

- The setting dulls down the brightness of the diffraction spikes and parts of the central flares.
- Default is 2%.
- Range is 0 to 100%.

Central Obstruction

Sets the size of the Central Obstruction as a percentage of the full aperture.

- This setting greatly influences the diffraction pattern.
- A setting of 0% effectively turns the virtual scope into a refractor.
- Default is 25% [Newtonian], 0% [Refractor].
- Range is 0 to 100%.

Dispersion

Sets the strength of the dispersion - the subtle way the different wavelengths (colors) are diffracted differently.

- This setting effectively controls the strength of the 'rainbow effect' in the diffraction spikes.
- Note that stars that do not emit equal amounts of light in all wavelengths will generate 'rainbows' with different color balances.
- Default is 100%.
- Range is 0 to 100%.

Screws

Sets the number of evenly distributed screws around the edge.

- These screws block light in the optical tube assembly.
- This setting subtly affects the central flare's diffraction pattern.
- Default is 12 [Newtonian], 0 [Refractor].
- Range is 0 to 50.

Clips

Sets the number of clips present around the edge to hold the primary mirror or lens in place.

- This setting subtly affects the central flare's diffraction pattern.
- Default is 3 [Newtonian], 0 [Refractor].
- Range is 0 to 20.

Clip Size

Sets the size of the clips that hold the primary mirror or lens.

- This setting subtly affects the central flare's diffraction pattern.
- Default is 10mm [Newtonian], 0mm [Refractor].
- Range is 0 to 100mm.

Focuser

Sets the size of the optional focuser.

- This setting subtly affects the central flare's diffraction pattern - adding a faint but distinct spike.
- Default is None.
- Range is None, 1.25, 2 and 3 inch.

Focuser Travel

Sets the amount of inward travel of the focuser.

- This setting subtly affects the central flare's diffraction pattern - adding a faint but distinct spike.

- Default is 20mm [Newtonian], 0mm [Refractor].
- Range is 0 to 100mm.

Focuser Angle

Sets the angle at which the focuser is mounted on the optical tube assembly.

- This affects the angle of the faint but distinct spike caused by the inward travel of the focuser.
- Default is 120 [Newtonian], 0 [Refractor].
- Range is 0 to 360.

Screen 2 - Set the parameters of the image.

Mask

Set the star mask if not already set - Mask-Auto-Stars-Do-Keep

- Remove any noise that is also selected - by using the 'Shrink'-Grow buttons.
- All stars should be selected - not just a few.

Gamma Adjust

Allows adjustment of the gamma of the synthesized star layer.

- This is a basic but effective means of reducing the virtual stars' diffraction spike and central flare.
- Default is 1.00.
- Range is 0.00 to 5.00.

Core Whiteness

Sets the amount of color the 'Synth' module should pick up from the star.

- Default is 67%.
- Range is 0 to 100%.
- Note that this parameter is not real-time.

Blur

Sets the kernel size of the Gaussian blur that is applied to the virtual layer, prior to overlaying it.

- This should be close to the blur in the image caused by the atmosphere.
- Default is 1.0 pixels.
- Range is 1.0 to 5.0 pixels.

OTA Rotation

Sets how many degrees the point spread function (and thus the virtual stars and any diffraction artifacts) is rotated.

- Default is 0.
- Range is 0 to 359.
- Note that this parameter is not real-time.

Grow Mask

This is identical to the 'Grow' operation as found in the mask editor.

- This parameter is best increased if the healing stage of the 'Synth' algorithm is not getting rid of all traces of the original star.
- Default is 0 pixels.
- Range is 0 to 10 pixels.
- Note that this parameter is not real-time.

Image Diameter

Sets the angular size of the image in arc-minutes.

- This information is used to scale the point spread function appropriately.
- Default is 70 arc-minutes.
- Range is 1 to 301 arc-minutes.

- Note that this parameter is not real-time.

Overlay Mode

Sets one of seven algorithms to be used to overlay the virtual stars.
The range is:

- Hybrid, Remove Stars - Performs a mixture of addition and lighten-only, while removing the original stars.
- Lighten, Remove Stars - Performs a lighten-only operation, while removing the original stars.
- Addition, Remove Stars - Performs an addition operation, while removing the original stars.
- Hybrid, Keep Stars - Performs a mixture of addition and lighten-only, while keeping the original stars.
- Lighten, Keep Stars - Performs a lighten-only operation, while keeping the original stars.
- Addition, Keep Stars - Performs an addition operation, while keeping the original stars.
- Subtract, Keep Stars - Performs a subtraction operation, while keeping the original stars. This mode may assist in masking out stars completely, revealing detail otherwise drowned out by its glare.
- Default is 'Hybrid, Keep Stars'.

Blend Original

Sets how much of the original image should be blended in with the synthesized version.

- Default is 0%.
- Range is 0% to 100%.

Brightness Adjust

- Adjusts the brightness of the virtual stars.
- Default is 100%.
- Range is 0% to 400%.

Core Noisiness

This parameter attempts to compensate for noise in the stars' cores.

- This setting is best used if the core of the stars are not perfectly white and contain noise. If the latter is the case, such stars may be (incorrectly) split into multiple smaller stars.
- Default is 0%.
- Range is 0% to 25%.
- Note that this parameter is not real-time.

'Lens': Distortion Correction and Field Flattening

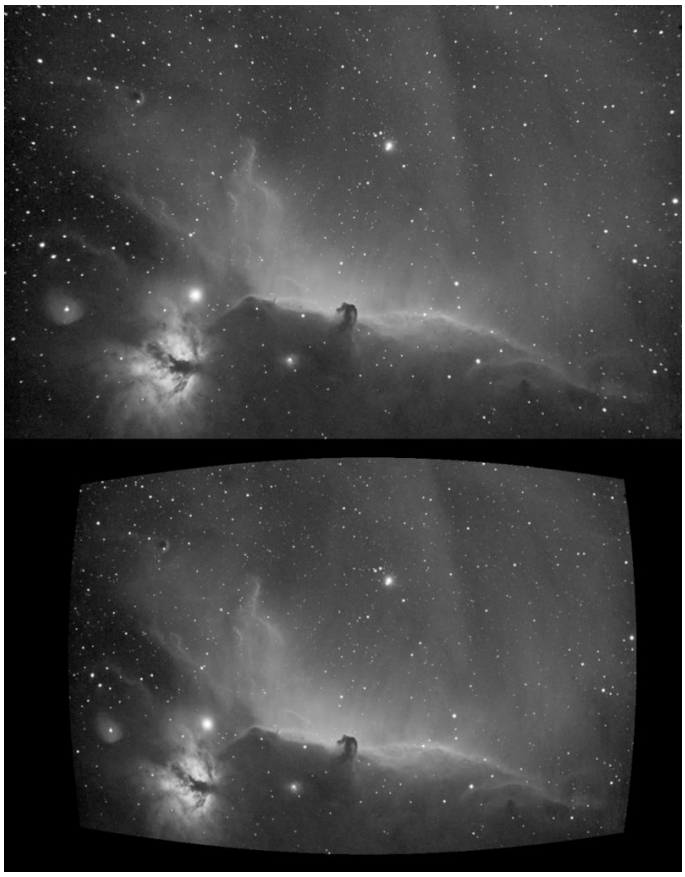

Top: source image (courtesy of Marc Aragnou), notice star elongation towards corners. Bottom: 'Lens' corrected image (without auto crop to show curvature).

The 'Lens' module was created to digitally correct for lens distortions and some types of chromatic aberration in the more affordable lens systems, mirror systems and eyepieces.

One of the many uses of this module is to digitally emulate some aspects of a field flattener for those who are imaging without a physical field flattener.

While imaging with a hardware solution to this type of aberration is always preferable, the 'Lens' module can achieve some very good results in cases where the distortion can be well modeled.

Wrap Up

Description

To correct common lens distortions such as chromatic aberration and coma. The 'Lens' module in StarTools corrects for 'simple' radial distortion, which can alleviate some of the worst effects of coma. It also allows for re-alignment of the Red, Green and Blue color channels to remove some of the worst effects of transverse (lateral) chromatic aberration. Axial (longitudinal) chromatic aberration can be controlled using the 'Filter' or 'Decon' modules.

When to use

- After the 'Bin' module (if used).
- Before the 'Crop' module - the 'Lens' module requires the centre of the image to be the same as when the image was taken.
- Stacking artifacts can be asymmetrical and require some re-centering of the image. This is easier to do in the 'Lens' module than by using the 'Crop' module.

What result to look for

When correcting coma:

- Stars at the edges and at the centre are round - any elongation has been compensated for.
- Diffraction spikes will, unfortunately, now be curved.

When correcting chromatic aberration:

- Star red and blue fringes are reduced or eliminated.

Ways of getting better results

To remove coma:

- Use a coma corrector/field flattener.
- Do this sort of adjustment to the subs before stacking if you can - as each sub will be slightly rotated.

To remove chromatic aberration:

- Use an Apochromatic refractor.
- Use a reflector.
- For any axial chromatic aberration - use the 'Decon' module or the 'Filter' module - which should remove any remaining blue fringes.

After Use

- Do 'Crop' if required - and continue with workflow.

Description of Controls

Auto 'Crop'

- Set if the 'Lens' module should automatically crop the image after correction.

Curvature Linked

This parameter defines if the field curvature of the red and the blue channel should be adjusted the same.

- For coma correction/field flattening both red and blue curvature should be the same.
- Changing this setting changes the Curvature Red and Curvature Blue to the same values.
- Default is 100.00%.
- Range is -100.00% to +300.00%.
- Change this value until the coma in the corners has been compensated for. Values of 150% are common.

Curvature Red

Sets the field curvature of the red channel independently of the blue channel.

- Default is 100.00%.
- Range is -100.00% to +300.00%

Curvature Blue

Sets the field curvature of the blue channel independently of the red channel.

- Default is 100.00%.
- Range is -100.00% to +300.00%

Center X

Sets the X-axis position of the centre of the distortion.

- Default is 0.0 pixels.
- Change this value if the image has been cropped unevenly or if the optical axis is offset.

Center Y

Sets the Y-axis position of the centre of the distortion

- Default is 0.0 pixels
- Change this value if the image has been cropped unevenly or if the optical axis is offset.

Red Shift X

Shifts the red channel along the X axis.

- Default is 0.0 pixels.
- Range -50.0 pixels to +50.0 pixels.
- Adjust this parameter in conjunction with Red Shift Y until the red fringe around the stars has been removed by moving it towards the centre of the star.

Red Shift Y

Shifts the red channel along the Y axis.

- Default is 0.0 pixels.
- Range -50.0 pixels to +50.0 pixels.
- Adjust this parameter in conjunction with Red Shift X until the red fringe around the stars has been removed by moving it towards the centre of the star.

Blue Shift X

Shifts the blue channel along the X axis.

- Default is 0.0 pixels.
- Range -50.0 pixels to +50.0 pixels.
- Adjust this parameter in conjunction with Blue Shift Y until the blue fringe around the stars has been removed by moving it towards the centre of the star.

Blue Shift Y

Shifts the blue channel along the Y axis.

- Default is 0.0 pixels.
- Range -50.0 pixels to +50.0 pixels.
- Adjust this parameter in conjunction with Blue Shift X until the blue fringe around the stars has been removed by moving it towards the centre of the star.

'Layer': Versatile Pixel Workbench

The 'Layer' module allows you to chain, mask and layer and apply countless of operations and filters.

The 'Layer' module is an extremely flexible pixel workbench for advanced image manipulation and pixel math, complementing StarTools' other modules. It was created to provide you with a nearly unlimited arsenal of implicit functionality by combining, chaining and modulating different versions of the same image in new ways.

Features like selective layering, automated luminance masking, a vast array of filters (including Gaussian, Median, Mean of Median, Offset, Fractional Differentiation and many, many more) allow you to emulate complex algorithms such as SMI (Screen Mask Invert), PIP (Power of Inverse Pixels), star rounding, halo reduction, chromatic aberration removal, 'HDR' integration, local histogram optimization or equalization, many types of noise reduction algorithms and much, much more.

Wrap Up

- To allow pixel manipulation in a very versatile but straightforward manner.
- The 'Layer' module performs Pixel math by means of graphical user interface. A huge range of operations is selectable, while multiple operations can be chained by means of a buffer.
- The effects of using this module can range from very simple to highly complex - and the effective use of the full range of this module requires a great deal of knowledge about the effects of these transformations and their combination.
- It is possible to cause clipping and artifacts with this module - which can negatively impact the image and subsequent processing. However, it is possible to make effective use of some of its power with a little basic knowledge.

When to use

- Usually it is best to use the 'Layer' module towards the end of the workflow after Tracking has been turned off and denoise has been done.
- Anything that causes artifacts (whether visible or not) like ringing, sharp edges, clipping, etc. will impact Tracking's ability to keep track of noise - so it is safest to use it after you have switched Tracking off.

- If you know that what you are doing with the 'Layer' module will not cause artifacts as described above then you can use it with Tracking on - and also the SMI and PIP techniques described below can be used safely with Tracking on.

What result to look for

Ensure that image layers are properly aligned.

Description of Controls

Mask

For general instructions on using masks see Mask

- The mask is applied to the foreground.
- Non-green parts appear in the result on the right as they are in the background image, unchanged.
- The effect of the 'Filter', 'Layer' Mode and the Brightness Mask Mode are all masked out.

Buffer Actions

The following Buffer Action buttons are at the top of the screen:

- Note: Fg=Foreground, Bg=background
- Open - Opens an image and loads it into the Foreground (centre image)
- Paste>Bg - Pastes Copy Buffer (i.e. the Image last Copied to the Copy Buffer) to the Background - greyed out if nothing in the Copy Buffer.
- Paste>Fg - Pastes Copy Buffer (i.e. the Image last Copied to the Copy Buffer) to the Foreground - greyed out if nothing in the Copy Buffer.
- Undo>Fg - Copies Undo Buffer (i.e. the old Image - as it was at the start of the previous module) to the Foreground.
- Undo>Bg - Copies Undo Buffer (i.e. the old Image - as it was at the start of the previous module) to the Background.
- Copy - Copies the 'Layer'ed composite result (right hand image) to the Copy buffer. If there is not an image in the right hand side StarTools may crash in some versions.
- Swap - Swaps the Foreground and Background images

'Layer' Mode

Defines how the foreground image should be layered on top of the background image.

This layering is applied after the 'Filter' has been applied to the Foreground.

- Blend - Copies the foreground over the background. Change the Blend Amount to change the Fg vs Bg.
- Lighten - Copies the pixels that are the lighter of the foreground and the background.
- Darken - Copies the pixels that are the darker of the foreground and the background.
- Multiply - Multiplies the background image by the foreground image.
- Add - Adds the foreground image to the background image.
- Subtract - Subtracts the foreground image from the background image.
- Difference - Calculates the difference (positive or negative) between the foreground image and the background image.
- HardLight - Dark pixels are darker, light pixels are lighter.
- Divide - Divides the background image by the foreground image.
- Brightness of fg - Applies the luminance ('brightness') information of the foreground to the background.
- 'Color' of fg - Applies the chroma ('color') information of the foreground to the background channel.

- Screen - Works like having two slide projectors projecting foreground and background images on top of each other.
- Power of Inverse - The blending algorithm used in the PIP algorithm.
- Desaturate fg (Average) - Projects a desaturated (i.e. black and white) version of the foreground over the background. The foreground is desaturated by using the average value of the combined red, green and blue channels.
- Desaturate fg (Luminance) - Projects a desaturated (i.e. black and white) version of the foreground over the background. The foreground is desaturated by using the weighted values for the red, green and blue channels corresponding to human-perceived luminance values (0.299, 0.587 and 0.114 respectively).
- Invert fg - Projects a negative version of the foreground on top of the background.
- 'Color' Extract fg - Generates an image devoid of luminance information (luminance information is set to unity), leaving only the (normalized) color information.
- Multiply Luminance - Multiplies the luminance info of the foreground image by the background image. Using this mode in conjunction with the 'Color' Extract fg mode allows us to recombine luminance and color information into one image again.
- Multiply Compensate Gamma - Multiplies the foreground and the background and then takes the square root.
- Distance - Calculates the magnitude as if the foreground and background were two vectors.
- Distance, MCG Hybrid - A hybrid of the Distance and Multiply Compensate Gamma modes.
- Overlay - Overlays the foreground image on top of the background image. It darkens the image but not as much as the Multiply mode.
- SoftLight - Dark pixels are darker, light pixels are lighter, edges softened and colors desaturated a bit.
- Multiply Foreground Only - takes the (filtered) foreground image and uses the Blend Amount as a multiplier - allows under- or overexposing an image - Over-exposing may clip.
- Default is 'Blend'

Blend Amount

Controls the relative intensity of the foreground image to the background image.

- Default is 100%. Range is 0% to 500%.
- For synthetic luminance - normally set so that the foreground is the same relative proportion as the components of the background.
- Examples of how to calculate the correct blend amount are shown in the Special Techniques section above.

Cap Mode

Defines how negative and over-unity values should be treated:

- Clip - Truncates the negative values to 0 and over-unity values to unity.
- Normalize - stretches levels to span the whole available dynamic range.
- Soft Clip
- Add 1/2 Unity + Clip - adds 50% of pure white to the signal, then truncates negative values to 0 and over-unity values to unity. This mode is particularly useful when using a layer mode or filter than may generate values that are centred around 0, such as the 'Subtract' layer mode or the Gaussian High-Pass filter.
- Default is 'Clip'

Offset X

Specifies the horizontal offset (in pixels) of the foreground relative to the background.

- Default is 0.0 pixels.
- Range is -50.0 pixels to 950.0 pixels.

Offset Y

Specifies the vertical offset (in pixels) of the foreground relative to the background.

- Default is 0.0 pixels.
- Range is -50.0 pixels to 950.0 pixels.

Brightness Mask Mode

Specifies an optional masking using the brightness information contained in both the foreground and background mode. This feature allows you to blend an image based on brightness. The available modes are:

- Off
- Where fg is dark, use bg
- Where fg is light, use bg
- Where fg is light & dark, use bg
- Where fg is grey, use bg
- Where composite is dark, use bg
- Where composite is light, use bg
- Where composite is light & dark, use bg
- Where composite is grey, use bg
- Default is Off

Brightness Mask Power

Sets the power that should be applied to a pixel in the brightness mask, giving you control over the range of very brightest (or darkest) pixels that will still impact the brightness mask blending procedure.

- A lower value lowers the Brightness Mask threshold.
- Default is 1.00.
- Range is 0.00 to 5.00 in 0.1 increments.

'Filter' Type

Specifies the type of filter to be applied to the foreground layer before layering:

- Gaussian (Fg) - Applies a Gaussian filter with a kernel size as specified in 'Filter' Kernel Radius
- Median (Fg) - Applies a median filter with a window size of (1+ ['Filter' Kernel Radius]x2)
- Mean of Medium Half (Fg) - Applies a 'mean of median half' filter with window size of (1+ ['Filter' Kernel Radius]x2)
- Minimum (Fg) - Applies a minimum filter with window size of (1+ ['Filter' Kernel Radius]x2)
- Maximum (Fg) - Applies a maximum filter with window size of (1+ ['Filter' Kernel Radius]x2)
- Lightness (Fg) - Applies a maximum and minimum filter with window size of (1+ ['Filter' Kernel Radius]x2) to the foreground layer. It then takes the mean of the minimum and maximum.
- Differential Adaptive Noise - Suppresses any noise increase in the foreground due to brightening.
- Min Distance to 1/2 Unity - chooses between foreground and background values based on which value is closest to 1/2 unity (gray). 'Filter' Kernel Radius controls the smoothness of the blend between the two.

- Max 'Contrast' - chooses between foreground and background values based on which adds most contrast to the image. 'Filter' Kernel Radius controls the smoothness of the blend between the two.
- Sobel - Performs a Sobel edge detection operation on the foreground image.
- Median Horizontal (Fg) - Performs a horizontal median filter with a horizontal kernel size of (1+ ['Filter' Kernel Radius]x2) pixels.
- Fractional Differentiation - Applies Fractional Differentiation Filtering on the foreground image. The 'Filter' Kernel Radius governs the v parameter, while alpha is fixed at 0.5. For some images this can show additional structural detail which otherwise would be hidden.
- Mode Approx. - Uses a fast method to approximate the Statistical Mode of a patch of pixels.
- Localized Histogram Equalize - performs Histogram Equalisation one patch at a time.
- Localized Histogram Optimize - calculate a non-linear transformation curve for each patch in such a way that the patch's histogram shape resembles a bell curve - This gives a much more natural look, rather than histogram equalization, and is very effective.
- Mean of Medium Half Distance Weighted - performs the same type of filtering as Mean of Medium Half - but weighted by distance.
- Local Maximum Entropy RGB Selection
- 'Gaussian High-pass filter (Fg)' - applies a Gaussian filter to the image, then subtracts that from the original. This leaves only high-frequency detail.
- 'Laplacian over Gaussian 0 crossing' - applies a Laplacian Filter for edge detection plus a Gaussian to reduce filter noise
- Default is 'Gaussian (Fg)'

'Filter' Kernel Radius

Controls a parameter of the selected filter:

- Kernel size for 'Filter' Types of: Gaussian (Fg), Median Horizontal (Fg)
- Window Size for 'Filter' Types of: Median (Fg), Mean of Medium Half (Fg), Minimum (Fg), Maximum (Fg), Lightness (Fg)

Smoothness of the blend for 'Filter' Types of: Min Distance of 1/2 Unity, Max 'Contrast'

- v Parameter for 'Filter' Types of: Fractional Differentiation
- Default is 1.0 pixels.
- Range is 1.0 to 51.0 pixels

Mask Fuzz

If a mask is used the Mask Fuzz parameter will ensure smooth, undetectable, transitions between the background and foreground images.

- This parameter specifies the kernel radius of an optional Gaussian blur, to be applied to the mask.
- Default is 1.0 pixels.
- Range is 1.0 to 101.0 pixels.

3D Stereo:
Plausible depth information synthesis for 3D capable media

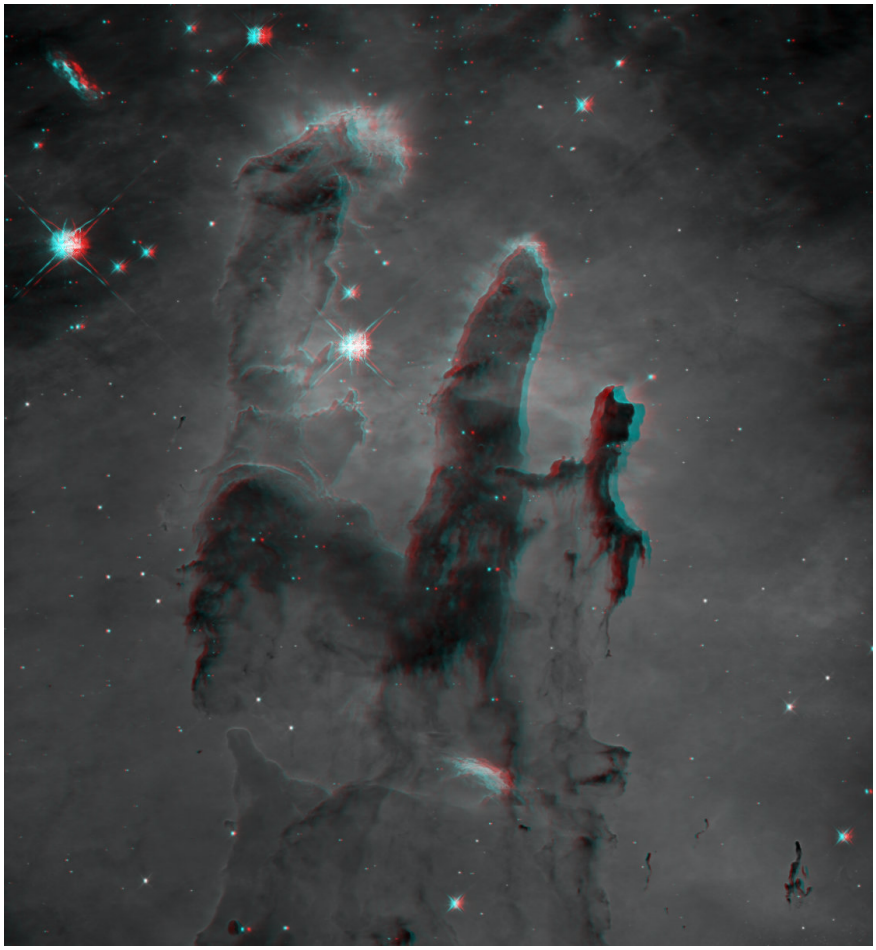

The 3D Stereo module can output images in various formats, for example in this red/cyan anaglyph format.

New as of StarTools 1.6 beta, is the 3D Stereo module. The 3D Stereo Synthesis module can be used to synthesize depth information based on astronomical image feature characteristics.

The depth cues introduced are merely educated guesses by the software and user, and should not be confused with scientific accuracy. Nevertheless, these cues can serve as a helpful tool for drawing attention to processes or features in an image.

Depth cues can also be highly instrumental in lending a fresh perspective to astronomical features in an image. The 3D Stereo Synthesis module is able to generate plausible depth information for most deep space objects, with the exception of some galaxies.

The module can output various popular 3D formats, including side-by-side (for cross eye viewing), anaglyphs, depth maps and self-contained web content HTML. The output images can be used on websites, with red/cyan glasses or within VR experiences and Facebook 3D photos.

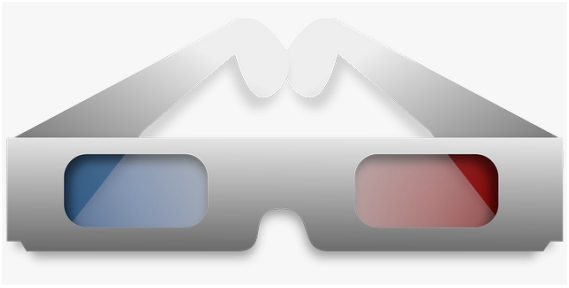

A pair of Red/Cyan glasses is a cheap but effective way of evaluating the output of the 'Stereo 3D' module.

Usage

Perceiving depth when using the module

Using the 'Stereo 3D' module effectively starts with choosing a depth perception method that is most comfortable or convenient.

By default, the Side-by-side Right/Left (Cross) Mode is used, which allows for seeing 3D using the cross-viewing technique. If you are more comfortable with the parallel-viewing technique, you may select Side-by-side Left/Right (Parallel). The benefits of the two aforementioned techniques is that they do not require any visual aids, while keeping coloring intact. The downside of these methods, is that the entire image must fit on half of the screen. E.g. zooming in breaks the 3D effect.

If you have a pair of red/cyan filter glasses, you may wish to use one of the three anaglyph Modes. The two monochromatic anaglyph modes render anaglyphs for printing and viewing on a screen. The screen-specific anaglyph will exhibit reduced cross-talk (aka 'ghosting') in most cases. An 'optimized' 'Color' mode is also available, which retains some coloring. Visual spectrum astrophotography tends to contain few colors that are retained in this way, however narrowband composites can benefit. Finally, a Depth Map mode is available to inspect (or save) the z-axis depth information that was generated by the current model.

Modeling and synthesizing depth information for astrophotography

The depth information generated by the 'Stereo 3D' module is entirely synthetic and should not be ascribed any scientific accuracy. However, the modeling performed by the module is based on a number of assumptions that tend to hold true for many Deep Space Objects and can hence be used for making educated guesses about objects. Fundamentally, these assumptions are:

- Dark detail is visible by virtue of a brighter background. Dust clouds and Bok globules are good examples of matter obstructing other matter and hence being in the foreground of the matter they are obstructing.
- Brighter areas (for example due to emissions or reflection nebulosity) correlate well with voluminous areas.
- Bright objects within brighter areas tend to drive the (bright) emissions in their immediate neighborhoods. Therefore these objects should preferably be shown as embedded within these bright areas.
- Bright objects (such as bright blue O and B-class stars), drive emissions in their immediate neighborhood and tend to generate cavities due to radiation pressure.
- Stark edges such as shock fronts tend to speed away from their origin. Therefore these objects should preferably be shown as veering off.

Tweaking the model

Depth information is created between two planes; the near plane (closest to the viewer) and the far plane (furthest away from the viewer). The distance between the two planes is governed by the 'Depth' parameter.

- The 'Protrude' parameter governs the location of the near and far planes with respect to distance from the viewer. At 50% protrusion, half the scene will be going into the screen (or print), while the other half will appear to 'jut out' of the screen (or print). At 100% protrusion, the entire scene will appear to float in front of the screen (or print). At 0% protrusion the entire scene will appear to be inside the screen (or print).
- The 'Luma to Volume' parameter controls whether large bright or dark structures should be given volume. Objects that primarily stand out against a bright background (for example, the iconic Hubble 'Pillars of Creation' image) benefit

from a shadow dominant setting. Conversely, objects that stand out against a dark background (for example M20) benefit from a highlight dominant setting.

- The 'Simple L to Depth' parameter naively maps a measure of brightness directly to depth information. This a somewhat crude tool and using the 'Luma to Volume' parameter is often sufficient.
- The 'Highlight Embedding' parameter controls how much bright highlights should be embedded within larger structures and context. Bright objects such as energetic stars are often the cause of the visible emissions around them. Given they radiate in all directions, embedding them within these emission areas is the most logical course of action.
- The 'Structure Embedding' parameter controls how small-scale structures should behave in the presence of larger scale structures. At low values for this parameter, they tend to float in front of the larger scale structures. At higher values, smaller scale structures tend to intersect larger scale structures more often.
- The 'Min. Structure Size' parameter controls the smallest detail size the module may use to construct a model. Smaller values generate models suitable for widefields with small scale detail. Larger values may yield more plausible results for narrowfields with many larger scale structures. Please note that larger values may cause the model to take longer to compute.
- The 'Intricacy' parameter controls how much smaller scale detail should prevail over larger scale detail. Higher values will yield models that show more fine, small scale changes in undulation and depth change. Lower values leave more of the depth changes to the larger scale structures.
- The 'Depth Non-linearity' parameter controls how matter is distributed across the depth field. Values higher than 1.0 progressively skew detail distribution towards the near plane. Values lower than 1.0 progressively skew detail distribution towards the far plane.

Exporting to 3D-capable media

Besides rendering images as anaglyphs or side-by-side 3D stereo content, the 'Stereo 3D' module is also able to generate Facebook 3D photos, as well as interactive self-contained 2.5D and Virtual Reality experiences.

Standalone Virtual Reality experience

The standalone VR experiences are compatible with anything from the latest headsets to the sub-$5 Google Cardboard devices.

The 'WebVR' button in the module exports your image as a standalone HTML file. This file can be viewed locally in your web browser, or it can be hosted online. It renders your image as an immersive VR experience, with a large screen wrapping around the viewer.

The VR experience can be viewed in most popular headsets, including HTC Vive, Oculus, Windows Mixed Reality, GearVR, Google Day Dream and even sub-$5 Google Cardboard devices.To view an experience, put it in an accessible location (locally or online) and launch it from a WebVR/XR capable browser.

Please note that landscape images tend to be more immersive.

Standalone interactive 2.5D web viewer

The 'Web2.5D' button in the module exports your image as a standalone HTML file. This file can be viewed locally in your web browser, or it can be hosted online.

Depth is conveyed by a subtle, configurable, bobbing motion. This motion subtly changes the viewing angle to reveal more or less of the object, depending on the angle. The motion is configurable both by you and the viewer in both X and Y axes. The motion can also be configured to be mapped to mouse movements.

A so called 'depth pulse' can be sent into the image, which travels through the image from the near plane to the far plane, highlighting pixels of equal depth as it travels. The 'depth pulse' is useful to re-calibrate the viewer's perspective if background and foreground appear swapped.

Hosting the file online, allows for embedding the image as an IFRAME. The following is an example of the HTML required to insert an image in any website:

```
<iframe scrolling='auto' marginheight='0' marginwidth='0' style='border:none;max-
width:100%;' src='https://download.startools.org/pillars_stereo.html?
spdx=4&amp;spdy=3&amp;caption=StarTools%20exports%20self-contained,
%20embeddable%20web%20content%20like%20this%20scene.
%20This%20image%20was%20created%20in%20seconds.%20Configurable,
%20subtle%20movement%20helps%20with%20conveying%20depth.' frameborder='0'></iframe>
```

The following parameters can be set via the url:

modex: 0=no movement, 1=positive sine wave modulation, 2=negative sine wave modulation, 3=positive sine wave modulation, 4=negative sine wave, 5=jump 3 frames only (left, middle, right), 6=mouse control

modey: 0=no movement, 1=positive sine wave modulation, 2=negative sine wave modulation, 3=positive sine wave modulation, 4=negative sine wave, 5=mouse control

spdx: speed of x-axis motion, range 1-5

spdy: speed of y-axis motion, range 1-5

caption: caption for the image

StarTools exports self-contained, embeddable web content like this scene. This image was created in seconds. Configurable, subtle movement helps with conveying depth

Facebook 3d photos

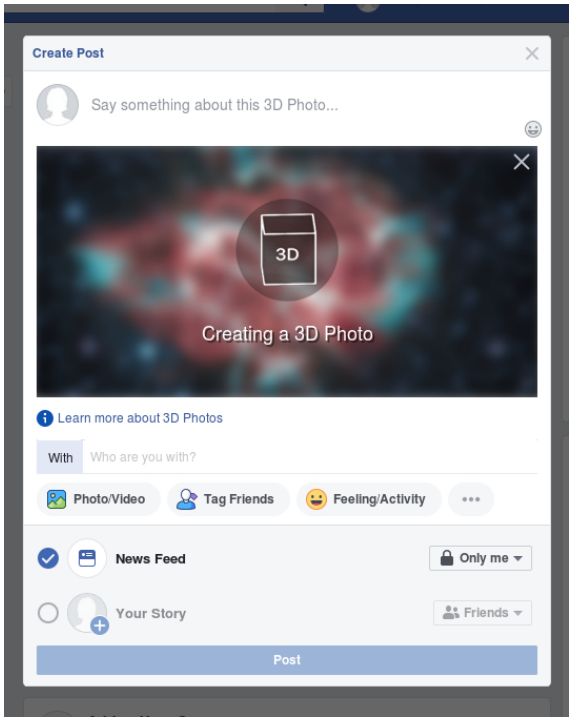

The 'Stereo 3D' module is able to export your images for use with Facebook's 3D photo feature.

The 'Facebook' button in the module saves your image as dual JPEGs; one image that ends in '.jpg' and one image that ends in '_depth.jpg' Uploading these images as photos at the same time will see Facebook detect and use the two images to generate a 3D photo.

Please note that due Facebook's algorithm being designed for terrestrial photography, the 3D reconstruction may be a bit odd in places with artifacts appearing and stars detaching from their halos. Nevertheless the result can look quite pleasing when simply browsing past the image in a Facebook feed.

3D-capable TVs and projectors

TVs and projectors that are 3D-ready can - at minimum - usually be configured to render side-by-side images as 3D. Please consult your TV or projector's manual or in-built menu to access the correct settings.

Wrap Up

Description

To synthesize depth information based on astronomical image feature characteristics to allow displaying 3D-like images in various ways. This can be an important tool in interpreting the subject - making us understand its structure - and guiding us to process the data to bring out that structure. This module synthesizes plausible depth information to allow images to appear three-dimensional.

Images can be created for viewing with: Red/Cyan glasses, VR Headset, Facebook 3D photo feature, or in standard web browser using the 2.5D option.

This depth information is not in the data but is suggested based on common rules and assumptions about gas and radiation behavior.

When to use

- When Tracking is Off.
- Standalone - on any image processed previously. No need to reprocess.
- At end of workflow.

What result to look for

- Image should give a sense of three-dimensional solidity.

Ways of getting better results

- Often less is more. Overdoing it may spoil the illusion
- For cross-eye or parallel viewing, portrait images are recommended
- For Virtual Reality and anaglyph viewing, landscape images are recommended

Description of Controls

Exporting Presets

In addition to the type of image being produced which is set by the Mode control.
The image can be output to support different 3D representations using these buttons:

- Web 2.5D - Exports as an HTML file for viewing using a browser.
- WebVR - Exports as an HTML file for viewing using a Virtual Reality headset.
- Facebook - Exports as an JPEG file for viewing using the Facebook 3D Photo feature.

Simple L to Depth

Does simplistic depth mapping - assumes lighter stuff is nearer than darker.

- Larger values mean brightness will be taken more as an indication of depth.
- 0% means that brightness is not taken as an indicator of depth.
- It is usually sufficient to leave this at 0% and rely on the Luma to Volume control.
- Default 0%.
- Range from 0% to 100%.

Mode

This sets the output mode. Select based on the type of output needed for the way it is to be viewed.

- Side-by-Side Right/Left (Cross) - For Cross-eyed depth perception without aids. The left hand image is for the left eye, the right hand image is for the right eye.
- Side-by-Side Right/Left (Parallel) - For Parallel depth perception without aids. The right hand image is for the left eye, the left hand image is for the right eye.
- Anaglyph Mono (screen) - For use with red/blue filter glasses - Mono image that minimizes ghosting on sRGB calibrated screens.
- Anaglyph Mono (print) - For use with red/blue filter glasses - Mono image suitable for printing.
- Anaglyph 'Color' - For use with red/blue filter glasses - Keeps some color tones intact.
- Depth Map - Outputs the depth map. White is the near plane, black is the far plane.
- Default is 'Side-by-Side Right/Left (Cross)'.

Depth

The 3D effect is obtained by modeling two planes. A near plane closer to the viewer and a far plane further from the viewer.

- The Depth parameter defines the distance between the planes.
- The smaller the value the less pronounced will be the 3D effect.
- The larger the value the more pronounced will be the 3D effect.
- Default is 21 pixels.
- Range from 0 to 42 pixels.

Protrude

The Protrude parameter defines how much of the scene appears to be in front of the image and how much behind.

- At 0% the entire scene will appear to be inside the screen or print.
- At 100% the entire scene will appear to be in front of the screen or print.
- Default 0%.
- Range from 0% to 100%.

Highlight Embedding

This controls whether highlights appear embedded within the larger structures around them.

- Bright objects often cause emissions around them. It makes sense to embed these.
- The higher the value the more embedded the highlights are.
- Default 25%.
- Range from 0% to 100%.

Min Structure Size

Controls the size of the smallest details that are acted on.

- Smaller values are best for wide fields with lots of small detail.
- Larger values are best for narrow fields with many larger scale structures.
- Larger values may take longer.
- Default 1.6 pixels.
- Range from 1.0 to 5.0 pixels.

Structure Embedding

This controls whether small-scale structures should appear embedded or intersect larger scale structures.

- Lower values mean more structures float in front of larger structures.
- Higher values mean more smaller structures intersect larger structures.
- Default 25%.
- Range from 0% to 100%.

Depth Non-Linearity

This controls how matter is distributed across the depth.

- A value of 1.0 distributes matter evenly.
- Values greater than 1.0 progressively skew distribution towards the near plane.
- Values less than 1.0 progressively skew distribution towards the far plane.
- Default 1.00
- Range from 0.00 to 5.00.

Luma to Volume

This controls how brightness acts as a predictor for volumetric detail.

- Larger 'shadow dominant' values (slider to left) give larger dark structures more volume. Objects against a bright background.
- Larger 'highlight dominant' values (slider to right) give larger bright structures more volume. Objects against a dark background.
- Default 50% highlight dominant.
- Range from 100% shadow dominant to 100% highlight dominant.

Intricacy

This controls what size detail is distributed in depth.

- Lower values show more depth changes for larger scale structures.
- Higher values show more depth changes for the finer detail.
- Default 50%.
- Range from 0% to 100%.

Background Information

Mask Module

Mask Blink control

- To stop the Mask blinking three times as a reminder it is set, put a file called 'noblink' in the StarTools executables folder. Instead of blinking it will just show the mask once for a few seconds.

Creating a mask in GIMP 2

You can create a mask in another program and import it into StarTools as a Tiff file.
So you can use brushes and other tools in PhotoShop or GIMP to create the mask, save it as a tiff, and then Open it in StarTools.
Here are some instructions for creating a Mask in GIMP 2 (These instructions have been tested on version 2.8):

1. Load a saved version of the image into GIMP 2.
2. In the 'Layer' menu - select 'New Layer' - Fill Type: 'Transparency' - this creates a new 'Layer'.
3. If the 'Layers' dialog is not visible open it using Ctrl-L.
4. If the 'Tools Options' dialog is not visible - open it using the Windows menu - Dockable Dialogs - 'Tool Options'.
5. From the Tools menu select 'Paint Tools' - then select Paintbrush (or perhaps another tool e.g. Airbrush).
6. In the 'Tool Options' dialog adjust the 'Brush' settings - such as Size and other settings as needed.
7. In the 'Layers' dialog select the new 'Layer'.
8. Select (in black) the elements of the image you want masked using the brush - or do the opposite and invert it later.
9. When done, select the original image layer - right-click and select 'Delete layer'. This leaves the mask layer.
10. In the Image menu, select 'Flatten Image'
11. In File menu - select 'Export As...' - Select 'File Type (by extension)' - 'TIFF Image', name the image, then press 'Export'.
12. In StarTools, click the 'Mask' button at the top of the main screen.
13. Click 'Open' and select the newly created mask. The mask should open with the black parts now the inactive (transparent) parts. Invert if needed.

Compose Module

Colours and Filters

Eyes Response to Light

The generally accepted range of visible light is from 400nm-700nm although this varies from person to person and with different light levels. or more information see Eye spectral response.

- Photopic: 50% 510nm-610nm, 5% 450nm-660nm
- Scotopic: 50% 455nm-550nm, 5% 410nm-600nm

RGB Filters

The RGB filters commonly used have a spectral response as follows:

- R: 600-700nm
- G: 500-600nm
- B: 425-500nm

Narrow-band imaging

Using different narrow-band filters to create three data sets and combining them using LRGB channels can have advantages:

- Allows imaging through a lot of light pollution (other than LED broad spectrum light pollution).
- Nebulosity appears more detailed.
- Stars seem less bloated.
- Colours are artificial - so can be chosen to highlight features - e.g. Hubble Palette better differentiates SII and Ha.

Common narrow band data collected is:

- SII - Red - 671.9 nm and 673.0 nm
- Ha - Red - 656.3 nm
- NII - Red - 654.8 nm and 658.3 nm
- OIII - Cyan - 495.9 nm and 500.7 nm
- Hb - Blue - 486.1 nm
- OII - Near U-V - 372.6 nm and 372.9 nm

Three colour narrow-band images:

Common combinations for assigning the three bands to the R,G and B channels are:

- SII : Ha : OIII (R:G:B) - Hubble palette.
- Ha : OIII : SII (R:G:B) - Canadian-France-Hawaii Telescope palette.

Bi-colour narrow-band images:

For a number of reasons, including reduced imaging time needed over 3-colour images, bi-colour imaging can be an attractive approach.

Choosing the right target is the key to getting a satisfying result. See the external article Narrowband Bicolor Palette Combinations.

Bi-colour colours - Getting the colours 'right' can be tricky. Areas where the two sets of data are the similar strengths look pale and washed out. See this post:

Style and LRGB Emulation in Bicolor

There are many possible bi-colour combinations of assigning the two sets of narrow-band data to the R,G and B channels - common ones are (R:G:B):

- Ha : OIII : OIII
- Ha : Ha+OIII : OIII
- Ha : (Ha+OIII)x1.5 : OIII - Steve Cannistra's technique - see description in Special Techniques section above.
- Ha : Hax0.4 + OIIIx0.6 : OIII
- SIIx0.5+ Hax0.5 : Hax0.4 + OIIIx0.6 : OIII

Synthetic H-beta

Wherever Ha appears there is often also Hb. You can capture the Hb with a filter but you can also easily create a synthetic Hb.

- See Synthetic RGB from [SII], Halpha and [OIII] Emission Line Data and Narrowband imaging with only H-alpha and OIII filters
- The ratio of Ha:Hb emission is normally about 3:1 for emission nebulae and 6:1 for planetary nebulae.
- The shorter wavelength Hb has a higher extinction so the perceived amount of Hb is lower.
- Common to Combine OIII and Synthetic Hb into the blue channel as follows: OIII x 0.85 + Ha x 0.15 for an emission nebula with variation depending on the source.
- For a list of targets see the CN list of advanced narrow band imaging targets.

'AutoDev' Module

The 'AutoDev' module, along with the 'FilmDev', 'Contrast' and 'HDR' modules, provide automated stretching at all levels of detail. These modules replace the traditional curve-based adjustment of stretching.

Dark Anomaly Headroom and more..

Dark Anomalies

These are - in StarTools - pixels (or clumps of pixels) that are darker than the 'real' background. E.g. they don't describe real signal, but rather some form of signal obstruction or hindrance (dust, dead pixel, tree, stacking artifacts etc.).

A number of modules in StarTools really, really hate dark anomalies as they rely on taking measurements of the 'real' background for various purposes. Dark anomalies can really throw off these measurements. Therefore filters like the Dark Anomaly filter exists for filtering these anomalies out if they are small, or masks for bigger areas/clumps.

The question in some modules is then, how should a module treat this anomalous data once it has been 'worked around'/'ignored'? Should it just clip it to black/0? Maintain the original value? Something in between?

The 'Shadow Dynamic Range Allocation' parameter in 'Contrast' module does exactly this: it controls the headroom (dynamic range) allocate to dark anomalies. No headroom = clip to black, full headroom = keep as-is.

Usually dynamic range taken up by dark anomalies is better allocated to real detail, but the choice is yours.

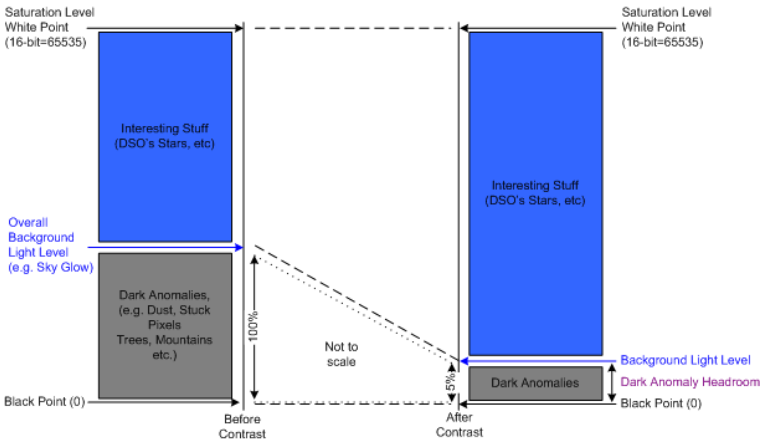

Dark anomaly dynamic range is 'crunched'

(the 5% vs 100% dark anomaly headroom is not to scale)

You can see that by crunching down the dynamic range allocated to dark anomalies, the average background level now sits much closer to black. This effectively makes dark anomalies harder to see, while re-purposing the freed up dynamic range for the 'Interesting stuff'.

Modules that don't explicitly offer a dynamic range control over dark anomalies use 100% for any remaining dark anomalies after any filtering. E.g. they largely retain the dynamic range allocated to dark anomalies for any large anomalies not caught by small anomaly filtering.

To be clear, the global background level is being reduced to the (squashed) amount of the brightest (approaching background level in brightness) dark anomaly. Dynamic range below that level is reserved to still describe the dark anomaly, however now there is much less headroom - it is 'squashed' - to do that.

'Lens' module

Optical Distortion

- The 'Lens' module in StarTools uses Brown's distortion model to correct for 'simple' radial distortion, which can alleviate some of the worst effects of coma. Ref: Wikipedia article on distortion in optics.
- Image distortion may be off-centre for a number of reasons, such as the camera not sitting perfectly flush/perpendicular to the imaging plane, or movement between sub-frames affecting the stacking process.

Chromatic aberration

For a general introduction to the subject see the article: Chromatic aberration.
This can be split into two types:

- Transverse or Lateral chromatic aberration - color shifts due to the different channels not quite lining up - this is the type that the 'Lens' module addresses. With transverse chromatic aberration the color fringes are on one axis.
- Axial or Longitudinal chromatic aberration - caused by different wavelengths/ colors coming into focus at slightly different distances - this can be mitigated using the 'Decon' or 'Filter' modules. With axial chromatic aberration the color fringes are all around the object.
- For further information see this article on Chromatic Aberration.

'Bin' Module

Oversampling

- An image is oversampled if the image resolution is greater than the image detail available. The extra resolution doesn't contribute to improved detail.
- For example, if you print a 2x oversampled image and print the same image after 2x binning (so it is not oversampled) enlarged to the same size - the detail visible will be the same in both.
- As a general rule, if the smallest unsaturated star occupies 3 or more pixels in any direction you can bin the data some more without losing detail.
    - This rule applies to monochrome cameras and also for OSC/DLSR cameras where the image has been stacked with a sufficient number of sub-frames which have been dithered.
    - The rule needs to be adjusted for less good OSC/DSLR images. At the extreme, for a single OSC/DSLR sub-frame, a star will be spread across a minimum of 3 pixels by the debayering algorithm - so the rule in this case if a star occupied 3 pixels it would no longer be oversampled.
- The major benefit of oversampling is that, using the right algorithm, you can bin it to improve the signal to noise ratio (SNR) of the image.
- It can also be used by deconvolution routines - so it may be beneficial to leave a bit of oversampling for use by the 'Decon' module.
- If we assume we have perfect focus and optics then to get the best information from the seeing we need to have an image scale of about half the seeing.
- Image scale (arcseconds per pixel) = 206 * camera pixel size (microns) / Focal length (mm)
    - For many DSLR camera and telescope combinations the image scale is around 1 arcsec/pixel or less.
    - E.g. Canon 1100D (5.2um pixels) - with 200mm f5 reflector = 1.07 arcsec/ pixel, with 80mm f6 refractor = 2.23 arcsec/pixel

- Seeing:
    - An Average night at an average site - 4-5 arc seconds
    - A Good night at an average site - 2-3 arc seconds
    - The Best conditions at the best sites - 0.5-1 arc second

As a result, for many conditions we meet, stars can be spread over a number of pixels. Luckily we can take advantage of this wasted resolution for the conditions by binning and improving the SNR.

'Bin' vs Deconvolution

The 'Bin' algorithm can take advantage of oversampling to improve signal to noise. The deconvolution algorithm used by the 'Decon' module can take advantage of the oversampling to recover detail. So how should we decide the amount? Should we leave some oversampling for deconvolution to use?

- 'Binning more will improve the SNR and allow 'Decon' to go 'deeper' and successfully deconvolve features that are darker/buried. Other modules will also benefit from the improved SNR.
- 'Binning less will accomplish the opposite (and the higher noise will impact other modules as well) but finer detail may be recovered and described (i.e. kept at higher resolution) in areas where signal is high.

So - it all depends - on whether you can get more detail using 'Decon' that justifies the degradation of the results by not getting every last bit of SNR available. For most noisy images binning to use all the oversampling will give best results.For higher quality images leaving some oversampling may be beneficial. Initially probably the best way of deciding is to try it. See what results you can get when you go through the workflow (including 'Decon') when keeping some oversampling - as to what you can get using the same workflow by binning as much as you can.

Make sure that the overall image has a high enough resolution - with the prominent stars keeping their roundness and the prominent detail not artificially jagged and angular. High quality printed images are generally about 300 pixels per inch - although 240 pixels per inch may be sufficient.

Hardware binning vs. StarTools binning

Hardware binning :

- Reduces noise by reducing the amount of read noise e.g. 2x2 binning means 4 pixels are read at once - so a quarter of the read noise.
- It is fixed at capture time - there is no way of reducing it later.
- May give better results where read noise is dominant over shot noise - such as with very faint objects under really dark skies.

StarTools software binning:

- Read noise is the unchanged because the pixels are read individually
- Shot Noise is reduced - Example: 2x2 binning - signal increases by 4x but noise increases by sqrt(4)=2 so SNR increases by 2 - assuming read noise is negligible
- Allows fractional binning (e.g. 2.1x2.1)
- Allows experimentation after capture
- Better results when there is skyglow or light pollution (i.e. shot & other noise is dominant over read noise)
- If in doubt - use StarTools binning unless target is a faint object under a really dark sky.

'Wipe' Module

Dark Anomalies

Every night sky image that is still linear has a minimum background light level caused by the sky glow. This is represented by a certain pixel value. If there are pixels with values below this then they don't represent something in the sky. These values are either caused by hot or stuck pixels, dust or scratches, noise, or something terrestrial like a tree or mountain. StarTools calls these Dark Anomalies.

Dark Anomalies stop some modules from working correctly so they need to be identified so they can be managed.

Dark Anomaly Headroom

The 'Wipe' module looks for the global background level and subtracts that from the signal to reveal more detail. The only time this is a problem is when there is no global background - i.e. when you are not looking at the sky (trees & mountains or dust) or when there is a faulty pixel. In this case the subtraction would produce a negative value and so clip to 0.

To avoid this, Dark Anomalies are detected and some dynamic range is allocated for these.

The background level is raised for the whole image so that when the signal level is subtracted from it is always greater than 0. The amount the background level is raised by is called the Dark Anomaly Headroom.

In the 'Wipe' module the Dark Anomaly Headroom allows all the dark anomalies to be mapped directly- without modification. The 'Contrast' module allows some compression of the Dark Anomaly Headroom.

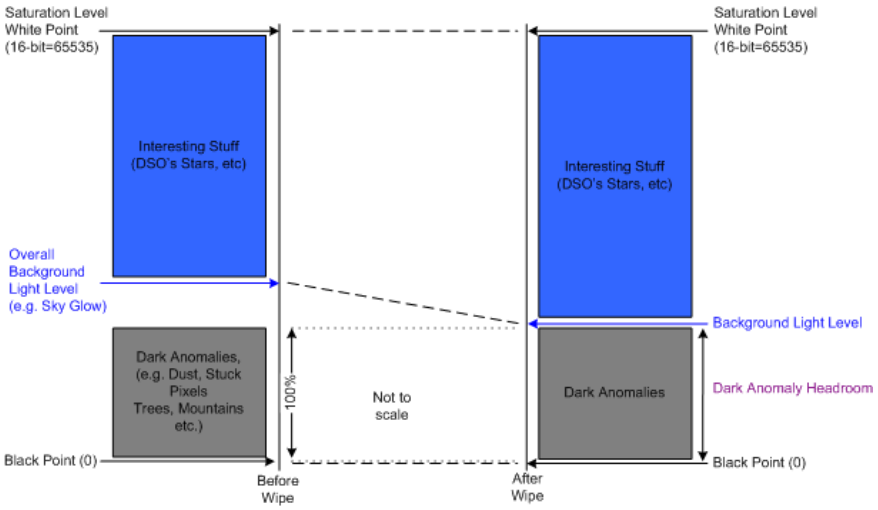

'Wipe'

Dark Anomaly 'Filter' and Mask

'Wipe' tries to find this background level so as to identify the minimum and maximum pixel values that represent light from the sky. Dark Anomalies are dips below the minimum level which would cause 'Wipe' to set an artificially low background level.

In order for StarTools to find the correct background level it uses two methods to identify and ignore Dark Anomalies

- Mask - Larger Dark Anomalies, such as mountains or trees, as well as dust bunnies can be identified using a Mask.

- Dark Anomaly 'Filter' - The Dark Anomaly 'Filter' lets 'Wipe' automatically ignore small clusters of pixels with very low values when deciding the background level. The value of the Dark Anomaly 'Filter' sets the maximum pixel cluster size that will be ignored. Set to 1 pixel to ignore single pixels such as cold or stuck pixels. Set higher to ignore clusters caused by noise or small dust specks or scratches.

'Color'/Luminance button

In 'Compose' mode the 'Wipe' module works on the luminance and color datasets separately but in parallel. You can see the results of each by pressing this button.

Calibration Modeling

The 'Wipe' module models and corrects for three stages of calibration:

- Synthetic bias/darks modeling and correction (subtraction)
- Synthetic flats modeling and correction (division)
- Gradient modeling and correction (subtraction).
- Any issues selected in the parameters and/or detected are modeled during the correct stage and feeds into the next.

'Synthetic Bias/Darks Modeling

Modeling and correction for Bias and Darks issues is new to 'Wipe' in v1.7 and makes the Band module obsolete. A number of different sensor defects can be detected and repaired.

'Contrast' Module

The 'Contrast' module uses a similar algorithm to that used by the 'Wipe' module. As a result the effect of the 'Wipe' module can be to reduce or eliminate the effectiveness of the 'Contrast' module.

The Image Scale that 'Contrast' works on

The 'Contrast' module works on medium-to-large areas. However terms like 'medium-to-large' and 'medium-to-small' are actually (and unhelpfully) image dependent.

It roughly depends on the entropy of the image at different scales (roughly corresponding with 'busyness' at different scales, which itself roughly corresponds with the angular size of the image).

As a result, the same settings for 'Contrast' may yield different effects for different images.

If the entropy at large scales is low (for example the presence of a 'fog' amongst a busy star field), 'Contrast', regardless of (most) settings will have the effect of lifting that fog.

If that fog did not exist and we just had a busy star field, 'Contrast' would appear to do very little.

The Dark Anomaly 'Filter' in the 'Contrast' Module

This works exactly the same as in - for example - the 'Wipe' module. However, its effect is much more pronounced as this is the only module that is intended for stretched data (and its effects are much easier to see). Somewhat similar to the 'Wipe' module, the 'Contrast' module models large scale light distribution and subtracts this from the image to reveal smaller scale detail.

What happens to the image when more is subtracted than what the local background (as determined by the filter) allows for, is determined by the Dark Anomaly Headroom setting?

Without any headroom allocated (e.g. 0%), the Dark Anomaly 'Filter' clips pixels to 0 (pure black) if more is subtracted than there is signal. The only time where there is not enough signal, is when we are dealing with a dark anomaly. Any other time the global background should provide enough signal for subtraction without clipping.

To avoid clipping clipped pixels are not set to 0, but instead the global background is raised (for the whole image!) to accommodate the subtraction. The every undershoot is, however scaled by the Headroom parameter. This means that the global background level may be raised less than 100%, thereby 'squashing' undershoot (e.g. darker-than-background pixels) dynamic range, while leaving more dynamic range for the pixels that are brighter than the background.

You can see, for example that varying the Dark Anomaly 'Filter' from 1px to 20px (or beyond) has no effect when Headroom allocation is set to 100%. This makes sense, as the background is simply raised/restored to what it was before the subtraction (e.g. 100%).

'HDR' Module

Old vs New HDR

- Old 'Detail Size' is related to new 'Context Size' but 'Detail Size' increased processing linearly not exponentially.
- There is a square root relationship. An old 'Detail Size' of 1000 is roughly equivalent to a new 'Context Size' of 32.
- In the old module the highest 'Detail size' was equivalent to a very low 'Context Size' in the new module.
- The old module offered a Gamma Correction algorithm or a Detail Boost algorithm
- The new module combines both Gamma Correction and Detail Boost

Performance

Comparative tests done with STReplay show that changes to Context Size greatly affect the processing time for this module.

In a simple test the time taken for the processing to stop (wheel disappears) relative to a Context Size of 50 were as follows:

CODE: SELECT ALL

Context Size	50	40	30	25	20	10
%Processing time	100	50	15	6	4	2

Effect of Quality Setting

Approximate Relative Processing Time (Medium =100%): Low 98%, Medium 100%, High 220%.

This may vary with environment/settings.

Improving Performance

If image oversampled consider binning

Reduce Context Size to about 30. Increase from there as needed.

Reduce Quality to 'Medium' or 'Low' and see if you still get acceptable results.

'Sharp' Module

The 'Sharp' module uses Wavelet (aka Laplacian) 'Sharpen'ing techniques to enhance the detail in the image. See How Unsharp Masking and Laplacian 'Sharpen'ing Work

Scale and regions affected

- The scale is based on the number of pixels - so if you use the same settings on an un-binned image and the same image 50% binned they will have different scales affected.
- For each scale increment the size increases exponentially, based on Structure Size. For Small the next scale size is 2x the current one, For Medium it is 2.75x and for Large it is 3.5x.

- If you want to see what size a particular scale affects - set that Scale to 100%, the Amount to 1000%, and all other scales to 0%. Do Before/After to see the scale affected.

Scale aware processing

Here's generally how to do it:

- Make two (or more - depending on how many scales you want) copies of the image.
- Blur one with, for example, a Gaussian blur.
- Subtract the blurred image from the non-blurred image. The blur will have killed most of the fine detail (it's a blur after all), so subtracting the blurred image from the non-blurred image logically leaves you with just the fine detail.
- You can now manipulate that fine detail - for example normalize it, use brightness/contrast operations on it, etc.
- Add it (the modified fine detail) back to the blurred (coarse detail) image and voilà: you have an image that has had just a particular scale manipulated.

You could even do all this in something like The GIMP or Photoshop with the caveat that some additional trickery is required to manipulate and visualize pixels with a negative value. Now, nothing of course stops us from grabbing that fine detail we isolated, perform the same trick on that, so that you can start getting 'bands' of different detail sizes.

The radius of the kernel you use for the Gaussian blur in this case defines the image scale peak response.

'SVDecon' Module

Deconvolution and Oversampling

Deconvolution is used to undo the blurring effect of an unstable atmosphere and an imperfect optical train, however this can amplify noise and introduce artefacts including ringing. It works best on data which is oversampled and has a high signal to noise ratio.

However, data that is on the cusp of being oversampled, where faint stars are spread over 3 pixels, may still benefit from a small amount of deconvolution. Every optical system, no matter how expensive, spreads a point light over multiple pixels to a degree (see Airy disk). SVDecon can reverse this spreading as well - just take it easy - it is very easy to overdo this.

For further details regarding oversampling, binning and deconvolution see the Bin module background notes.

Deconvolution and Binning

SV Decon should definitely still be able to yield good results and meaningful improvements on binned data. Even for data that is undersampled. The only potential issue you may be a few more samples are required (if using sampling), as reconstructing PSFs from undersampled data is a lot harder (providing more samples essentially has the same effect of "drizzling").

Apod Mask Color Indicators

The colors give an indication of the quality of the signal (linearity, SNR and gradient slope). The better the signal, the less noise the sample will have and the better (more cleanly) it will describe the PSF.

Deconvolution and Singularities

Singularities in the data are those areas where there is a discontinuity in the valid data - where the valid data is missing - such as in the saturated white cores of stars. These areas normally cause bad ringing artifacts.

StarTools uses a novel de-ringing algorithm which ensures stars are protected from the Gibbs phenomenon (also known as 'panda eye effect'), while actually being able to still coalesce singularities, such as over-exposed white cores of stars, into point lights.

Lunar/Planetary/Solar

If the image is of Planetary and Lunar and other non-DSO subjects
Selecting this mode clears the apodization mask
'Synthetic Iterations' are increased to 50x
'Synthetic PSF Radius' is increased to 10.5 - something more suitable to high magnification.
'Sampled Iterations' are automatically set to 'Off' (0), and the sampling mode disabled, when this is selected.
Star samples, if any, would not help for these sorts of images where there is a large starless area in the centre of the image.

Sample Stars and deconvolution

Sample Stars: Good candidates for sample stars have the following characteristics:

- Are set in an even background.
- Are neither over-exposed or dim.
- Ideally their profile covers most of the linear dynamic range of the image.
- Ideally they are located towards the centre of the image.

Spatial variation of the Point Spread Function

The spatial variance of the Point Spread Function can be due to any issue which has deformed the stars in a non-uniform way - such as:
field rotation
field curvature
tracking error
coma
camera mounting stability
some other issue which has deformed the stars in a non-uniform way
Apodizaton mask
Identifies the places where the apodization function is applied.
The apodisation function acts to remove ringing artifacts.

Tracking

Tracking in StarTools is the name given to the way in which StarTools gathers information about the signal and its evolution through different modules.
It provides each module with as much information as possible to allow it to get the best results.
Each module can:
Understand how each pixel has been modified by previous modules.
Influence data earlier in the processing chain (e.g. linear data) and re-apply the modifications made since then.
Use the information gained from previous modules to understand how the signal has been changed and where the noise is.

Some benefits of Tracking:

Deconvolution only works on linear data - but the SVDecon Module is used after the data has been stretched.
In the SVDecon module we look at the result of a stretched and processed image and apply deconvolution on the linear data and watch its effect on the processed image.
Noise reduction is applied at the end of processing where Tracking has gained the most information about noise.
Noise reduction is automatically targeted most at the areas where it is needed.

Lunar/Planetary/Solar

The SVDecon module behaves slightly differently with Lunar, Planetary and Solar images when it comes to reconstructing highlights.

With these images if a reconstructed highlight requires more dynamic range it is allocated it. The reconstructed highlights are not allowed to over-expose. The dynamic range of the complete image is adjusted to accommodate the new highlights.
With DSOs, if a reconstructed highlight is not given any more dynamic range and is allowed to over-expose.
Also, these images do not require an aggressive deringing strategy.
See also StarTools help on Lunar, Planetary and Solar

Models of the atmosphere

The way a point source has its light scattered around its actual location is called a Point Spread Function (PSF)
Deconvolution does its best to model this PSF and then reverses it to get back to the original.
Over the years models have been developed for the PSF of atmospheric blurring.
Five of these models are available to select from in the SVDecon Module.

- Gaussian (Fast)
  - Uses Gaussian distribution to model atmospheric blurring
  - Model used previous to Startools 1.6
  - Fast
- Circle of Confusion
  - Models the way a lens focuses the light of a point source assuming no atmosphere
  - Suitable for images taken outside the Earth's atmosphere.
- Moffatt Beta=4.765 (Trujillo)
  - Uses Moffat distribution to model atmospheric blurring
  - Uses a Beta factor 4.765.
  - Recommended by Trujillo et al (2001) [PDF] as best fit for prevailing atmospheric turbulence theory
- Moffatt Beta=3.0 (Saglia, FALT)
  - Uses Moffat distribution to model atmospheric blurring
  - Uses a Beta factor of 3.0.
  - This is a rough average of the values tested by Saglia et al (1993)[PDF]
  - It corresponds with findings of Bendinelli et al (1988)[PDF]
  - As a result of studying the Mayall II cluster - Implemented as the default in the FALT software at ESO.

Moffatt Beta=2.5 (IRAF)
Uses Moffat distribution to model atmospheric blurring
Uses a Beta factor of 2.5.
As implemented in the IRAF software by US National Optical Astronomy Observatory.
Selection of Beta for Moffat PSF
A Moffat PSF with a value of 10 is similar to a Gaussian PSF.
Pixinsight defaults to a Beta of 4
Trujillo, Saglia, Bendinelli and others all did evaluations and came to answers varying from 2.5 to 4.765
Take your pick - try each one out and see what you think.The 'Decon' module behaves slightly differently with Lunar, Planetary and Solar images when it comes to reconstructing highlights. With these images, if a reconstructed highlight requires more dynamic range it is allocated it. The reconstructed highlights are not allowed to over-expose. The dynamic range of the complete image is adjusted to accommodate the new highlights. With DSOs, a reconstructed highlight is not given any more dynamic range and is allowed to over-

expose.Also, these images do not require an aggressive deringing strategy so 'Deringing' can be set to 0%.

Tracking

Tracking in StarTools allows the 'SVDecon' module to do things that are normally impossible!

- Deconvolution only works on linear data - but the 'Decon' Module is used after the data has been stretched.
- In the 'SVDecon' module we look at the result of a stretched and processed image and apply deconvolution on the linear data and watch its effect on the processed image.

'Flux' Module

The way Fractal 'Flux' works is by analyzing the image larger scale structures and gas flows and then tries to apply what it has learned about the larger structures to the smaller structures.

In wide fields, the gaseous structures start behaving/looking materially different than at smaller scales (kind of like how cloud patterns seem to look different from the surface of the earth than they look from high up in outer space where you can see swirls and distinct low/high pressure areas, cloud fronts, etc.), so using very large structures (thousands of light years across) to base smaller structures on when the difference is big, starts becoming harder. Therefore the results will start looking less plausible.

'SuperStructure' Module

Adding Life - Modeling Light Diffraction

The 'SuperStructure' module tries to compensate for the fact that some sorts of processing can cause bright objects to lose their 'glow'. It assumes the light went through a perfect lens with a circular aperture and produced an Airy Disk. The 'SuperStructure' module tries to recreate the original Airy Disk by taking the energy represented by the sample pixels and redistributing that energy into an Airy disk. This is done using a point spread function that produces an Airy Disk. The calculated Airy Disk is then added back into the original image.

Airy Disk

The Airy Disk is defined by the following settings:

- Glow Threshold - Defines how bright a pixel needs to be before it is included in the Airy Disk diffraction model. This defines the area that is to be redistributed.
- Airy Disk Radius - Sets the radius of the point spread function - dictates how the resultant Airy Disk looks.
    - slow (small diameter, more 'blur', showing individual rings and a big bulging center) - suits best for closeups
    - fast (large diameter, less 'blur', showing diffraction patterns) - suits best for widefields
- Airy Disk Sampling - defines the accuracy with which the point spread function defines the Airy disk. Greater accuracy takes longer.

Adding the Airy Disk back into the image

The calculated Airy Disk diffraction model is added back into the image in two steps.

- Creating a composite of the model and the original image
- Selectively merging the composite with the original image.

The process is controlled by the following settings:

- Compositing Algorithm - this defines how the contribution from a pixel in the Airy Disk is combined with the contribution of the corresponding pixel in the image (e.g. Screen, Power of Inverse, Multiply, Gamma Correct) to create a composite.
- Detail Preservation - Defines how to decide between the detail in the image and the Airy Disk, or the proportions of both. This is done by selecting the algorithm to use in merging the composite of the model and the original image back into the image. Takes the Detail Preservation Radius as a parameter.
    - Off - None of the detail is preserved - the area to be redistributed is completely replaced by the new Airy Disk.
    - Linear Brightness Mask - uses a brightness mask with a linear relationship between input and masking - and with the composite of Airy Disk model and original as input - so, as the composite gets dimmer it masks out more of the composite and uses more of the original instead.
    - Max 'Contrast' - chooses between the composite and image detail values based on which adds most contrast with its neighboring pixels. Takes the Detail Preservation Radius as a parameter.
    - Min Distance to 1/2 Unity - chooses between the composite and the image detail values based on which value is closest to 1/2 unity (gray). Pixels closest to gray tend to show the most human detectable detail. Takes the Detail Preservation Radius as a parameter.
- Detail Preservation Radius - used with binary either/or Detail Preservation algorithms to smooth the transition. As you move away from where a processed (composite) pixel is selected, a greater proportion of the relevant non-processed (original image) pixel is combined until there is no processed pixel contribution. This parameter defines the radius of that transition. Used in the Detail Preservation algorithms 'Min Distance to 1/2 Unity' and 'Max 'Contrast".

'Color' Module

Star Colors

- Stars radiate similar to a black body - with the color made from a continuous spectrum - but with absorption lines.
- The peak output wavelength depends on temperature - appearing as red, orange, yellow, white & light blue
- They are not the colors of the spectrum - there is no green, indigo (dark blue) or violet

CLASS	TEMP K	APPARENT COLOR	NOTES
M	<3,700	Orange-Red	Old stars
K	<5,200	Yellow-Orange	
G	<6,000	Yellow-white	
F	<7,500	White	
A	<10,000	White-Blue	
B	<30,000	Blue-white	
O	>30,000	Light Blue	Young stars and star-forming regions

Gas Colors

Excited gases produce emission spectra - Narrow band

ID	ELEMENT	IONISATION	COLOR	WAVLENGTH (NM)
HII	Hydrogen	Single	Red-Pink	656.3 nm
HA+HB	Hydrogen	Single	Purple (Hb Blue)	656.3 486.1 nm
NII	Nitrogen	Single	Red	654.8,658.3 nm
SII	Sulphur	Single	Red	671.9,673.0 nm
OIII	Oxygen	Double	Cyan	495.9,500.7 nm
OII	Oxygen	Single	Near U-V	372.6,372.9 nm

'Real' colors in Astronomical Images

Checking colors are ‘real’:

- Full range of star colors - Older (red, orange, yellow) and younger (white and blue)
- Gas clouds - HII regions Red (Ha) or Purple (Ha & Hb), SII and NII Red, OIII cyan

There is an interesting discussion of the issue of 'real' colors in astrophotography in the article by Jerry Lodigruss 'Color' in Astronomical Images.

Also there is another related discussion on Starizona True 'Color' Imaging.

This discusses ways of getting a proper color balance including using a white (G2V) star.

G2V vs Other sampling techniques

The article 'Getting the colors right in your astrophotos' discusses other approaches as well as discussing the idea that a sun-centric view of color balance (Anthropocentrism) does not make sense, so methods other than using G2V stars to color balance should be used.

This PixInsight forum post 'About our color calibration methodology' describes the PixInsight approach to color balancing - which they call 'spectrum-agnostic' or 'documentary' calibration methods.

They 'try to apply a neutral criterion that pursues a very different goal: to represent a deep sky scene in an unbiased way regarding color, where no particular spectral type or color is being favored over others'.

They advocate using light sources which include a good range of stellar populations and spectral types. Examples they suggest are:

- A nearby galaxy with negligible red shift.
- A sampling of a large number of stars - by averaging a sufficiently representative sample of stars there is no bias towards one color.

Color Balancing and Light Pollution 'Filter's

- Traditional light pollution filters filter out the band of wavelengths that some artificial lights (e.g. low pressure sodium lights) produce.
- This means there is part of the spectrum missing (in the orange-yellow region) which means it is not possible to color balance using this data - see the Special Techniques section for a way of dealing with this.
- With the rise of more broadband lighting (LED, HPS and Metal Hydride) this approach is also becoming less effective.
- Here is an Interesting approach to using narrow band filters to overcome light pollution Using Photometric 'Filter's to Overcome Light Pollution.
- Instead of a single LP filter, it uses 3 narrowband filters Ha (656nm x 20nm), sYel (550nm x 19nm) and sV (410nm x 16nm) assigned to R, G and B. The 3 channels are color balanced by calibrating against a white source.

Factors affecting color balance

A large number of factors can affect the white balance (interstellar, terrestrial and equipment) and not all of these have a uniform affect across the image. This needs to be taken into account when choosing reference points.

- Red Shift (varies across image)
- Emission gases - causing enhancement in specific wavelength (varies across image)
- Interstellar extinction - reddening due to scattering off interstellar dust or other matter. (varies across image)
- - Absorbs and scatters blue more than red
- Atmospheric Extinction - scattering by atmospheric dust (varies with altitude of target)
- - Scatters blue more than red - reasonably predictable - lowest at zenith, increasing with zenith angle
- Transparency - atmospheric moisture content, pollution, ash etc. - can cause color shifts (varies across image)
- Telescope Optical path (uniform across image) - including:
- RGB 'Filter' transmission profiles
- Light Pollution 'Filter'
- Sensor sensitivity profile
- Intensity variation over time - affecting images made with RGB filters - since R, G and B are captured at different times
- Preprocessing (uniform across image)

'Color' Module Demystified

This section is to demystify what exactly StarTools is doing to your colors and how the Style and LRGB Method Emulation let you wield awesome power, virtually without having to lift a finger.

First up is the 'Style' parameter. Ever since Tracking (and by extent, non-linear processing) was introduced in 1.3, a number of completely new and innovative abilities were bestowed upon StarTools' modules.

We're all familiar with StarTools detail aware noise reduction, deconvolution after stretching, and lately the intelligent 'detail aware' wavelet sharpener. Around 1.3.2 the 'Color' module was revamped with a new way of 'Color' compositing, separating luminance and color processing completely. At the time it seemed like the logical way to implement it, having the Tracking data at our disposal, but it turned out to have some more interesting attributes.

This algorithm, which is now embodied in the 'Scientific ('Color' Constancy)' setting of the 'Style' parameter, is a new, more scientifically accurate way of representing color in your images. 'Color' perception is a difficult subject and there are no right answers to what constitutes 'the right way' to present color in deep space images. I'd love to elaborate on this, but we'll save that discussion for another time. What does exist, however is a 'worse' way of presenting color in deep space images and it is, unfortunately, surprisingly common.

Take this image of the Orion nebula for example

It's not the prettiest image, but there's nothing majorly wrong with that. Or so you'd say... Indeed, this is typical of the best that people using other software come up with in terms of color. But here is the rub:

This is the exact same image, same color calibration, just stretched less.

The core now appears green (which, incidentally, is correct; OIII emissions are dominant), you might also notice that the stars that are visible are much bluer.

So how come, in the previous image, the core was nearly white? The answer is that this is the result of (erroneously) stretching the color data along with the luminance (brightness) as well, which is strangely very much common practice, even in other software that prides itself on scientific accuracy. 'Color' is a product of discrepancies in the red, green and blue channels. Stretching these channels non-linearly deforms, stretches, squashes and skews these discrepancies based on their in-channel brightness. The result is that hues and saturation levels vary wildly as the user starts stretching the data in order to recover more luminance detail.

Of course, the notion that things out there in space magically change their color based on how we stretch our image would be a fairly ludicrous proposition. Yet, sadly, here we are, with a majority of images out there suggesting that this is the case.

Now look at the two images again. The core is clearly green in the second (less stretched) image, while the core in the first image is so pale that we might have missed it is actually green (indeed many astrophotographers throw their hands up in the air and just depict it as white). However, notice the area adjacent to the core at 3 o'clock (in the second image it is completely absent). It is decidedly more green. Could the two areas be of similar chemical make up?

In fact they are! It is just neigh impossible to see, as stretching the luminosity component has (for no good reason at all) drained the core of its color, compared to the area adjacent to it.

And here is where StarTools' Tracking aided color compositing algorithm comes in:

Now let's ignore for a minute that the colors might appear a little oversaturated. In the interest of fairness I applied the same settings for image #1 and image #3. I had to be a good bit more aggressive with the saturation to show any color in #1.

What we're seeing here is 'color constancy'. No matter how the image was stretched, the colors are absolutely comparable between areas that vary wildly in their dynamic range. The area adjacent to the core at 3 o'clock is the exact same green as the core. Also spare a moment to look at the stars. They now show color, no matter how dull or bright they are. The full visible black body emission spectrum is covered and what's most important - temperatures are completely stable and independent of how dull or bright they were recorded and subsequently stretched. This is because StarTools stays true to the color data as initially recorded, undoing all the stretching (thanks to Tracking!) to the luminance data to arrive at the 'virgin' colors as they were recorded.

This is how color in space should be presented from a scientific point of view - reproducibility no matter the exposure time or who processed it, with true separation from the way luminance was processed. Faint Andromedas should produce the same colors as bright Andromedas. Short exposure star fields should produce the same colors as long exposure starfields, etc.

Below is another striking example. Notice how star colors are recovered until well into the core of the globular cluster, while the viewer is able to objectively compare star temperatures all the way; colors are constant and comparable throughout the DSO.

Now, as of 1.3.5 there is actually a way to 'dumb down' (which perversely was a lot of work to implement and get right!) color compositing to the levels of other software. StarTools is all about enabling the user, not about pushing philosophies, so if you want to create images with the 'handicapped' traditional look you now can, using the 'Artistic' setting for the 'Style' parameter.

As a StarTools user you may have felt that your colors were never quite like 'the other guys'; more vivid, just more colorful. And you would have been right. You may have secretly liked it but you may have been uneasy being proud of the result as it is 'different' to what a lot of other folks have been producing (your Andromeda may have had a 'strange' yellow core, your M42 may have had an 'odd' green core, your M16 may have had many more tints of red and pink, or your stars seemed almost too CGI-pretty to be true). Now you now why, and now you know you can be just as proud at your images - perhaps even more, as they have been more scientifically correct/valuable to look at than the other guy's.

The 'LRGB Method Emulation' setting

Over the past two decades, astrophotographer have struggled with compositing color (e.g. combining separate luminance and color information) images, leading to a number of different LRGB compositing methods. I won't go into them here, but all attempt to invariably 'marry' luminance and color information. The different techniques all produced different hues and saturations for the same data. And what is true for all of them is that they involved a lot of layering and pixel math to produce.

Because StarTools separates color from luminance until the very last moment (that moment being when you use the 'Color' module), really any method of color compositing can be chosen for combining the data at that moment. The 'LRGB Method Emulation' simply lets you choose from a number of popular compositing techniques at the click of a button. You will notice that all these different settings have a subtle (or sometimes not-so-subtle) impact on the hues (and sometimes saturation and luminance) that are produced. It's just that easy.

Even if your data was from a non-LRGB source, it still works as StarTools will have still separated color from luminance during processing by creating a synthetic luminance frame from your RGB-only data where needed. Conversely, as you may have seen in the latest

YouTube video, processing LRGB data is now no different from processing DSLR or OSC data. Luminance is separated anyway! The only difference is that you load your data through the LRGB module. It's a big step towards a unified yet extremely powerful workflow for all data sources.

Channel mapping and matrices

A non - optimal solution in 'Pixel Math' solution would be to feed a Ha multiplier greater than 1.0 was into the R channel, with the OIII subtracted from it. From signal processing point of view this would open up clipping scenarios (e.g. what if a Ha multiplier > 1.0 was fed into R and there was no O-III in that area to pull it down from above 1.0? Unless some sort of normalization is applied afterwards…), this would result in signal clipping. Plus it makes things rather intuitive for someone trying to figure out what is going on in an area. Hue creation from tristimulus values doesn't really call for (or work with) subtraction, unless perhaps some sort of out-of-gamut color space is being targeted (which doesn't really make sense in this context IMO).

In ST, it is really (almost) as simple as following the assignment laid out in the Matrix parameter options.

Let's take, for example "SHO 40SII+60Ha,70Ha+30OIII,100OIII"

What's happening internally is:

1. The assumption is made that your data was imported as SHO:RGB in the compose module (e.g. as it says on the butttons at the top in that module).
2. Therefore, ST assumes that S-II currently purely lives in the red channel, Ha currently purely lives in the green channel and O-III currently purely lives in the blue channel.
3. The Color module then scales the R, G and B channels (aka the S-II, Ha and O-III signal respectively) by the multipliers defined by Red/Green/Blue Bias Reduce/ Increase. In other words, by modifying those Red/Green/Blue Bias controls, you are directly throttling the pure S-II, Ha and O-III signals before any remapping is done.
4. Note that the S-II, Ha and O-III signals are normalized after the multiplication so that no channel will ever exceed 1.0 (or in other words, a massive pull on the green slider will result in the other channels really just "shrink" by the weighted inverse). For example RGB 0.6:12:3 would become 0.05:1.0:0.25 after normalization. This is now your input signal for the remapping.
5. Finally, the remapping is performed. In the case of our example "SHO 40SII+60Ha,70Ha+30OIII,100OIII", the remapping becomes:
    - *Rnew = 0.4*Rbiased + 0.6*Gbiased*
    - *Gnew = 0.7*Gbiased + 0.3*Bbiased*
    - *Bnew = 1.0*Bbiased*

In the case of StarTools that is not entirely the end of the story, however, as, only the resultant coloring is used. The brightness component is unceremoniously thrown away, while the separate brightness/luminance from the (parallel-processed) detail/luminance dataset is adopted instead. How exactly the luminance is integrated with the coloring is dependent on the chose LRGB Method Emulation (see docs or in-app help item).

The remapping, however, is quite transparent and you can, for example, calculate from the mapping equation what hue a pure S-II, Ha or O-III signal will have. For example, a pure Ha signal (1.0) in our example would yield:

- R = 0.4 * 0 + 0.6 * 1.0 = 0.6
- G = 0.7 * 1.0 + 0.3 * 0 = 0.7
- B = 0

You can plug that into, say an online RGB to HSL converter (multiply by 255 as this one expects an 8-bit value for each R, G and B value). For the resultant RGB 153, 179, 0, you will get a Hue of 69 degrees, which looks like this. In other applications (e.g. PixInsight), many (not all) people would now freak out because that's green! (ew!) They would now likely resort to tweaking the equation (in essence picking another hue for Ha), or resorting to SCNR/green killing tools (selectively modifying pixels above a threshold of green), rather than do the logical/"correct" thing and tweaking the input signal. In StarTools this is as easy as chaining the Green Bias Reduce/Increase to a point where that yucky green is overwhelmed by the other S-II and O-III signal.

So you get something nicely balanced like this even though Ha is (of course) dominant:

That smidgen of green dominance at 9 o'clock is enough to let the viewer know that that area is indeed strongly Ha dominant compared to other areas.

Crucially, the input signal is just attenuated (or boosted) in relation to the other three signals, but is not harmed in any other way that is either destructive, unpredictable or non-replicable across other datasets/objects/gear. The holy grail of color rendering for documentary photographical purposes is:

to be able to get comparable coloring across different objects with comparable emissions and to be able to replicate comparable coloring by different people with different gear, however with comparable filters

Narrowband imaging is really no exception. Sure, you have leeway in the hues, but the way these hues behave and the story they tell about the same object should be consistent.

To make a long story short, it's not so much about the pixel math / compositing equation - it just establishes the hue (which is best kept as simple and predictable as possible to aid the viewer seeing the three different emissions). The important bit is the throttling of the input signal to the equation to get draw the viewer's attention to a specific feature/ emission via color.

'Denoise' Module

StarTools denoise techniques

By doing denoise late in the workflow - Tracking has had time to follow noise evolution over most of process and identify areas prone to noise, allowing noise reduction to target these areas.

Identifying Detail in an Image

- The traditional method of identifying detail in an image is to use a mask - either based on luminance or created manually. To avoid this Startools identifies detail automatically by using a technique called Scale Correlation.
- Usually, when there is a correlation between image elements over multiple scales it indicates important detail in an image. This is how the 'Denoise' module identifies detail. It can then provide the control to protect this detail from the denoising algorithm.
- By making the 'Denoise' module scale aware it allows the comparison of elements at different scales. Looking for correlation between elements at different scales enables identification of likely detail.
- The number of scale levels which the algorithm tries to correlate dictates the smallest detail that is identified - and therefore may be protected. If the scale is too small it is possible, under certain conditions, that you start to protect noise that is mistaken for detail. That is why we control the depth of the search for detail in the scale levels.

Wavelet 'Denoise'

- Wavelet scale extraction - classifies features and structures into 5 different size scales.
- Noise removal is done by an enhanced wavelet denoiser - removes features (such as noise) based on their size.
- Noise grain caused by shot noise exists at all scale levels - becoming less noticeable as size increases.
- 'Denoise' aggressiveness at each scale is adjustable using the Scale parameter.
- Global noise reduction (i.e. not scale-specific) is done by the Brightness/Color detail loss setting.

Scale Correlation

StarTools looks for inter-scale pattern/structure correlation to identify image detail.

- Correlation is higher in areas that look 'busy' - this is normally associated with image detail.
- Correlation is low in areas that have little change such as large, smooth, gas clouds.
- Scale Correlation removes the need for a mask to protect image detail from noise reduction.
- Where noise does not exhibit a Poisson distribution it may exhibit scale correlation - which can cause noise to be mistaken for detail.
- To avoid this, reduce the depth of correlation using the Scale Correlation parameter.

Noise Sources

Shot Noise

- Caused by the random arrival of photons.
- Proportional to the square root of the intensity of light falling on the pixel.
- Independent of other pixels.
- Poisson distribution.
- Reduced by stacking multiple sub-frames.

Read Noise

- Caused by random variations in the current in the equipment electronics.
- Mainly thermal noise - temperature dependent.
- Independent of the amount of light falling on the pixels.
- Gaussian distribution.
- Dominant at low intensities.
- Reduced by stacking multiple sub-frames.
- Noise reduction of (Gaussian) Read Noise is done differently from the noise reduction of the (Poisson) Shot noise.

Dark Current Noise

- Caused by the dark current - which increases linearly with time and exponentially with temperature.
- Poisson distribution.
- Independent of the amount of light falling on the pixels.
- Reduced by cooling of the sensor.

Quantization Noise

- Derived from quantization error in A-D converter.
- Depends on the number of bits.
- Can be intensity dependent.
- Small for A-D converters of 12 bits or more.

Visible Noise Characteristics

Salt and Pepper Noise

- Descriptive of noise where there are bright pixels in dark regions (salt) and dark pixels in bright regions (pepper).
- Random errors of large variation.
- Caused by bit errors, A-D errors, electronic interference.
- Normally removed by use of dark frames or median filtering.

Fixed Pattern Noise

- Descriptive of noise which is distributed in a fixed pattern.
- e.g. Row or column patterns.
- May be caused by small differences in the characteristics of pixels.
- May be due to debayering issues in color cameras.
- Normally removed by using bias frames, or dithering.

Evaluating Results

When working on an image over a period of time it is difficult to ensure the changes you make are moving to a 'better' overall image.

- It is important to try and see it as if through the eyes of someone seeing the image for the first time.
- This is particularly difficult when trying to get the optimum noise reduction.
- One useful technique is to look at the image at different levels of zoom to see what effect your changes have at different scales.

Walking Noise

Walking noise (which is indeed correlated as its presence correlates with neighboring pixels) happens when a pixel is a bit "hotter" or "colder" than its neighboring pixels. E.g. it produces a little higher or lower brightness when given the same signal. In a single frame this shows up as a single pixel, but as the field of view slightly shifts around due to small tracking error as the earth rotates, that pixel will sit in a slightly different spot each time (usually slowly moving away as the field of view drifts due to tracking error). When the stacker stacks it all together this will cause a small trail.

Even modern camera's sensors might have visible pixel sensitivity deviation in a single frame. Stacking however, produces a dataset of much more bit depth. Because random noise from other sources (which would at least partially mask the sensitivity deviations in a single exposure) will be reduced by the stacking process, the pixel sensitivity deviations may become obvious using an aggressive stretch.

Dithering introduces deliberate big shifts in the field of view, so that any hot or cold pixel will move a lot and will never sit at the same position. This breaks up the trail and, more importantly gives the stacker a much easier job detecting the hot or cold pixel as an outlier. The result is a vastly cleaner dataset without walking noise trails.

'Heal' Module

Removing Star Trails in Comet Images

If you want to replace the star trails in an image of a comet with point stars you can use the 'Heal' module.

1. Take a sequence of comet images for stacking.
2. Stack the comet images - using the comet as the alignment point. Use sigma clipping removes some if not all the stars. This gives the clear image of the comet with blurred (or reduced) stars.
3. Re-stack using the stars as the alignment points. Use sigma clipping to reduce the comet. This gives a clear image of the stars with a fainter blurred image of the comet.
4. Create a completely starless image using the clear comet image - as described in 'Creating a Starless Image' above.
    - 'Keep' the comet-only image when you are happy with it.
5. Create a stars-only image:
    - Launch the 'Layer' module.
    - Click the preset 'Undo->Bg' - this sets the background (left panel) to what is in the undo buffer - which is the original image.
    - Set the 'Layer Mode' to 'Subtract' - and the result (right panel) is the extracted stars.
    - 'Keep' the resulting stars-only image.
    - Save the stars-only image - we will need it later.
6. Process the comet image:
    - Load the comet image.
    - Process the starless comet image as you want (e.g. 'Sharp', 'SuperStructure', 'Denoise', 'HDR' etc.) until you are happy with it.
7. Merge the stars back in:
    - Launch the 'Layer' module - This loads the comet image in the foreground (centre panel).
    - Click 'Open' and select the saved stars-only image - This will load into the foreground (centre panel) and the comet image becomes the background (left panel).

    - Set the 'Layer Mode' to 'Add'.
    - The resulting merged image is in the right panel.
    - Click 'Keep' to keep the merged image.

You should end up with an image where the comet is clear and in focus and so are the background stars.
A variation of this is to take an image of the stars with a shorter duration and use this as the basis for the stars-only image. This will reduce the brightness of the background stars and so highlight the comet.

'Repair' Module

The 'Repair' module allows for the correction of more complex aberrations than the much less sophisticated 'offset filter & darken layer' method, whilst retaining the star's exact appearance and color.

The repair module comes with two different main algorithms:

- The 'Warp' algorithm uses all pixels that make up a star and warps them into a circular shape. This algorithm is very effective on stars that are oval or otherwise have a convex shape.
- The 'Redistribution' algorithm uses all pixels that make up a star and redistributes them in such a way that the original star is reconstructed. This algorithm is very effective on stars that are concave or otherwise heavily distorted, such as stars that look more like patches of light without a distinct core which cannot be repaired using the 'Warp' algorithm.

'Layer' Module

PixInSight nearest equivalent: - Gradient'HDR'Composition, Morphological 'Filter's, PixelMath

Aligning two or more Images

In cases where multiple images are combined it is important that the images are exactly the same size. If they are not you will get the error message: 'Dimension differ from already loaded file'. If you do, check the dimensions of individual data files by loading them in StarTools individually. The dimensions are listed to the right of the file name at the top of the screen.
Also, in most cases it is important to ensure all the images are aligned the same.
To ensure multiple stacked images are aligned you can use one of the sub-frames as a reference frame.
This is described for Deep Sky Stacker (DSS) and Regim in the Background Notes in LRGB Module Use notes.
When processing the images in StarTools before combining, make sure you 'Bin' and 'Crop' each image the same amount.

Pixel Maths

In the 'Layer' module the resulting pixel value depends on three functions:

- 'Filter'(mode,blend amount, Fg,Bg) - calculated first.
- 'Layer' Combination(mode, kernel radius, Fg,Bg) - calculated using the result of the filter.
- Brightness mask(mode, power, Fg,Bg) - overall mask applied to result.
- This result can be used as the input and processed again.

'Filter' Types

See the 'Filter Type' parameter for a large selection of different filters.

The 'Filter' Types in StarTools are similar to some of the 'Filter' Types in Photoshop and GIMP or the Filters in PixInsight. They fall into three broad groups:

- Noise reduction - Gaussian, Median, Mean of Medium Half, Differential Adaptive Noise, Mean of Medium Half Distance Weighted.
- Bring out detail - Minimum, Maximum, Lightness, Local Histogram Equalize, Local Maximum 'Entropy' RGB Selection, Sobel, Fractional Differentiation.
- Blending Foreground and Background - Min Distance to 1/2 Unity, Max 'Contrast' - these can be used to create High Dynamic Range composites from 2 images with different exposure lengths.

'Layer' Modes

The 'Layer' Modes in StarTools specify how the background and foreground are combined. They are similar in effect to some of the GIMP 'Layer' Modes or the Photoshop Blending Modes or PixInsight Blending Modes. See also the Wikipedia article on Blend Modes.

'Layer' Modes can be split into the following general groups:

- Lightening - Lighten, Screen, Add, Power of Inverse
- Darkening - Darken, Multiply, Multiply Luminance
- Making Dark Darker, Light Lighter - Overlay, SoftLight, HardLight
- Increasing Saturation - 'Color' Extract fg
- Desaturation (Black and White) - Desaturate fg (Average), Desaturate fg (Luminance)
- Combining L and RGB - Brightness of fg, 'Color' Of fg
- More Extreme Darkening - Subtract, Difference, Divide

'Stereo 3D' Module

The benefits of 3D representations

One of the benefits of this module is to encourage new ways of looking at a subject. It is possible to make assumptions about the target which may be wrong - such as:

- Dark objects are in the background
- Light objects are in the foreground

This are assumptions we make from our experience of our own world - but this is based on reflection of a bright light source.

However with DSO's the following may be true:

- Sources of reflected light are less commonly visible than emission sources.
- Dark areas may be in the foreground - blocking the light coming from elements behind.
- Bright DSO cores are often embedded in, surrounded by, and shrouded by, bright gases.

Understanding how the object looks in 3D may help you make better decisions when processing:

- Highlighting local detail at the expense of the shroud of gas surrounding it - giving a false impression.
- The way you process dark areas may make a DSO with a dust cloud in front of it look like a DSO with a hole in it.

3D viewing without aids - Free Viewing

There is a useful web site with information about 3D images here: Basic 3D Viewing Terms

You need to train your brain to see using Freer viewing methods - see here for a general discussion: The logical approach to seeing 3D pictures

The main ways of seeing 3D images without aids are:

- Cross-viewing or The Cross-eyed Method - This link describes how to see images using Cross-viewing: The Single Finger Method
- Parallel-viewing or The Parallel Method - This link describes how to see images using Parallel-viewing: Parallel-viewing Stereo Photographs

Ways of identifying depth

The 'Stereo 3D' module gives us a number of ways of identifying depth based on common rules and assumptions about gas and radiation behavior:

- Simple L to Depth - Assumes the lower the luminance the further away it is.
- Luma to Volume - Shadow dominant e.g. Pillars of Creation , Highlight dominant e.g. M20.
- Highlight embedding - bright stars with emissions gas around them - embed.
- Structure embedding - small structures embedded in larger ones.

Types of 3D images

- Anaglyph - Creates a 3D image by 'Layer'ing of different colors to create a stereoscopic image. Requires color filter glasses to view.
- Lenticular - Places thin lenses over the image to restrict the view of each eye to a particular part of the image.

Depth Map

The Depth Map tells how far away from the viewer each pixel is.

- Can be used to convert a 2D image to a 3D image.
- Facebook and Photoshop can use a depth maps to create a 3D image.

Appendix

Recommended ISO for DSLR cameras

There is an optimal ISO value for each DSLR, where your specific sensor provides the optimal balance between read noise and dynamic range.

ISO in the digital domain is unfortunately much misunderstood. The most important thing to understand is that picking an ISO value does not - in any way - make your digital camera's sensor more or less sensitive to light. A sensor's ability to convert incoming photons into electrons is fixed. The article 'What is the best ISO for your DSLR for astrophotography?' by Chris van den Berge goes in more depth.

For the purpose of astrophotography then, your camera will have an ISO value that is optimal for this type of photography. This section contains a number of suggested ISO values for popular DSLR models from popular vendors. These values are based on data from Photons to photos, sensorgen.info (now defunct), DxOMark and dslr-astrophotography.com.

Please note that these are suggestions and you may wish to do more research and/or try one above the suggested setting.

Canon

MODEL	ISO
Canon 100D / Rebel SL1	400
Canon 1000D / Rebel XS	200
Canon 1100D / Rebel T3	1600
Canon 1200D / Rebel T5	800
Canon 1DX	3200
Canon 1D Mark II	400
Canon 1D Mark II-N	400
Canon 1D Mark III	800
Canon 1D Mark IV	1600
Canon 1Ds Mark II	400
Canon 1Ds Mark III	800
Canon 20D	1600
Canon 30D	800
Canon 350D / Rebel XT	400
Canon 400D / Rebel XTi	200
Canon 40D	800
Canon 450D / Rebel XSi	400
Canon 500D / Rebel T1i	1600
Canon 550D / Rebel T2i	800
Canon 50D	400
Canon 5D	800
Canon 5D Mark II	1600
Canon 5D Mark III	3200
Canon 600D / Rebel T3i	800
Canon 60D / Canon 60Da	800
Canon 650D / Rebel T4i	800
Canon 6D	1600
Canon 700D / Rebel T5i	800
Canon 70D	1600
Canon 7D	1600
Canon 7D Mark II	1600
Canon 80D	200

Nikon

MODEL	ISO
D3100	1600
D3200	100
D3300	400
D5000	400
D5100	200
D5200	800
D5300	200
D5500	100
D600	200
D610	400
D700	800
D7000	100
D7100	200
D750	100
D800	200
D810a	200

Sony

MODEL	ISO
Sony A3000	800
Sony A5000	200
Sony A5100	400
Sony A6000	400
Sony A7	800
Sony A7 II	800
Sony A7 III	1600
Sony A7R	200
Sony A7R II	800
Sony A7R III	1600
Sony A7R IV	800
Sony A7S	3200

Deep Sky Stacker Settings

Deep Sky Stacker Settings

Deep Sky Stacker (FREE) remains a one of the most popular pre-processing applications for Windows. Stacking and saving your data with these settings is essential to getting good results from StarTools.

When applying the important pre-processing do's and don'ts when using StarTools with any stacker, you will want to configure Deep Sky Stacker specifically in the following manner.

- Choose No White Balance Processing in the RAW/FIST
- Choose Bilinear Interpolation for the Bayer Matrix Transformation algorithm
- Save your final stack as 32-bit/channel integer files, with adjustments not applied.
- Stack with Intersection mode - this reduces (but may not completely eliminate) stacking artifacts
- Do not choose Drizzling
- Turn off any sort of Background Calibration

With all the above settings made, you can then safely stack and (assuming you used a DSLR or OSC) import your dataset into StarTools as 'Linear, from OSC/DSLR with Bayer matrix and not white balanced'.

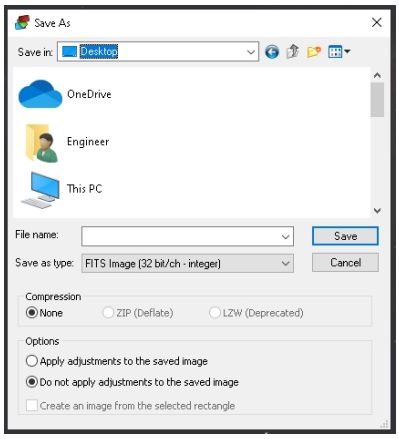

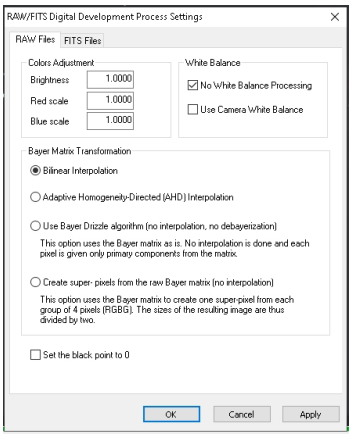

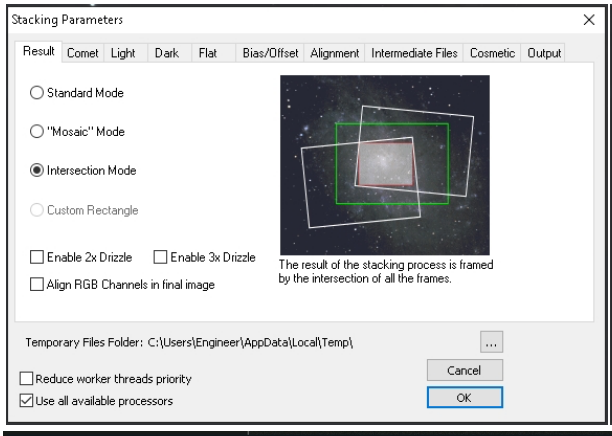

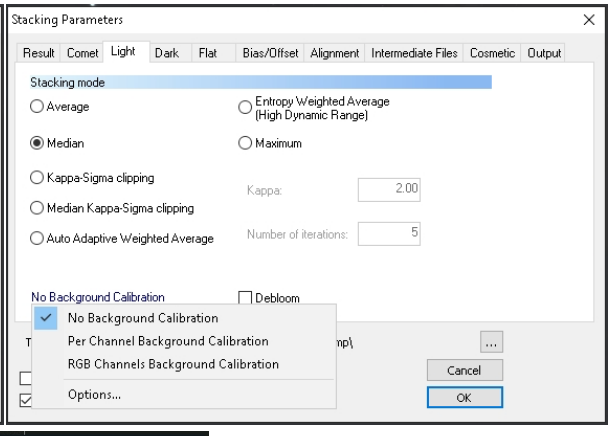

AstroPixelProcessor Settings

Astro Pixel Processor ("APP") is a paid stacking solution for Intel-based Windows, macOS and Linux OS's. Stacking and saving your data with these settings is essential to getting good results from StarTools.

In addition to the important pre-processing do's and don'ts, these are the settings in the APP tabs that need to be used to optimize datasets for StarTools:

0) / RAW/FITS: This tab is only relevant for DSLR / OSC camera users. Use the default settings for DSLR/OSC

- provide the Bayer pattern of your camera manually if it's not supported
- 'Adaptive Airy Disk' seems to produce a slightly better result over 'Bilinear'.
- leave Camera White Balance unchecked

1) / LOAD: use the default settings

2) / CALIBRATE: use the default settings except:

- Disable 'adaptive pedestal / reduce Amp glow' if possible. Amp Glow should be removed by DarkFrames
- Do not use 'remove light pollution' (leave unchecked)

3) / ANALYSE STARS: use default settings

4) / REGISTER: use default settings

5) / NORMALIZE: default settings except:

- Disable "neutralize background". This will perform normalization for proper outlier rejection but skip the background calibration which should be performed in StarTools by WIPE (under Tracking)
- As of APP 1.083 it will be possible to disable normalization completely but this is not recommended because this would impact outlier rejection.

6) INTEGRATE: default settings.

- APP is using an Auto mode and will choose integration mode based on the number of lightframes loaded. You may also chose settings manually - please refer to APP hover-over instructions for guidance.
- Refrain from using Local Normalization Correction and Multi Band Blending unless absolutely needed (if you need to stack images from multiple nights for example).
- do not use drizzle modes

The integration dataset(s) will be saved as a 32-bit FITS file(s) named "(Name)-(Channel)-session-1.fits". These may be loaded into StarTools.

9) TOOLS: Do not use this tab - these tools would introduce severe unlinearity and ST Tracking will have an issue with these. Use the ST 'Wipe' tool to perform background color calibration and light pollution removal instead.

If stacking multiple mono datasets for use in a composite, make sure to use one set's finished stack as a reference to stack the others with; StarTools's Compose module requires every dataset to be the same dimensions. Aligning remaining channels against an initial channel during stacking is particularly important to ensure consistency of point spread functions across channels; do not align finished stacks against each other after stacking.

Finally, do not pre-composite in APP, but use the Compose module in StarTools instead.

ASTAP settings

These setting should yield a suitable stack under most circumstances.

ASTAP ("Astrometric STAcking Program") is a FREE, competent, actively developed stacker, available for all platforms. Stacking and saving your data with these settings is essential to getting good results from StarTools.

Of particular importance is switching off "Auto levels" in the Stack method tab. Any sort of color calibration should be avoided. Once turned off, you will be able to import your ASTAP-exported dataset into StarTools with the second option.

Besides calibration and stacking, do not perform any further operations on the resulting dataset in ASTAP.

If stacking multiple mono datasets for use in a composite, make sure to use one set's finished stack as a reference to stack the others with; StarTools's Compose module requires every dataset to be the same dimensions. Aligning remaining channels against an initial channel during stacking is particularly important to ensure consistency of point spread functions across channels; do not align finished stacks against each other after stacking.

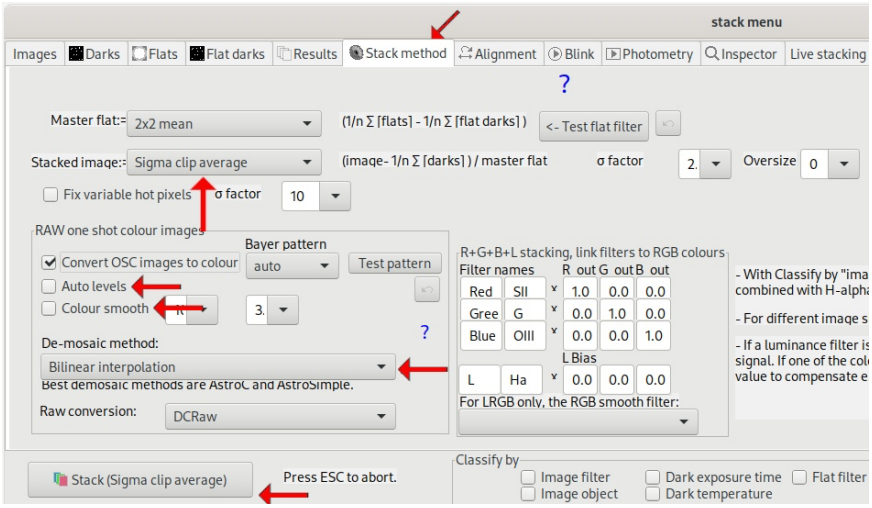

PixInsight User Information

If you're coming from post-processing in PixInsight, then there are a number of important things that are different/new in StarTools:

Signal evolution 'Tracking' plays an extremely important role in StarTools. This makes workflows much less linear and allows for StarTools' engine to 'time travel' between different versions of the data as needed, so that it can insert modifications or consult the data in different points in time as needed ('change the past for a new present and future'). It's the primary reason why there is no difference between linear and non-linear data in StarTools, and you can do things in StarTools that are otherwise quite destructive (like deconvolution after stretching your data) in PixInsight. If you're not familiar with Tracking and what it means for your images, signal fidelity and simplification of the workflow & UI, please do read the above manual section.

As of version 1.5, StarTools is able to process luminance - whether synthetic, real or a combination thereof - and chrominance (color) data in parallel. Apart from important enhancements in signal fidelity, this means that processing luminance and chrominance data separately (and combining afterwards) is no longer required, nor considered best practice. (video examples see ST homepage)

The time-shifting freedom Tracking affords, makes some operations best applied at entirely different times than you might be used to:

Noise reduction is always applied at the very end of your processing when switching Tracking off.

'Color' calibration can be applied at any time, but is preferably applied at the end of your processing, as StarTools is able to completely negate the adverse effects of stretching (global and local), yielding perfect color constancy (e.g. no hacks or scripts such as 'Repaired HSV Separation' are necessary).

As of the 1.6 alpha versions, extensive support for narrowband coloring schemes, means you can completely change the color presentation of your narrowband composite with a single click at any time during your processing, without having to re-composite.

Deconvolution is preferably applied after stretching (and possibly even after other local detail enhancement). This is because 'Decon' takes into account how signal has evolved in the stretched data when performing regularization of the linear equivalent.

Traditional histogram stretching and white/black point setting is completely abstracted away. StarTools considers such tools sub-optimal and archaic approaches to image processing.

Many tools seems simpler but, in reality, are more powerful. Data mining and analysis does away with parameters that can be objectively determined, or (derivatives of) parameters that you already specified in other modules earlier. As such, truly destructive settings are neigh-impossible to achieve (for example StarTools refuses to clip your data) and objectively optimal constraints are automatically determined from prior input. Luminance masks or local supports are automatically generated and updated as needed without user intervention, given the program keeps the necessary statistics and image versions/instances available at all times (Tracking).

StarTools does not treat processing as a linear, object-oriented process (e.g. an implied collection of independent processes), because real image processing carried out by a human emphatically isn't a linear process. In StarTools, as long as Tracking is on, processing steps depend on each other and reuse each other's data, findings and previous user settings. StarTools is all about refinement and shaping of data, where you revisit, rather than reapply modules. A linear, object oriented approach is therefore not possible or even desirable if signal fidelity, noise mitigation and ease of use are paramount. As such you will not find the equivalent of 'process containers', undo histories or similar in StarTools. Instead you'll find a single undo buffer and a 'Restore' button, which will let you negate the contributions that certain operations made to the currently visible end result, regardless of sequence. Processing sessions in StarTools are very short as a result; it is a converging process and 'overcooking' an image is not a thing.

While some prefer scientific purity of an image (the author included), others like to emphasize aesthetics. StarTools does not try to prescribe what's right and what's wrong in this regard. As such StarTools contains tools like the 'Heal' Module, the 'Synth' module, the 'SuperStructure' module (and others) that can be (ab)used to help the user make educated guesses about what plausible data would look like, based on what is already in the image and based on what the user knows is definitely not in the image. Contrary to PixInsight, we take no stance in what should or should not

be allowed by including or excluding these tools in the software. Some people like diffraction spikes on their stars, some people don't like stars at all. Some people like to be able to remove dust donuts or fix CCD blooming issues. You decide.

PixInsight	StarTools	Notes
ABE	'Wipe'	Using 'Wipe' without a mask functions similar to ABE - a gradient/bias model is created based on all pixels in the image.
DBE	'Wipe'	Using 'Wipe' with a mask functions roughly similar to DBE. If a large area of the image is covered by a DSO or nebulosity, it is recommended to mask out this DSO or nebulosity, as 'Wipe' may include aspects of the DSO or nebulosity in the background/bias model.
BackgroundNeutralization	'Wipe'	At the same time 'Wipe' flattens the field and removes gradients, 'Wipe' also calibrates the background and removes bias in the color channels. Note that this does not necessarily mean that the background becomes a nice neutral grey. Rather it makes sure that the background is luminance bias-free across all color channels. Subsequent color calibration will modify and neutralize the chrominance (color) information.
Assisted'Color'Calibration	'Color'	While in PixInsight it is mandatory that color calibration is performed on the linear data at the start, color calibration in StarTools may be performed at any time, preferably closer to the end of your workflow. This is to achieve pure, more scientifically correct color constancy (a more 'compromised' PixInsight style of coloring is also selectable if you're feeling nostalgic). White point references may be selected with a mask or through clicking on the image.
'Color'Calibration	'Color'	see above
'Repair'ed HSV Separation	'Color'	see above, (StarTools retains color information, rather than trying to recover it from neighboring pixels)
AutoHistogram	Develop	Develop automatically sets the white and black point.
AdaptiveStretch, MaskedStretch	'AutoDev'	StarTools employs an enhanced (and much faster) algorithm that yields optimal global dynamic range assignment (without masking artifacts), while being easier to control. Noise may be rejected for optimization and a Region Of Interest may also be specified.
HistogramTransformation	obsolete	In StarTools histogram transformations for global dynamic range assignment are considered obsolete and sub-optimal tools. Use 'AutoDev' and Develop instead and optimize local dynamic range subsequently with 'Contrast', 'HDR' and 'Sharp'.
ScreenTransferFunction	obsolete	In StarTools the distinction between linear and non-linear data does not exist. Use Develop or 'AutoDev' to redo your stretch as many times as you like.
LocalHistogramEqualization	'HDR', 'Contrast', 'Sharp'	LocalHistogramEqualization results can be approximated with the more flexible 'HDR', 'Contrast' and Wavelet 'Sharpen' modules.
ATrousWaveletTransform	'Sharp'	In StarTools 'Sharp' does not clip, allows control over interscale-fighting and is noise-aware.
'HDR'WaveletTransform	'Contrast', 'HDR'	In StarTools, 'Contrast' nor 'HDR' clip the image and work locally instead of globally, leading to more optimized results.

Gradient'HDR'Compression	'HDR'	
Gradient'HDR'Composition	'Layer'	The Min Distance to 1/2 Unity and Max 'Contrast' filters in the 'Layer' module will create High Dynamic Range composites from 2 images with different exposure lengths.
RangeSelection	Mask (Auto)	In StarTools the Mask editor's 'auto' feature allows for a great variety of automated ways of selecting pixels, including range selection and multi-step mask addition/ overlaying.
ChannelExtraction	Compose	The Compose module performs channel extraction and even interpolation and re-weighting into new luminance masters.
ChannelCombination, LRGBCombination	Compose, 'Color'	In StarTools the Compose module stays active in the background, allowing StarTools to process luminance and chrominance data separately, yet in parallel. The 'Color' module as of 1.6 allows for complete change of color presentation at any time, without having to re-composite
Morphological 'Filter's	'Layer'	See 'filter type' parameter for a large selection of different filters.
PixelMath	'Layer'	In StarTools, pixel math is performed in the 'Layer' module by means of graphical user interface. A huge range of operations is selectable, while multiple operations can be chained by means of a buffer.
Deconvolution	'Decon'	In StarTools, deconvolution is noise-aware and is able to generate its own de-ringing mask. De-ringing will still try to coalesce singularities. Deconvolution is StarTools is geared towards dealing with atmospheric and diffraction-limited data only.
ACDNR, TGV'Denoise', MURE'Denoise'	'Denoise', 'Denoise' 2	In StarTools, noise reduction is applied at the very end when Tracking is switched off. Due to StarTools' noise evolution Tracking noise reduction will much more targeted than any particular 2D noise reduction routine or combination thereof. No (sub-optimal) luminance masks or local supports are required.
Star De-emphasis (Adam Block)	'SuperStructure' Isolate preset	The 'SuperStructure' module's Isolate preset Isolates large scale structures and pushes back busy star fields.

OS and Hardware Requirements

Operating systems

StarTools works on all 32-bit (NT-based) and 64-bit versions of Windows. This means that StarTools runs on Windows NT 4, Windows 2000, Windows XP, Windows Vista, Windows 7, Windows 8 and Windows 10.

StarTools works on macOS versions from 10.7 onwards.

StarTools should works on 32-bit or 64-bit Linux distributions with X11, GLIBC 2.15 and Zenity.

What are the minimum recommended specifications to run StarTools?

The minimum specifications to run StarTools depends mostly on the resolution the data you intend to process.

Low resolution data sets (for example from a 1MP CCD or Webcam) may be processed successfully on a Pentium 4 with 512Mb RAM.

High-resolution data sets, such as those from a DSLR typically require at least 4Gb of RAM.

8GB should usually be enough to process images ~12MP in size comfortably.

For best results, 16GB and a modern 4-core CPU are recommended, in addition to running from a RAM disk (or alternatively a Solid State Drive).

As of version 1.7 StarTools is fully GPU accelerated. Heavy arithmetic is offloaded to any OpenCL 1.1 compliant GPU present in your system. Significant processing speed can be seen on even modest, older GPUs.

Regardless of your machine's specification, consider binning your data if your data is oversampled.

StarTools is an unapologetic memory and resource hog. When it comes to low-signal astrophotography, 'regular' old algorithms are not sufficient and more rigorous, brute force approaches are warranted to make the most of your prized signal. StarTools implements these more advanced algorithms in spades, augmented at every step by its pervasive signal evolution Tracking (which keeps tabs on per-pixel signal vs noise evolution). As such, all components of your system will be given a workout and all CPU cores and threads are loaded up where possible. The features and flexibility you get in return though, is absolutely worth the CPU, memory and storage resources.

The StarTools 1.6 development versions implement a great number of optimizations that may alleviate some pressure on some systems (in addition to new features and other improvements). You may wish to give these a spin if you haven't done so.

Thank you also for sharing your specs; it's very helpful.

8GB is a the lower end of what is needed to run StarTools, but should usually be enough to process images ~12MP in size comfortably.

Intel's older U series of CPUs (i3, i5 and even i7) are only 2-core, 4-thread parts with very low clock speeds. They are great for keeping power consumption down and yield excellent battery life. Unfortunately they accomplish this by being low on raw processing grunt. It's not the end of the world - it just means longer processing times. Indeed binning will dramatically (exponentially) cut processing times, as does judicious use of previews. Be sure to try the 1.6 versions. The i5 4200U is one of the first CPUs to support AVX2, so try the new (in 1.6) AVX2-specific executable for a small speed boost as well. Like the i5, the i7-6500U is also a 2-core, 4-thread part and (according to benchmarks) should yield a modest 20% performance improvement over the older i5.

Laptops with 4-core Intel QM-series processors or the more recent 4-core AMD Ryzen U processors are a lot more powerful. They yield performance that is roughly twice as fast their older siblings.

An SSD vs mechanical drive will make less of a difference while within a module. Provided no virtual RAM is used (e.g. you have enough physical RAM), most disk access happens in between modules (after 'Keep'-ing an image). This disk access is caused by the Tracking data mining, which back and forward propagates the changes through time. With larger datasets this means routinely shifting around many gigabytes of data. In 1.6 this disk access is somewhat reduced and compression is applied.

Low-cost GPU solution (current as of late 2020)

With 1.7 now supporting comprehensive GPU acceleration, I've been testing and evaluating many different discrete and integrated GPUs in the process.

Many of you already have a capable discrete GPU in your system. However if you rely on your CPU's iGPU, or your GPU is quite old (e.g. released before ~2013, particularly from NVidia), or your GPU is a budget model (e.g. just a basic display adapter.), then you may wish to read on.

The crypto-mining boom saw the release of a number of cut-down GPUs, lacking display outputs. These were designed for maximum compute power at minimum power draw, in order to maximize crypto currency returns vs power bills.

With the mining craze having subsided, as well as with more capable (profitable) GPUs being released over time, these older GPUs are now effectively e-waste; they don't work as display adapters and, today, consume too much power vs crypto currency returns. However, these cards still provide very good compute performance at very low power draws.

This means that these cards are:

- Abundant on sites like eBay (typically buying multiple will yield discounts too)
- Cheap
- Relatively powerful
- Fit in most systems without the need for a power supply upgrade

As said, these cards complement your current display adapter by taking over compute duties. Your current display solution will continue working as normal. These cards will not be suitable for things like gaming or stream encoding, but will only be suitable for accelerating compute workloads.

The card in question is a P106-090:

You will need an empty, full-size, PCie slot on your motherboard (e.g. this empty slot should look the same as the one your current display adapter - if any - resides in);

The card will require a 6-pin PCIe power connector. Most power supplies provide one, but if not (or you current display adapter has them all in use), then cheap 12V molex or SATA adapter cables exist to take power from that. My P106-090 came with 12 molex adapter already provided.

Prices

Pricing as of the time of this writing is ~$50 USD (e.g. roughly the price of a StarTools license). I would probably not pay significantly more than that, unless other options in your area/country are limited. There are other models than the P106-090 that have better specs, however their price/performance is IMHO too high to be worth it. In most cases you are better off buying a "regular" used GPU, which can do display duties as well.

Risks

These cards have been used before. If you decide on procuring such an ex-mining card, you will be doing so at your own "risk" (I would estimate it reasonably a low risk, however, and my experience buying one through AliExpress has been good).

Performance

This card is essentially a cut-down Nvidia GTX 1060 3GB, though seeing compute performance more in the range of a GTX 1050 Ti. For a rough idea of where your current GPU solution ranks, you can have a look here. Please note that this a *very* rough estimate.

Mac OS Hints

MacOS crashes

There are two types of crashes:

- An Application crash produces an error warning including a button providing a crash report of the terminal. With such report, debugging StarTools is possible
- A system crash usually happens instantaneously with no apparent reason, usually while operating the GUI. It may not be easily reproducable. Also it does neither produce error alerts nor crash reports. Such crashes usually indicate an OS issue.

One major source of such crashes is an incorrect default command line shell (terminal). While introducing Catalina, Apple switched the standard terminal shell from bash to zsh. It is highly recommended to change the default terminal shell to zsh when upgrading to Catalina, Big Sur or Monterrey MacOS versions to avoid such system crashes using a procedure described here: https://support.apple.com/de-de/HT208050

Why does the StarTools crash under heavy load on my older macOS device?

If StarTools appears to be unstable on your older (up to 2014) macOS device, particularly when using bigger datasets, then this may be due to an underpowered iGPU solution. Particularly the second and third generation Intel-based macOS devices are equipped with minimal GPU acceleration.

In such cases, the intergrated GPU may get overwhelmed and time out causing a watchdog to reset the graphics driver. If this is the case, then the best course of action is to force StarTools to use the CPU, rather than the GPU. To do so, create an empty file named 'openclforcecpu.cfg' (case sensitive - e.g. all lower case) in the StarTools folder.

Unquarantining Star Tools

Users may have to "unquarantine" StarTools, before the OS allows it to run. Alternatively StarTools can be launched via control + clicking (right clicking) on the application, Show Package Contents, navigating to Contents/MacOS and clicking on the application.

The following command unquarantines StarTools:

```
xattr -dr com.apple.quarantine StarTools.app
```

See the screenshots for more information, or download this detailed guide on homepage / download.

Windows & Linux OS Hints

Why does the GPU version appear unstable on my Windows machine?

If your older or lower-powered GPU or iGPU appears to be unstable on your Windows operating system in StarTools, and you think it may be struggling with any larger datasets you give it, then the issue may be caused an unsuitable Timeout Detection and Recovery (TDR) allowance.

TDR is a feature that is meant to prevent GPU "hangs". If a task "hangs" the GPU for longer than 2 seconds, the TDR kicks in and will reset the GPU driver.

This Windows default behaviour is not suited for compute-heavy tasks as found in StarTools. Fortunately, it can be corrected by making modifications to the default 2 second timeout value.

https://www.pugetsystems.com/labs/hpc/Working-around-TDR-in-Windows-for-a-better-GPU-computing-experience-777/

Why does StarTools crash in some demanding modules like SVDecon?

If you find StarTools sometimes crashes under heavy load in demanding modules like SVDecon, this may be due to your Operating System being incorrectly configured to provide enough virtual memory.

Ideally, you should configure your system to provide "unlimited" virtual memory. However, if this is not possible or desirable, a good rule of thumb is to make sure your operating system can provide at least 2x-3x the amount of physical memory as additional virtual memory (see tutorial for Windows, or install a package like SwapSpace on your Linux distro).

Star Tools Changelog

1.0.1

* Initial Release

1.0.3

* Made changes to Windows and Linux eventloop to reduce CPU usage in UI thread.

1.0.6

* Made some multi-core CPU specific optimizations
* Optimized Gaussian filter code

1.0.8

* Added second pass function to heal module for more accurate star removal.
* No longer crashes when 'resources' file is missing, instead shows error message.

1.0.10

* Added refractor and newtonian presets to 'Synth"s PSF generator per feature request.

1.0.13

* Added PNG read support
* Overhauled 'Color' module and added new, more effective color balancing algorithm.
* Added pre-SSE3 legacy 32-bit executables to distribution.
* Small UI tweaks.
* Small performance tweaks.

1.0.14

* Added 32-bit signal path version to distribution for faster performance and lower memory usage on legacy systems.

1.0.19

* Revamped 'Color' Module
    - Added average RGB readout to color module for color correction guiding.
    - Removed superfluous per-channel saturation controls.
    - Rewrote
* Made fixes to 32-bit signal path version and made this version the standard for 32-bit systems.

1.0.24

* Revamped 'Layer' Module
    - Added new modes Screen, Power of Inverse (PIP), Desaturate Average, Desaturate Luminance and Invert
    - Added luminance masking (prerequiste for SMI).
    - Added new Median filter type.
    - Added filter type selection (Gaussian or Median).
    - New scrolling behavior
    - Now capable of a huge amount of added functionality; SMI, PIP, many new noise reduction methods)
* Revamped 'Sharp' Module
    - New highly effective general purpose StarTools 'Sharp' algorithm.
* Enhanced StarTools Definition algorithm.
    - Added new De-ringing option.
    - Removed Highpass algorithm in lieu of new Deringing option (De-ringing converts Unsharp Mask into Highpass).
* Added Sobel edge detection to Mask's Auto feature.
    - Rewrote median filter code.
* Added presets to 'HDR' module.
* Added presets to 'Wipe' module.
* Normalization filter in Levels now only applies to bottom end.

1.0.26

* Added new Definition/ST 'Sharp' hybrid sharpening algorithm to 'Sharp'.
* Improved 'Wipe"s pre-filtering.
* Improved 'Contrast"s pre-filtering.
* Added Minimum filter to 'Layer'.
* Added Maximum filter to 'Layer'.
* Added 4th order (e.g. steeper response) option to Luma Mask generation in 'Layer'.
* Added 'Color' Extract mode to 'Layer'.

1.0.28

* Added new Wavelet module for wavelet sharpening
* Added histogram to 'Color' module
* 'Color' ratio in color module no longer affected by luminance retention
* Added lightness filter to 'Layer' module
* Added Multiply Luminance mode to 'Layer' module

1.0.29

* Added new 'Repair' module for star repair and rounding
* Fixed major bug in 'Synth' (didn't work properly on bigger stars)
* Made numerous UI fixes and enhancements across several modules
* Added icon to Linux version

1.0.31

* Added new Deconvolution module
* Fixed small bug in 'Heal' module
* Fixed bug in 'Layer' module (would ignore background & foreground swap when using mask and/or creating luma mask)
* Fixed multi-threading issues on all platforms
* Fixed memory zeroing bug in Gaussian filter

1.0.33

* Added preview cropping mode to 'Decon' module
* Fixed bug in multi-threading code on Windows
* Added presets to 'Bin' module
* Removed desaturate mode from Levels module in lieu of noise filter kernel width control

1.0.34

* Added Reset button to Wavelet sharpening
* Enhanced zooming logic
* Enhanced UI layout
* Rewrote 'Crop' module

1.0.35

* Reduced size of some controls
* Fixed major bug when opening B&W files or files with alpha channel
* Added linked slider to 'Lens' module

1.0.36

* Fixed bug in Levels module ('Color' Noise reduction).
* Fixed issue with +/ buttons working in increments of 1/10th.
* Fixed some auto-fitting logic.
* Added frame around image to make image limits better visible.
* Fixed some gaussian radius parameters specifiying kernel size (double radius) instead of radius.

1.0.37

* Added grayscale hot pixel detection to mask module's auto feature.
* Upgraded 'Repair' module with 'Redistribution' algorithm.

1.0.40

* Fixed memory leak in 'Color' module.
* Made small optimization to Levels module.
* Made small optimization to Mask module.
* Made some minor other small fixes.

1.0.52

* Made many performance optimizations (multi-core and non-multicore) across all modules.
* Fixed some usability issues.
* Added more graceful save file handling.
* Added partial mask detection to 'Wipe'.

1.0.53

* Made additional performance and memory optimizations across all modules.
* Made some fixes to the UI and help text.
* Fixed crash in Windows version in alpha gap interpolation during load file.
* Improved 'Wipe"s noise reduction filter.

1.0.54

* Added dead pixel detection for mono images.

1.0.55

* Fixed path rememberance on Windows.
* Fixed file rememberance on Linux.
* Fixed incorrect FITS loading for images that have negative values.
* Increased precision of the Digital Development paramater in the Levels Module.

1.0.61

* Added edge cropping to debayering routine.
* Replaced memory hungry gradient heal algorithm.
* Reintroduced top-end normalization in Levels module.
* Added Horizontal artifact detection in Auto Mask generation.
* Added Vertical artifact detection in Auto Mask generation.
* Added 'Dust & Scratches' detection in Auto Mask generation.

1.0.63

* Revamped 'HDR' module with new algorithm.

1.0.64

* Minor UI enhancements and fixes.

1.0.65

* Added Stacking artifact detection and warning.
* Added Stacking artifact auto-crop to crop module
* Fixed some minor UI issues.

1.0.66

* Optimized 'Decon' module.
* Fixed potential memory leak in 'Contrast' module.
* Fixed some further minor UI issues.
* Fixed bug in 'Decon' module.
* Enhanced vignetting handling in 'Wipe'
* Change parameter labeling and behavior in 'Wipe' and 'Contrast' into something more humanly comprehendable.

1.1.89

* Fixed auto file naming logic.
* Made preview area selection in 'Decon' more user friendly.
* Optimized Mask application across all modules.
* Added wait cursor to progress dialog.
* Added new Mean of Median Half filter to 'Layer' module.
* Noise filter in Levels module now uses Mean of Median Half filter.
* Added 'Make 'HDR" layering mode in 'Layer' module.
* Optimized 'Decon' module further.
* Added 'Difference' layering mode in 'Layer' module.
* Replaced 'HDR' module and algporithms with new 3rd generation 'Optimize' module.
* Added indicator that 'Do' button needs to be used to update image with non-realtime parameter changes.
* Now graying out buttons that are not applicable.
* Changed 'Synth' module workflow.
* Made Luma masking in 'Layer' module more flexible.
* Fixed Sobel (Edge) detection threshold response.
* Added graying out and message when preview is out of date.
* Added descriptive label to 'Crop' module.
* Fixed bug in Windows version that would cause screen corruption after prolonged use.
* Fixed bug in Linux version that would make ST unresponsive for a small amount of time when switching focus and window needed exposure.
* Made numerous enhancements to UI and workflow.
* Added negative curvature to 'Lens' module.
* Graying out controls when not applicable.
* 'Contrast' module now preserves luminance and doesn't modify chrominance.
* Levels module now called Develop.
* Added option to calibrate top end against stars in Develop module.
* Now using EKOpath compiler for Linux 64-bit
* Added 'Non-linearity' parameter to 'Decon' module.
* Added 'Luma Mask' to 'Decon' module.
* Generic optimizations across all modules.
* New algorithm for horizontal and vertical artifact detection.
* Fixed glitch in single channel Gaussian filter.
* Added screenstretch facility.
* Added hints on appropriateness of linear (screen stretched) vs non-linear (developed) data.
* Added new Fractal 'Flux' module (automated sharpening, noise reduction and detail estimation).
* Added new Banding module.
* Added Radius (circle) mask generator.
* Added'Reject masked pixel' mode in lieu of De-Ringing in 'Decon'.
* Added Amp Glow preset to 'Wipe'.

1.1.90

* Fixed bug where Mask Fuzz would not enable in some modules after creating a mask in-module.

1.1.91

* Reverting back to GCC from EKOpath on Linux 64-bit due to shared object dependencies.

1.1.94

* Changed LRGB's RGB composition weighting.
* New font.
* Reduced button labels and made descriptions longer (full).

1.1.97

* Added new Tame algorithm to Optimize module.

1.1.98

* Made Digital Development parameter response exponential.
* Updated help texts.
* Including suggested processing flow PNG.

1.1.100

* Fixed bug in XML parser, causing problems mostly on Win32.
* Fixed overflow problem in Mask Fuzz on 32-bit systems.

1.1.101

* Fixed glitch in Mask editor when launched from 'Wipe'.

1.1.105

* Major memory optimizations - reduced memory usage by 30% in some cases.
* Minor speed optimizations across all modules.
* Small quality enhancement to Warp algorithm in 'Repair' module.
* Worked around GCC compiler optimization bug in bicubic scaler.

1.1.108

* Fixed minor glitch in Band module.
* Added splash screen.

1.1.109

* Some parameter changes in 'Wipe' module are now real-time.
* Fixed memory leak in Band module.

1.1.112

* Fixed instability bug in 'Wipe' module.
* Fixed instability bug in 'Synth' module.
* 'Off' setting in Stretch Mode now really means off.

1.1.113

* Added healed deconvolution option to allow for deconvolution with a healed copy.
* Auto Mask generator now responds the same irrespective of screen stretch.
* Made memory optimization to Mask editor.
* Added Undo buffer modulation to Fractal 'Flux' module; it is now possible to apply automated Deconvolution.
* Added option in Auto Mask generator to add or subtract newly generated mask from old mask.
* Updated noise mask algorithm.

1.1.114

* Fixed bug where option in Auto Mask generator to add or subtract newly generated mask from old mask was not working as expected.

1.1.115

* Fixed bug which would render some colors incorrectly when screen stretch is active.

1.1.116

* Fixed broken real-time updating of parameters in 'Wipe' module.

1.1.117

* New 'Wipe' Top-end treatment mode ("Wipe' 2.0').
* Removed Min Of Max.
* Added Differential Adaptive Noise filter to 'Layer' module.

1.1.118

* Added crosshairs to 'Crop' module.
* Added reset button to Develop module.

1.1.120

* Changed Bottom End Clipping in 'Wipe' module to become a continuum between Clip < 0 and Add Headroom
* Added Multiply with Gamma Compensation to 'Layer' module.
* Improved 'Add Detail' algorithm in 'Flux' module.
* New 'Life' module.
* Revamped 'Color' module (luminance retention, added 'Sample Mask' button for white balance correction).
* Added 'Flatten' preset to Optimize module.
* Enhanced core algorithm of Optimize module.

1.1.122

* Fixed bug where differing dimensions in 'Layer' module were allowed
* Fixed bug where loading a color image after a b&w image would not enable luminance mixer controls in Develop module
* Revamped 'Wavelet+' module with new algorithm
* Added new 'Equalize' algorithm to Optimize module
* Fixed bug in 'Wipe' (Add Headroom and Clip <0 blending)
* Fixed memory leak in Wavelet
* Enhanced Differential Adaptive Noise Reduction

1.1.123

* Optimized Wavelet+ module
* Re-wrote Life module
* Changed 'Wipe"s 'Noise 'Filter" parameter to 'Dead Pixel 'Filter".

1.1.125

* Fixed bug in Life module in 32-bit versions

1.1.126

* Free Luma power selection in 'Layer' module
* Make 'HDR' is now 'Min Distance to 1/2 Unity' filter in 'Layer' module
* New Max 'Contrast' 'Filter' in 'Layer' module
* Fixed bug in 'Decon' where Mask Fuzz would have no effect in some cases
* Fixed bug in Band module where selecting Horizontal on first launch removed 'please click Do' message
* Added 'intelligent' mode to Boost module
* Addedd 'Mask Fuzz' to Life module
* Added 'select similar color' and 'select similar brightness' brushes to mask editor
* Added Isolate preset to Life module

1.2.147

* New 'Denoise' module
* 'Sharp' module now is new Wavelet 'Sharpen'er
* Fixed bug in rotation and mirroring of images with alpha channel
* Updated Differential Adaptive Noise algorithm
* Made fix to 'Wipe"s Dead pixel filter
* New Dark Noise headroom added to 'Wipe' in lieu of old Bottom End Treatment algorithm
* Clearer font rendering
* Added Debloom algorithm to 'Repair'
* New 'Heal' algorithm
* Made enhancements to 'Repair' algorithm
* Revamped Development module
* New noise suppression in Dev module
* Fixed glitch in 'Contrast' module when using noise filter
* Added horizontal Median filtering to 'Layer' module

* Added 3rd source reference luma masking to 'Layer' module
* Added Fractional Differentation filter to 'Layer'
* Brightness Mask Power selection across all modules
* New Logging module
* UI optimizations
* Improved mask healing quality in 'Decon' module
* Warning in decon module when operating on stretched data
* Revamped 'Contrast' module
* More consistent wording across all modules
* More consistent parameters across all modules
* Can now select items from drop down menu
* Reset button in 'Layer' module
* Fractional Differentiation filter in 'Layer' module
* Algorithm speed ups of 10% across the board (new compilers)

1.3.158 Alpha

* New content aware heal algorithm
* Fixed reading of 32-bits TIFF
* 'Color' module now has maximum RGB display mode
* 'Mask' button now lights up when mask active
* 'Do' button now lights up when appropriate
* New 'Compensate Gamma' option in 'Contrast' module
* New auto development feature in Develop module
* Completely new real-time Optimize module with new 'Reveal Core' algorithm
* Completely new 'Heal' module with content-aware inpainting features
* Cross Operation Noise Reduction And Diffusion (CONRAD) groundwork
* Completely new Deconvolution module with first-in-class deringing algorithm and new CONRAD-driven regularization
* New 'Structural Emphasis' feature in 'Denoise' module
* More optimisations leading to algorithm speed of up to 50%
* Faster 'Bin' algorithm
* Fixes to 'Repair' module to counter artifacts between close stars
    * More multi-core optimisations
    * Improved Halo Killer in Boost module

1.3.159 Alpha

* Fixed non-functioning 'Repair' module
* Fixed button-lighting bug
* Fixed 'Heal' module 'darker-than-value' bug in visual stretch mode
* Moved auto develop feature next to Develop parameter and renamed to 'home-in'
* Fixed Optimize normalisation behavior

1.3.160 Alpha

* Made enhancements to 'Repair' module's Redistribution module
* Fixed bug in 'Heal''s 'Must be darker' behavior

1.3.163 Alpha

* Lots more work on Tracking (formerly CONRAD)
* New Autodev module
* 'Wipe' module now tracks
* 'Decon' module now tracks
* Develop module now tracks
* Autodev module now tracks
* 'Bin' module now tracks
* 'Crop' module now tracks
* 'Denoise' module now tracks
* Linear vs non-linear data now abstracted away from user
* Removed Stretch module
* Tracking now swaps to harddrive insted of internal memory
* Tracking (formerly CONRAD) now enabled and ready for testing
* Fixed severe bug/instability in Gaussian filter code causing problems mainly on 32-bit versions and MacOSX (Thanks TimN!)
* Updated help documentation
* Some memory and speed optimisations

1.3.165 Alpha

* Fix for 32-bit Autodev
* Fix for 32-bit 'Denoise'
* Fixed bug in 'Denoise'
* Hourglass/wait icons to indicate binning and cropping in respective modules while in tracking mode

1.3.166 Alpha

* Misc bugfixes
* Further updated help documentation
* Curve now tracking
* Curve now has redo dialog
* 'Wipe' now tracking
* Autodev now tracking
* Autodev now has redo dialog
* Remove global histogram equalization from 'Layer' module (needs rewrite)
* Made fix to last scale of Wavelet 'Denoise'
* 'Denoise' is now preview-only during tracking
* 'Denoise' can only be used upon switching of tracking

1.3.168 Alpha

* Made optimizations to tracking

1.3.169 Alpha

* Fixed bug in 'Decon' where returning to decon would not start preview rendering
* Added option of de-noising decon detail

1.3.170 Alpha

* Fixed crashing bug in 'AutoDev' and Curve
* Fixed bug where modules would still be disabled after terminating Tracking

1.3.171 Alpha

* Limited support for reading Hubble Space Telescope extended FITS files
* Fixed bug where 'Wipe' would claim image was unstretched even though user indicated it was

1.3.173 Beta

* Added 'Rotate' module
* Fixed bug in Adaptive Noise filter in 'Layer' module
* Added Noise Suppression to 'HDR' module
* Optimised Reveal DSO algorithm in 'HDR' module
* Now feature complete (entering Beta)

1.3.174 Beta

* Fixed a race condition in the Develop and Auto Develop modules

1.3.175 Beta

* Moved tracking swap files to tmp folder on Linux and MacOSX (fixes crash when using Tracking feature in some verisons of MacOSX)

1.3.176 Beta

* Fixed 2 memory leaks in Autodev module
* Fixed small memory leak in 'HDR'
* Fixed memory leak in 'Decon'

1.3.179 Beta

* Fixed MacOSX stability issues by reverting back to O2 level of optimisation

1.3.180 Beta

* Fixed more MacOSX stability issues by initializing threading mutex differently
* Added CPU core detection for multi-threading support (no longer hardcoded to 4)

1.3.182 Beta

* Small multi-core optimisation to 'Decon', 'HDR' and de-noise module

1.3.184 Beta

* Fixed bug in rotate where tracking color data would be garbled
* Fixed bug in 'Wipe' module where stretched color would not be correct
* 'Color' module now using CIELAB for saturation

1.3.185 Beta

* Made optimisations to the 'Color' module
* Fixed jitter in statusbar on multi-core machines
* Made fix to 'Sample' function to not allow sampling of pixels that could potentially be clipping
* Fixed further bug in 'Wipe' module where stretched levels would not be correct
* Reverted some modules that don't need it to old, faster desturation
* Made Unsharp 'Sharp' Mask Voreinstellung ist in 'Flux' module
* Removed defunct "Heal' Masked Pixels' from 'Sharp' module

1.3.186 Beta

* Improved anti-aliasing filters
* Major bug fixes to tracking's detail and levels retention
* Added reminder dialog to 'Wipe' to restretch
* Added reminder dialog to De-Noise to reset mask

1.3.187 Beta

* Small fixes to in-app help and dialog information
* Fixed bug in MaxRGB button behaviour

1.3.188 Beta

* New 'single pixel off' brush in mask editor

1.3.189 Beta

* Fixed some typos in help
* Improved 'AutoDev' algorithm to never blow out stars
* Recommending Autodev for non-linear data only

1.3.190 Beta

* 'Crop' now remembers settings if next crop is same dimensions
* Small fix to luminance loading in LRGB
* LRGB now asks whether data is linear

1.3.192 Beta

* Made fix to TIFF loader to load first image encountered, rather than last

1.3.194 Release Candidate

* Minor texts updates
* Fixed bug in 32-bit 'Contrast' module
* Fixed bug in 32-bit multiplatform system
* Fixed bug in 'Sharp' module

1.3.195 Release Candidate

* Fixed banding issue in 32-bit version of 'Wipe' module

1.3.199 Release Candidate

* Made StarTools netbook friendly
* Fixed overflow issue in Wavelet 'Sharpen' module on 32-bit versions
* Fixed overflow issue in 'Wipe' module on 32-bit versions
* 32-bit now using minimum Pentium 4
* Found malloc bug in GCC 4.7.2 on 32-bit, adjusting optimisation flags to circumvent

1.3.200 Release Candidate

* Enhanced 'HDR' module statistical analysis accuracy
* Added Local Histogram Equalize and Optimize to 'Layer' module
* StarTools log file now also written on MacOSX

1.3.201 Release Candidate

* Fixed instability bug in 'Denoise' module

1.3.203 Release Candidate

* Update in-app documentation for 'Repair' module
* Update in-app documentation for 'HDR' module
* Update log parameters for LRGB module
* Update log parameters for 'Decon' module
* 'Layer' module now available during Tracking mode, but usage not recommended unless expert (added popup)

1.3.206 Final Release

* Accuracy fix to 'Decon"s de-ringing
* Warning on too large images for 32-bit version
* Overflow fix in Wavelet-'Denoise'
* Added Intelligent Despeckle denoising augmentation to Wavelet-'Denoise'
* Fixed crash when clicking keep before result is rendered in 'Sharp' module

1.3.208 Maintenance Release

* Backported 1.4 improved color calibration in 'Wipe' for noisy images (using alternative dark anomaly filter)
* Backported 1.4 bug fix in scaler on 32-bit versions that would introduce artifacts in some cases in various modules

1.3.1.212 Maintenance Release

* Backported 1.4 new mask tools by popular request (lassoo, freehand on/off, line)

1.3.2.215 Maintenance Release

* Backported 1.4 new color module
* Fixed some typos in help text (many thanks to D. 'Sharp'e!)

1.3.2.216 Maintenance Release

* Downgraded optimisation level for older GCC compilers in attempt to enhance stability

1.3.3.222 Maintenance Release

* Backported 1.4 'HDR' enhancements
* Backported 1.4 star mask creation during Tracking mode using linear data
* Made bug fix to 'Decon' module (infinite loop in some corner cases)
* Mask Fuzz in 'Sharp' now Voreinstellung ists to 8.0 during Tracking mode

1.3.3.223 Maintenance Release

* Create 'Decon' preset in Mask Auto module
* Fixed some help text

1.3.3.224 Maintenance Release

* Tapering off PSF in Life and 'Synth' modules to mask FIR window boundaries

1.3.5.227 Alpha

* Fixed small memory leak in median filter
* Made 'Color' module more intuitive
* Improved 'Wipe"s handling of edge cases
* Improved Autodev
* Improved De-Noise module
* Added JPEG export
* Added preliminary JPEG import
* Added automated mask generation for 'Decon'
* Added Restore feature

1.3.5.229 Alpha

* Fixed bug that would trip up the denoise module on MacOSX
* Added warning dialog if Tracking cannot swap to drive (out of disk space)
* Changed Develop's Digital Development response

1.3.5.233 Alpha

* Added dynamic range cruncher to FITS code for 32-bit versions
* Simplified 'Wipe'
* Fixed path bug on Windows
* Added DSLR/OSC flag setting for better noise tracking

1.3.5.236 Alpha

* Reduced number of dialogs for DSLR/OSC flag setting
* Small bug fixes

1.3.5.238 Alpha

* Made fix to outside ROI response in 'AutoDev'
* Improvement to energy redistribution in 'Denoise' module
* Fixed bug/glitch in 'Decon'
* Added fully original to Tracking and Restore
* Improved denoise module (added read noise compensation)

1.3.5.243 Beta

* Improved 'Color' module and added color styles

1.3.5.244 Beta

* Fixed small glitch in color module (human vision style)
* Made some small fixes to 'Wipe' module and added courtesy 'AutoDev' feature

1.3.5.246 Beta

* Improved precision for 32-bit FITS files on 32-bit versions
* Made fix to courtesy 'AutoDev' feature in 'Wipe' module
* Fixed glitch in Mask auto generator sometimes not initialising
* Cancel button now showing in Restore feature

1.3.5.247 Beta

* Made fixes to 'Color' module
* Made fixes to 'Wipe' module (healing)
* Made fixes to 'HDR' module (color retention), presets
* Made fixes to 'Decon' module (artifacts under some conditions)
* 'Lens' module now works while Tracking is active
* Windows version now puts .trk files in executable's directory

1.3.5.248 Beta

* Added LRGB Emulation to 'Color' module
* Made further fix to 'Color' module (Style)
* Added back auto sampling in 'Color' module
* Added Intelligent 'Sharpen'ing to Wavlet 'Sharpen' module
* Removed obsolete 'channel' in Wavlet 'Sharpen'
* New Read Noise Compensation algorithm in 'Denoise' module
* Really fixed glitch in Mask auto generator sometimes not initialising

1.3.5.249 Beta

* 'Layer' module now uses CIELab space for inheriting chroma or luma of foreground
* Improved Read Noise Compensation algorithm in 'Denoise' module

1.3.5.250 Beta

* Fixed glitch in Read Noise Compensation algorithm in 'Denoise' module (ringing around bright stars)

1.3.5.251 Beta

* Fixed another glitch in Read Noise Compensation algorithm

1.3.5.255 Beta

* Fixed rare glitch in FITS 32-bit Integer loading code
* Fixed showing 'demo user' even when registered on some windows installs

1.3.5.258 Beta

* Updated more in-application help
* Fixed bug in 'Denoise' module when used in non-tracking mode

1.3.5.259 Beta

* Improved brush behavior in Mask editor

1.3.5.260 Beta

* Improved 'HDR' and 'Sharp' effectiveness

1.3.5.261 Beta

* Marginal improvement to 'Wipe' color handling of brighter objects
* Improved de-ringing algorithm in 'Decon'
* Changed styles to 'scientific' and 'artistic'

1.3.5.262 Beta

* Fixed bug in 'Sharp' that would cause mask editor to malfunction when launched from 'Sharp'

1.3.5.263 Beta

* Tracking files are now being compressed and take up less space
* Reverted back to mean of median half filtering for 'Wipe"s dark anomaly filter

1.3.5.264 Beta

* Fixed two buffer overflow bugs

1.3.5.265 Beta

* Fixed Tracking corruption bug

1.3.5.267 Beta

* Change message in 'Color' module
* Improved read noise removal algorithm

1.3.5.268 Beta

* Removed auto-normalization upon loading
* Set 'AutoDev' ROI influence to 15% by Voreinstellung ist
* Improved Tracking precision under heavy gradients and light pollution conditions

1.3.5.269 Beta

* Further improved read noise removal algorithm

1.3.5.272 Beta

* Fixed bug in denoise preview when using read noise removal algorithm

1.3.5.273 Beta

* Switched to GCC 4.8 for Windows compiler

1.3.5.275 Release Candidate

Fixed bug in 'Restore' function

1.3.5.276 Release Candidate

* Increased ratio range in 'Color' module
* Enabled 'smart' dynamic range optimisation in FITS loading routines if range information is lacking (e.g. HSTLA data)

1.3.5.278 Release Candidate

* Made TIFF loading more robust against incorrect numStrips count (Lykeos)

1.3.5.279 Release Candidate

* Enabled 64-bit FITS loading
* Fixed bug in 'Restore' function
* Added date/time to log

1.3.5.281 Release Candidate

* Added some help text

1.3.5.283 Release Candidate

* Fixed small bug in 'Denoise' where nothing would render after popup unless a parameter was changed
* Fine tuned precision in some of the Tracking code (impacts various modules)
* Added grain size calibration (previously Redistribution Kernel) to setup screen in 'Denoise'
* Major improvement to Read Noise denoising algorithm
* Major improvement to general 'Denoise' algorithm
* 'Denoise' may now detect and fix small ringing artifacts introduced by 'Decon'

1.3.5.284 Release Candidate

* Fixed creation of small 'dots'/artifacts in the 'Decon' module when dealing with very noisy data.
* 'AutoDev' yields better results when redoing stretch after many different operations
* Renamed 'Decon' preset in Auto Mask generator to 'Fat Stars' to avoid potential confusion
* Slightly enhanced results in decon when working with color data

1.3.5.285 Release Candidate

* Fixed crashing bug in 'Color' module when clicking Sample under some circumstances (thanks to Alivinillo for reporting this).

1.3.5.286 Release Candidate

* Improved FITS loading behaviour and interpolation of missing data, anomalous data and NaN pixels

1.3.5.288 Release Candidate

* Exposed Multiply Foreground 'blend' mode in 'Layer' module
* Exposed Local Maximum 'Entropy' 'Filter' in 'Layer' module
* Pressing 's' key creates screenshot

1.3.5.289 Release

* Significant improvements to Tracking code (increased accuracy of noise detection in the various modules)
* Significant improvement to 'Wipe' module results

1.4.295 Alpha

* Significantly enhanced noise reduction in 'Denoise' module
* Great number of architectural changes and preparatory work under the hood
* Linux version no longer needs GTK, uses Zenity instead
* MacOSX 64-bit version

1.4.296 Alpha

* Small fixes to 'Denoise' module
* Fix to 64-bit Windows version which would corrupt the image during some operations (ex. Life module)

1.4.297 Alpha

* Fixed bug where 'Synth' module would not allow progressing beyond PSF generation

1.4.299 Alpha

* Made stability improvements to MacOSX version
* Implemented fallback solution for Linux distros without wmctrl

1.4.300 Alpha

* Fixed memory leak in 'Denoise'
* New psychovisual correlation algorithm in 'Denoise'
* Improved Read Noise compensation

1.4.301 Alpha

* Made stability and bug fixes to MacOSX version

1.4.302 Alpha

* Fixed small 32-bit TIFF loading bug
* Improved noise floor handling in 'Denoise'
* New Reveal algorithm in 'HDR' - needs work
* Improved color handling for non-color calibrated OSC/DSLR data
* Rewrote file dialogs for MacOSX *again*

1.4.303 Alpha

* Improved 'Wipe' results
* Fixed memory corruption bug in file dialog code
* Stabilised file dialog code on MacOSX

1.4.304 Alpha

* Fixed bug in 'Wipe' on 32-bit versions
* Enabled O2 optimisations on MacOSX

1.4.312 Alpha

* Fixed bug in 'Synth' module on 32-bit versions
* Added noblink option for user Xiando (gets migranes from the blinking)

1.4.315 Alpha

* Fixed 'white dots' bug in 'Decon' module

1.4.316 Alpha

* Fixed Tracking bug where previously deconvolved data would revert back to non-deconvolved if Develop/'AutoDev' launched after decon
* Fixed bug in LRGB showing old 'stretched/unstretched' dialog
* LRGB module now fills mask by Voreinstellung ist instead of clearing it

1.4.322 Alpha

* Now handling multiple screens with different resolutions on Windows correctly
* Clicking outside modal dialog with only one action dismisses dialog and executes action (if any)
* Fixed crashing bug in 'Denoise' when clicking 'Back' before first result has rendered

* Fixed crashing bug in 'Bin' when clicking 'Back' before first result has rendered
* Made some changes to UI appearance
* Set Default ist 'strength' parameter for Reveal algoprithm to 1.2

1.4.324 Alpha

* Fixed bug where screen would not update in some cases when parameter changed (introduced in 1.4.322)
* Fixed bug where Max 'Contrast' 'Filter' in 'Layer' moudle would not work
* Increased max. blend/blur radius for Max 'Contrast' 'Filter' and Min Distance to 1/2 Unity 'Filter'

1.4.325 Alpha

* Fixed crashing bug in 'Bin' and Wavelet 'Sharpen'ing modules when pressing back button before initial image has loaded.
* macos Sierra compatibility

1.4.326 Alpha

* < macos Sierra compatibility reinstated (in addition to Sierra)

1.4.327 Alpha

* Fixed sporadic crashing bug during mask save

1.4.328 Alpha

* Removed unused Mask button in Band module
* Fixed bug in 'Decon' lunar/planetary mode
* Added micro-scale wavelet detail regularization to 'HDR' module.

1.4.329 Alpha

* Loading TIFF no longer normalises data
* Made resources file more attribute agnostic
* Added high-DPI functionality (create empty file called highdpi in STarTools folder to activate)

1.4.330 Alpha

* 'AutoDev' RoI can now be set via sliders
* 'AutoDev' RoI now available in log
* Fixed reference to 'Decon' mask preset in 'Decon' module
* Fixed Temporary 'AutoDev' in 'Wipe' module not being mentioned correctly in log
* Improvement to 'Wipe' handling of edges

1.4.331 Alpha

* Put upper limit on amount of pixels per star for Redistribute algorithm in 'Repair' module
* Enabled subsampling control in 'Repair' module at all times
* Fixed LRGB module crash
* Logging more parameters and actions

1.4.332 Alpha

* Further improvement to 'Wipe' handling of edges
* Logging more parameters and actions

1.4.336 Beta

* Fixed 'Layer' reset button resetting offset filter incorrectly
* Improved 'Layer' module offset filter
* Fixed slight color tinge being introduced into mono images by 'Denoise' under certain circumstances
* Fixed memory freeing bug in 'Layer' module's Max 'Entropy' filter
* Preliminary work for hotkeys

1.4.337 Beta

* Fixed slight color tinge being introduced into mono images by various modules under certain circumstances

1.4.340 Beta

* Now logging mask as BASE64 PNG
* Created mouse-in interface for scripting

1.4.341 Beta

* Now logging image size changes

1.4.344 Beta

* Fixed bug in 'Layer' module's offset filter sometimes Voreinstellung isting to 0.5px, rather than 0px
* Fixed bug in LRGB module where RGB blur could not be used while still processing

1.4.345 Beta

* First cut of ST Replay interface

1.4.346 Beta

* Added 'are you sure?' dialog when closing StarTools

1.4.347 Beta

* Further improvement to Tracking accuracy in 'Denoise' and 'Decon', in processing flows where 'Wipe' was used

1.4.350 Beta

* Fixed bug in 'Decon' that prevented color images from deconvolving correctly

1.4.352 Beta

* Made improvements to RGB Ratio and 50/50 'Layer'ing more useful in 'Color' module
* Made improvements to Artistic mode in 'Color' module

1.5.355 Release Candidate 1

* Fixed bug in FITS import module affecting datasets with NaN entries (e.g. Hubble Legacy Data)
* Removed FITS data interpolation for 'corrupt' (out-of-bounds) pixels
* Improved 'Binning function performance for integer factors (e.g. 25%, 50%, 33.333%)
* Improved 'Decon' quality with more sophisticated stretch modeling of Tracking data
* Life module improvements to luma/chroma inheritence (also impacts Isolate preset)
* Life module Voreinstellung ist settings now Voreinstellung ist to Isolate preset
* Fixed small bug in 'AutoDev' that would yield slightly incorrect results in color image in some edge cases
* Fixed bugs in 'AutoDev' and 'Wipe' that would incorrectly designate a mono image as a color image
* Automatically adding TIFF extensions when saving image originally imported as FITS
* Added 'pre tweak' and 'post tweak' views to 'Denoise' for easier evaluation of changes
* Rename to 'Grain Size' to 'Grain Dispersion' in second 'Denoise' module screen (still inherits Grain Size value as a starting point)
* Removed Noise propagation suppression in 'HDR' module when tracking is on (better results are achieved through subsequent 'Denoise')
* Fixed bug in 'Denoise' module that would yield 'purple' coloring in some cases (even in mono images)
* Changed Read Noise Compensation to 'Non-linear Response <' in 'Denoise' module, and replaced algorithm
* Changed some Defaults in 'HDR' module
* Tweak to 'AutoDev' algorithm weighting
* Added 'Hubble' preset to 'Color' module

* Replaced "Color' cast' preset by 'Narrow band' preset in 'Wipe' module
* Tweaked 'Decon' auto mask generator to no longer pick up less bright nebulosity detail
* Revamped LRGB module
* Fixed border rendering glitch of image fit-to-screen
* Added hotkeys for zoom and common functions (-, =/+ and 0 keys for zooming, O for open, S for save, D for done, K for keep, M for mask, ESC for back/cancel/OK, ENTER for OK/Done, B for 'blink' before/after)
* Added screenshot saving functionality (X key)

1.5.356 Release Candidate 2

* Made improvements to zoom code
* Updated some dialog text
* Moved home screen module icons to reflect common workflow from left to right, top to bottom
* Added new 'Entropy' module
* Renamed LRGB module to 'Compose'
* Added 'pre tweak' and 'post tweak' views to 'Decon' for easier evaluation of changes
* Made improvement to 'Decon' de-ringing
* Made improvement to 'Denoise' precision and handling of parameters (Smoothness response is now different and more controllable)
* Fixed RoI message in 'AutoDev'
* Made UI easier to see with some contrast enhancements
* Added 'pre tweak' and 'post tweak' views to 'HDR' for easier evaluation of changes
* Made improvement to Tracking code handling of clipping highlights
* Added 8-bit PNG saving
* Made 'Denoise' chroma denoising more aggressive

1.5.357 Release Candidate 3

* Fixed normalization bug in 'Shrink'/'Shrink' module
* Renamed 'Shrink' module to 'Shrink'
* Added Luminance/'Color' toggle to 'Wipe' module when in 'Compose' mode
* Added message to 'Wipe' module
* Made 'AutoDev' always on in 'Wipe' module

1.5.358 Release Candidate 4

* 'Lens' module did not take chroma into account while in Tracking mode
* Optimized 'Wipe' when mask is used
* Fixed bug in 'Denoise' using wrong color space for gain map
* Made further improvements to 'Decon' noise grain suppression

1.5.359 Release Candidate 5

* ST Replay integration enhancements

1.5.360 Release Candidate 6

* Fixed bug in 'Entropy' blue channel mode (O-III)
* Added presets to 'Entropy'
* Added more helpful descriptions to channel selection control
* Optimized 'Decon' deringing
* Added denoising to 'Entropy' module if Tracking is off
* Fixed issue with 'holes' in over-exposing star cores when using more aggressive 'Denoise' settings

1.5.361 Release Candidate 7

* Small tweak to 'Decon' deringing
* No longer applying energy redistribution (Grain Dispersion) in 'Denoise' to chroma information
* Inital Grain Dispersion selection screen in 'Denoise' now shows black&white (luminance) only
* Using globally remembered noise filter settings (shared by Dark Anomaly 'Filter' in 'Wipe' and 'Contrast', Ignore Fine Detail < in 'AutoDev' and Grain Size in 'Denoise')

1.5.363 Release Candidate 8

* Small tweak to 'AutoDev' Ignore Detail Smaller Than < behavior
* Tweak to Hubble preset in 'Color' module
* Added Legacy preset to 'Color' module
* Added Constancy preset to 'Color' module
* Made Cap Green into a slider
* Made fix to MaxRGB mode
* Max Red Bias Reduce/Increase now set depending on image
* Made fix to 'Entropy' module
* Improved FITS import for Hubble Legacy Archive data

1.5.364 Release Candidate 9, Final Release

* Added histogram to Develop module

1.5.365

* Fixed incorrect labeling in log
* Made minor updates to help
* Fixed glitch in Develop module with global dark anomaly filter setting
* Fixed bug in 'Contrast' module not weighing color channels correctly
* Whitelisted Nebulosity 32-bit float FITS if min/max not specified in header

1.5.366

* Moved Cap Green application to earlier point in the signal flow
* Added warning message to 'Wipe' if color result has not been evaluated by user while in Compose mode
* Data was bayered and not white balanced now uses Compose mode

1.5.367

* Fixed stack overflow in 'Entropy' module on macOS

1.5.368

* Added Detector Gamma and Shadow Linearity controls to 'AutoDev'
* Fixed crash bug in Develop module when using 100% Develop
* Fixed color loading bug for RGB-only data
* Windows NT 4.0 compatibility

1.5.369

* Fixed another crash bug in Develop module when using gamma 0.0
* Fixed bug when loading 32-bit rational TIFF (e.g. DSS)
* Made some text changes
* Tweaked visual Ha preset in 'Entropy' module to be brighten only

1.6.373 Alpha

* Optimized overall Tracking storage usage and processing times
* Added camera model selection in 'Color' module (accessible when data is not white balanced DSLR w/ Bayer matrix data)
* Added narrowband blend selection in 'Color' module
* Added duoband blend selection in 'Color' module
* Added channel remapping selection in 'Color' module
* Added duoband and SHO:OHS presets to 'Color' module

* Added highlight fix option in 'Color' module

1.6.374 Alpha

* Improved 'Decon' deringing of small stars
* Fixed edge cases in 'Decon' where 'Decon' would generate pixel artifacts
* Fixed bug in 'Lens' module introduced during Tracking optimization in 1.6.373

1.6.376 Alpha

* Fixed bug in handling edge cases in 'Decon' where 'Decon' would generate pixel artifacts
* Fixed bug in histogram handling code on 32-bit versions

1.6.378 Alpha

* Fixed memory leak in 'Wipe' module
* Small improvement in overall handling of Tracking fidelity
* Enhanced detail recovery in 'Sharp' module
* Fixed bug where small screens would not show 'Rotate' functionality in main screen
* New Enhance Deringing algorithm in 'Decon'
* MaxRGB mode now shows linear color dominance in 'before' mode

1.6.380 Alpha

* New 'Sharp' module
* Added Solar False 'Color' to 'Color' module Matrix presets
* Added Strategy presets to denoise
* Fixed memory overflow bug in 'Shrink' module

1.6.382 Alpha

* CPU cache-friendly data structure optimizations across the board
* Multi-threading optimizations across the board
* Removed camera-specific white balance reset
* New Grain Equalization 'Denoise' module ("'Denoise' 2")
* Fixed High DPI issue on Windows
* New mask generation and handling during 'Sharp'
* Tweaked weighting algorithm in Compose module
* Fixed "type of data" popup logic/behaviour for different cases (e.g. Compose module)
* Added conservative auto-mask generation
* Added shrink and grow-to-ratio presets and reset functions to 'Crop' module

1.6.383 Alpha

* Made some tweaks to 'Decon' and 'Color' presets
* Changed parameter names of 'Denoise' 2 and added a "Mode" parameter to switch between statistical and psychovisual noise reduction
* Retina screen improvement on macOS

1.6.386 Alpha

* New 3D Stereo synthesis module

1.6.387 Beta

* Fixed bug in Develop where Skyglow was sometimes ignored
* Whitelisted NOAO IRAF 32-bit float FITS min/max handling
* Added custom PSF by guide star functionality to 'Decon'
* Fixed bug where 'Decon' would reset parameters upon returning from mask editor

1.6.388 Beta

* Whitelisted DeepSkyStacker 32-bit float FITS min/max handling
* Automatically disabling Channel Interpolation parameter in Compose module if all color channels are accounted for
* Fixing some typos, wording, naming
* Added "support" functionality in main screen
* Help functionality in modules now opens relevant web content in addition to regular off-line content
* Added Circle of Confusion primary PSF option to 'Decon' module
* Added dynamic PSF mode to 'Decon' module
* Added intra-iteration back and forward propagation to 'Decon' module
* Added more guidance and warning messages for dynamic PSF mode to 'Decon' module

1.6.389 Beta

* Changed links to live digital footprint (including Notifications link)

1.6.390 Beta

* Fixed rounding issue in 'Bin' module
* 'AutoDev' now remembering RoI
* Logging parameters of 'Synth' module properly
* Fixed crash upon immediately using 'Color' module first thing when image is mono

1.6.391 Beta

* Fixed crash upon immediately using 'Color' module without Tracking

1.6.392 RC1

* Dialog, help tweaks

1.6.393 RC2

* Fixed glitch bug in 'Crop' module when choosing fixed aspect ratio preset grow buttons
* Added 'Color' button for help cropping image based on chroma information (courtesy temporary 'AutoDev' applied)

1.6.394 RC3

* Improved star shapes and depth rendering somewhat for Web 2.5D in 'Stereo 3D' module
* Whitelisted Fistwork 32-bit float FITS min/max handling
* Added crosshair in 'Decon' star selection
* Added snap-to-star in 'Decon' star selection
* Setting Dynamic PSF mode and quality to high by default if sample star selected in 'Decon'
* Increased default Mask Fuzz in 'Decon'
* Changed Dark Anomaly Headroom default to 50% in 'Contrast' module

1.6.395 RC4

* Added crosshair in 'Decon' star selection
* Added snap-to-star in 'Decon' star selection
* Setting Dynamic PSF mode and quality to high by default if sample star selected in 'Decon'
* Increased default Mask Fuzz in 'Decon'
* Changed Dark Anomaly Headroom default to 50% in 'Contrast' module
* Fixed bug preventing help info displaying subsequent launches of 'AutoDev' and Develop
* [DOCUMENTATION] 'AutoDev'
* [DOCUMENTATION] 'Color' (OSCs)

1.6.396 RC5

* Fixed crash in Stereo module on Windows when saving VR scene

1.6.399 Final

* Fixed Windows 32-bit "illegal instruction" crash on older machines

1.6.400 MR1

* Backported 1.7 memory leak fix in 'Decon'
* Backported Windows >4GB file I/O handling fix
* Backported 'AutoDev' stability fix

1.7.401 Alpha

* 1.6.395 RC4 fork
* GPU acceleration rewrite of signal evolution Tracking engine backbone
* GPU acceleration rewrite of all modules
* UI overhaul (ongoing)
* Whitelisted CCDStack FITS import treatment
* Added Zenity and wmctrl dependency popup for Linux version
* Reverse sRGB stretch maintains stretch for visual result
* Rewrite of 'Denoise'
* Improved 'Denoise' and 'Denoise' 2 results
* Rewrite of 'Decon' with Poisson Maximum A Posterior ("MAP") (Hunt and Sementilli, 1992; and Hunt 1995)
* Forward port of 1.6.396 VR export fix for Windows
* Forward port of 1.6.398 fix for Windows 32-bit "illegal instruction" crash on older machines

1.7.402 Alpha

* macOS crashes fixed for AMD drivers (many thanks to J. Scharmann)

1.7.403 Alpha

* Added AutoMask function in Wavelet 'Sharpen'
* Fixed memory leak in 'Decon'
* Improved deringing behavior in 'Decon'

1.7.407 Alpha

* Added new AutoMask generation behavior in 'Decon' and 'Sharp'
* Enhanced mask generation functionality with signal evolution tracking
* Fixed edges cases, handling 0-size GPU data processing
* Now works on macOS with Intel UHD integrated graphics (many thanks to Cytan299)
* Now selecting discrete GPU where dual graphics are available

1.7.412 Alpha

* Fix for RX400/500 series on Windows (many thanks to Peter and Jim)

1.7.416 Alpha

* Middle mouse button click & drag now pans image
* Scroll wheel now zooms
* Left click now toggles between before/after
* Added walking noise reduction feature to 'Denoise' for users who cannot dither (e.g. EAA)
* Added noise grain retention parameter to 'Denoise'
* Tracking file I/O fixes on Windows for files > 4GB
* Fixed crash condition in 'AutoDev' when exiting module mid-processing
* Removed 'Denoise' and 'Denoise'2 from home screen, replaced it with Track

1.7.417 Alpha

* Fixed scrollwheel triggering clicks

1.7.418 Alpha

* Rewrote 'Wipe' for improved multi-core, GPU utilization and reduced memory consumption
* Rewrote 'AutoDev' components for improved GPU utilization
* Fixed auto-mask cancellation loop
* Sped up mask cancellation
* Fixed potential crash edge case when cancelling 'Decon' mid-processing
* Fixed small rendering glitch in 'Entropy' module when tracking is off
* Fixed bug in 1-bit PNG loader
* Ignore Secondary PSF if no star selected in 'Decon'
* Show mask coinciding with selected star as red, not coinciding as green in 'Decon'

1.7.419

* Fixed compatibility issues with some older GPUs and iGPUs
* Fixed instability edge case in 'Decon'
* Fixed instability edge case in 'AutoDev'

1.7.420

* Changed Track button to "Track/NR" to indicate noise reduction functionality
* Added dedicated Bi-color and DSLR/OSC composite modes to Compose module
* Changed 'Crop' preset labels from < and > to -- and ++
* Optimized 'Wipe' for GPU and better multi-core usage and made 'Wipe' real-time module; removed "Do" button
* Removed obsolete "Mode" parameter from 'Wipe' (already taken care of by luminance / chroma separation)
* Rewrote 'Shrink' module (now also availabe during Tracking)
* Change 'Color' Exclude in auto mask generator to "Include Only", added cyan and changed option descriptions to be more helpful
* Set Full Mask option shortcut when launching Life module
* Made some multi-core optimizations to Life module

1.7.421

* Fixed bug in Life module not acknowledging mask choice

1.7.422

* Turned off popup dialog effects
* Fixed rendering glitch when opening file dialog in Compose module
* Updated in-app documentation for Wavelet 'Denoise' (walking noise)
* Updated in-app documentation for 'Shrink'WALKING NOISE SIZ

1.7.424

* [DOCUMENTATION] 'Denoise', Walking noise
* Added SNR improvement label to 'Bin' module
* 'Wipe' re-write
* Mask editor now inherits 'Wiped image for improved visibility of signal
* 'Wipe' 'AutoDev' now mimics Autodev's default
* Added more presets to 'Wipe'
* Removed setup screen for 'Denoise'2
* Added Pre-tweak/Post-tweak functionality to 'Wipe'
* Removed now obsolete Band module
* Added device info to splash screen

* Fixed bug in 'Shrink' where 'Color' Taming would be appplied regardless of mask set
* Improved handling of areas protected by mask in 'Sharp' (e.g. widefields with many stars)

1.7.425 Alpha
* Fixed potential crashing bug in 'Wipe'
* "Before" (both Luminance and 'Color') in 'Wipe' now shows diagnostics 'AutoDev"d version of linear data (and no longer previously stretched image), for easier evaluation of before/after effects
* Fixed precision issue in 'AutoDev' GPU version (corrupted highlights in rare cases)

1.7.426 Alpha
* Fixed 'Wipe' before/after button label not changing
* Fixed bug in GPU version of 'Wipe' that would crash some drivers due to overload (capped number of iterations)

1.7.427 Alpha
* Fixed edge cases in 'Denoise' and 'Denoise'2 that could result in a crash or GPU overload
* Fixed stability bug in 'Denoise'2
* Fixed egde cases in GPU code to avoid overloads
* Added Adaptive Multi-Axis Bias Conservative and Aggressive modes to 'Wipe'

1.7.430 Alpha
* New Super Structure module superseding Life module
* Further GPU optimisations, including in 'Synth' module

1.7.431 Alpha
* Fixed small glitch in Super Structure when processing images that are exactly powers of 2 in size
* Fixed bug in Super Structure's Brightness preset (now selects at least one of the brightness modes)
* Correct version number on macOS

1.7.434 Alpha
* Fixed mask-related glitches in Super Structure
* Fixed some type-casting related possible GPU-acceleration glitches for some GPU architectures

1.7.436 Alpha
* Made Super Structure module less strenuous on more modest GPUs

1.7.438 Alpha
* Added "large UI" functionality (create empty file called largeui in StarTools folder to activate)
* Fixed issue on Linux/X11 where application would use width/height of full desktop, rather than window
* Fixed bug in Super Structure module that would not trigger re-synthesis of Airy disc model correctly when mask dialogs are used
* Renamed Develop module to Film Develop ('FilmDev')

1.7.439 Alpha
* Further reduced GPU load durations and memory requirements in Super Structure module for some GPU architectures

1.7.440 Alpha
* Fixed critical bug in GPU workload allocation code (not all work was always allocated depending on GPU) introduced in 1.7.439 Alpha

1.7.443 Alpha
* Rewrote 'Contrast' for improved multi-core, GPU utilization
* Star mask generation ("tracked") source now uses deconvolved source if available
* Added "Unaclibrated 2" preset in 'Wipe' module
* Added "Un-halo" functionality to 'Shrink' module
* Fixed rare edge case that would make Super Structure diffraction process fail and leave it with an unfinished model
* Fixed bug where scrollwheel use in module/screen that does not support zoom sometimes caused image to disappear
* Fixed bug on Windows where scrollwheel would not work when StarTools window is located on second monitor
* Fixed bug on Windows where mouse cursor would show incorrect icon when moving in/out of window

1.7.444 Beta
* macOS 11 / Big Sur compatibility

1.7.445 Beta
* Fixed 64FP GPU incompatibility in 'Contrast' module on some architectures
* Fixed small graphical glitch around buttons (showing through previous background)

1.7.447 Beta
* Fixed memory leak in Super Structure module
* Increased GUI library memory limit for enhanced stability
* Fixed out-of-bounds condition in CPU-only version of 'Decon' when using large (>15) kernels
* Buttons now respond to all types of mouse clicks (including scroll wheel up/down)
* Level setters now respond to scroll up events to increase/decrease values
* Updated some in-app doc, as well as log entry labelling
* Re-instated mouse click injection for ST Replay

1.7.448 Beta
* Fixed zeroing out bug in Super Structure, causing Airy Disk generation to generate garbage

1.7.449 Beta
* Fixed potential crash bug in 'Contrast' module

1.7.450 Beta
* Fixed small UI glitch in Compose module when update comes through and popover menu is up
* Updated Grain Equalization functionality in 'Denoise' and consolidated 'Denoise' and 'Denoise' 2 into one module as Unified De-noise
* Fixed control sizing and position on various modules

1.7.451 Release Candidate 1
* Fixed button sequence in 'Filter' module
* Removed some unused resources
* Fixed 1st scale in 'Sharp' module not doing anything
* Updated in-app help

* Made some small improvements and optimisations to mask handling in 'Wipe' module

1.7.453 Release Candidate 2

* Fixed bug in 'Wipe' module where clearing mask through popup would not trigger re-rendering
* Fixed gain out-of-bounds condition in 'Denoise' module, causing over-exposed star cores to become black

1.7.455 Release Candidate 3 / Final Release

* Fixed bug on Windows causing 'opencldeviceindex.cfg' (OpenCL device selection) to not work.
* Fixed in-app documentation/help not showing in 'Shrink' module
* Added 'openclplatformindex.cfg' (OpenCL platform selection) for systems with multi-vendor GPUs on Windows

1.7.456 Maintenance Release 1

* Added 'openclforcecpu.cfg' switch (OpenCL GPU/CPU selection) for users that wish to run the GPU version, but have systems with misbehaving OpenCL GPU hardware or GPU drivers (e.g. Apple)

1.7.457 Maintenance Release 2

* Fixed logging of Wavelet Library parameter in Fractal Flux module
* Added logging of response to Global Restretch question

1.7.458 Maintenance Release 3

* Fixed bug that would ignore "undo stretch" and would not undo stretch for linear counterpart of image

1.7.459 Maintenance Release 4

* Fixed bug in Restore functionality
* Fixed incorrect representation of SNR improvement in Bin module (kind thanks to Z. Walinski)

1.7.459.1

* Fixed typo in Mask Auto generator
* Fixed Fractal Flux Setup header not appearing in log

1.7.460 Maintenance Release 5

* Fixed bug in 32-bit float TIFF loading

1.7.461 Maintenance Release 6

* Fixed mouse cursor shape sometimes getting stuck when leaving/entering application window on Windows
* Possible fix for "freeze"(?) issue reported by some users in Compose module

1.8.503 Alpha

* New Spatially Variant PSF Decon (SVDecon) module
* New Narrowband Accent module (NBAccent) + related functionality across various modules and engine
* Added Color Filter Array (aka. "Bayer pattern") artifact filter to Wipe module
* Improved Denoise module
* Optimisations across codebase
* Improved Tracking quality across codebase (particularly large scale noise grain)
* Made Repair module Tracking aware (more work to follow)
* First universal binary build attempt for Apple Intel + Apple Silicon / native M1 ARM CPU

1.8.504 Alpha

* Fixed crashes in SVDecon module
* Further tweaks to Denoise module
* Improved SVDecon module UI
* Tweaked SVDecon PSF spatial variance modelling algorithm

1.8.505 Closed Alpha/Preview Alpha 3

* Improved Sharp module, with improved algorithm, new Protection parameter and enhanced deringing behavior
* Added "Add 1/2 Unity + Clip" Cap mode to Layer module
* Added "Gaussian Highpass filter to" Filter Type to Layer module
* Fixed small bug in NBAccent module that would incorrectly render over-exposing areas (e.g. star cores)
* Fixed samples temporarily disappearing in SVDecon module
* Fixed handling of dynamic range extension in SVDecon module
* Disabled NBAccent module if no accents loaded in Compose module

1.8.506 Public Alpha/Preview Alpha 1

* Fixed multi-threading instabilities in SVDecon
* Further improved deringing and singularity handling in SVDecon
* Optimized SVDecon by cacheing
* Zoom/position retention when switching from various module to mask editor
* Improved Denoise module Grain Size behavior matching with visual result
* Improved Denoise module chromatic (color) residual retention handling
* Heal module now available during Tracking (more work needed)
* Fixed glitch in Heal module code (causing aberrant red hue to come through)
* Added Circle brush (click & drag) to mask editor
* First cut of new dedicated Star Mask Generator ("AltStars" preset) in Auto Mask generator

1.8.510 Public Alpha/Preview Alpha 2

* Fixed GPU initialisation issue for some older GPU/iGPU architectures
* Made PSF bounding boxes in SVDecon a lighter blue to aid visibility

* Fixed(?) "random white/hot pixel" problem in SVDecon
* Now logging parameters of NBAccent setup
* now correctly logging various new parameters

1.8.511 Public Alpha/Preview Alpha 3

* Fixed UI becoming non-responsive to clicks on Windows/macOS in Compose module when loading images while import/compositing code is still running
* Fixed potential issue that may trigger spinlock-detection on some operating systems (now yielding more often during wait for threads to complete/abort)
* Screenshot ('x' button) functionality now enumerating cross-sessions, rather than per-sesssion
* Fixed crash when using presets in Auto mask generator when Tracking is off
* Now logging SVDecon PSF samples to log (as BASE64)
* Added BASE64 import of samples to SVDecon module

1.8.512 Public Alpha/Preview Alpha 4

* Fixed (dumbed down) color preservation in Contrast module
* Fixed disappeared Help button in SVDecon
* Fixed color artifacts sometimes introduced when using Mask Fuzz in Color module
* Fixed SV Decon PSF Load dialog not appearing until after completion of processing (abort early introduced)
* Fixed SV Decon PSF Load only working first time

1.8.515 Public Alpha/Preview Alpha 5

* New HDR module (implementing Spatial Contrast and Gamma Limited Local Histogram Optimisation), with more consistent results, better stellar profile protection, better multi-core and GPU utilization and preview support
* Now disabling exposure length level setters in Compose module for compositing modes that do not need/use them
* Tweaked Color Filter Array (aka. "Bayer pattern") artifact filter in Wipe module
* Added 1:1 ratio to Crop module
* Fixed bug in Tracking (affecting decon and mask generation) in CPU versions of the alpha
* Fixed label in Wipe's Luminance/Color/Narrowband button not changing correctly, depending on what datasets are available
* Cleaned up various other glitches in Wipe

1.8.516 Public Alpha/Preview Alpha 6

* Fixed HDR parameter logging
* Added "Quality" parameter to HDR module, allowing for lower quality (precision) local gamma correction solving

1.8.518 Beta 1

* Changed parameter display in HDR module to include area size

* Reduced number of memory (re)allocations in SVDecon
* Added warning message to Color module if Wipe has not been used

1.8.518 Beta 2

* Fixed exposure times sliders re-enabling when mode does not need them

1.8.519 Beta 3

* Added Laplacian over Gaussian filter to Layer module (Scottk)
* Increased thread stack sizes on macOS and Windows

1.8.521 Release Candidate 1

* Added new Tracking-enhanced artifact suppression to SVDecon
* Fixed memory corruption bug (crash) in SVDecon when using previews
* Fixed 2 memory leaks in SVDecon
* Added NBAccent view to Crop module

1.8.522 Release Candidate 2

* Fixed memory leak in Wipe
* Launching mask editor from Wipe now passes chroma and NBAccent image

1.8.525 Release

* Fixed help documentation not appearing for Correlation Filtering in Wipe when pressing '?' next to it
* Minor tweak to auto mask generation

1.8.526 Maintenance Release 1

* Made HDR module default quality setting dependent on amount of detected CPU cores/threads (<=8 threads defaults to low quality)
* Corrected three matrix options in Color module (two mis-labelled, one incorrect matrix; thank you Mike in Rancho)
* Reformatted matrix options in Color module to mention S-II (if any) first, Ha (if any) second, O-III (if any) third

1.8.527 Maintenance Release 2

* Fixed glitch in HDR module in over-exposing areas under certain circumstances

Document Revisions

V1.0 - Initial Release for Star Tools 1.5

V1.1 - Rework for Star Tools 1.6

- 'Decon' Module Usage added *(Ivo)*
- 'Denoise' 2 Module added *(Ivo)*
- 3D Stereo Module added *(Ivo)*
- 'AutoDev' Usage amended *(Ivo)*
- 'Sharp' Module amended *(Ivo)*
- PixInsight User Information added *(Ivo)*
- OS and Hardware Requirements added *(Ivo)*
- Various corrections
- Improved Formatting
- Rearranged User Notes

V1.2 - Rework for Star Tools 1.6.392

- Updated all 'Descriptions of Controls' sections of the User Notes
- Added 2 workflow diagrams to Tutorials
- Added 'Aesthetics and knowing Your audience' to Life Module *(Ivo)*
- Added 'How to make big/fat/white stars appear tighter' as Usage to 'Shrink' Module *(Ivo)*
- Added 'Dark Anomoly Headroom and more..' as Background Notes to 'AutoDev'
- Added 'Airy Disk Radius' as Background Notes to Life *(Ivo)*
- Added 'Minimum distance to 1/2 unity' as Background Notes to Life *(Ivo)*
- Added 'Detector Gamma ' and 'Shadow Linearity' to 'AutoDev' *(Ivo)*
- Updated workflow diagrams in Tutorials for new 'HDR' and 'Decon' positions and new descriptions
- Changed workflow descriptions in User Notes for new 'HDR' and 'Decon' positions
- Changed order of Module and User Notes sections to match GUI and new workflow
- Updated 'Decon' Module *(Ivo)*
- Reworked User Notes for V1.6
- Added 'Deep Sky Stacker Settings' to Tutorials *(Ivo)*
- Added 'Starting with a good Dataset' to Tutorials *(Ivo)*

V1.3 - Rework for Star Tools 1.6.393

- Added Bookmarks
- Reformatting
- Added "Denoise' 2.0' to User Notes *(Guy)*
- Added 'Tracking' to Main Window use in User Notes (Guy)
- Included updates to 'Decon' in User Notes *(Guy)*
- Included updates to 'Sharp' in User Notes *(Guy)*
- Included updates to 'Wipe' in User Notes *(Guy)*
- Added 'Dark Anomalie Headroom' to 'Wipe' User Notes *(Guy)*
- Included updates to 'Color' in User Notes *(Guy)*
- Added 'Adding HA to OSC data' to Compose User Notes *(Guy)*
- Minor updates to 'Crop', 'HDR' and 'Main Screen' User Notes *(Guy)*
- Added 'Matrix correction and remapping' to 'Color' Module *(Ivo)*
- *Added 'Crop' (Ivo)*

V1.4 - Update for Star Tools 1.6.393

- Added 'Stereo 3D' to User Notes *(Guy)*
- Updated 'Background Notes of 'Life' in User Notes *(Guy)*

V1.5 - Update for Star Tools 1.6.396 RC5

- Added 'Recommended ISO for DSLR cameras' to Tutorials *(Ivo)*
- Added "Understanding 'AutoDev"s behavior" to '"AutoDev"' *(Ivo)*
- Added "OSC Instruments" to '"Color"' *(Ivo)*
- Added "On synthetic luminance generation" to "Compose" *(Ivo)*
- Added "Channel assignment and coloring" to "Compose" *(Ivo)*

V1.6 - Draft for ST1.7

- Added 'Mouse Controls' to Introduction *(Ivo)*
- Added new tutorial for Basis workflow in 'Tutorials' (*Guy/Jochen)*
- Added 'Low Cost GPU solution' *(Ivo)*
- Added 'AstroPixelProcessor Settings' to 'Tutorials' *(Jochen)*
- Added 'ASTAP Settings' to 'Tutorials' *( )*
- Updated 'Denoise' in Modules *(Ivo)*
- Exchanged 'Wipe' User Notes to V1.7 *(Guy)*
- Added "GPU Acceleration and Your Hardware" *(Ivo)*
- *Update in 'Denoise'*
- *Added to Introduction (Ivo)*
- Updated Modules & Features as per online documentation for: Mask, 'AutoDev', 'FilmDev', 'Wipe', 'Contrast', 'Sharp', 'Decon', 'Color', 'Filter', 'SuperStructure', 'Denoise', 'Shrink' *(Ivo)*
- Updated Introduction *(Ivo)*
- Added Full workflow and Special Techniques to 'Tutorials' (*Guy/Jochen)*

V1.7 - Final Draft for ST1.7

- Final Restructuring. Embedding Guy's User notes into module descriptions
- Added new section 'Background Information'
- Added 'Removing background noise using 'Flux' module' to 'Tutorials *(elpajare)*
- Included updates to 'Denoise' in User Notes *(Guy)*

- Included updates to 'SuperStructure' in User Notes *(Guy)*
- Included updates to 'Shrink' in User Notes *(Guy)*
- Included updates to 'AutoDev' in User Notes *(Guy)*
- Added 'Perform Denoise outside of Tracking' to 'Tutorials' *(Ivo)*

V1.8 - Draft for ST1.8

- 'SV Decon' replaced 'Decon' in Modules *(Ivo)*
- Added 'NBA' in Modules *(Ivo)*
- Updated 'Wipe' in Modules *(Ivo)*
- Updated 'Repair' in Modules *(Ivo)*
- Updated 'Layer' in Modules *(Ivo)*
- Updated Full Workflow Diagram in Tutorials *(J Scharmann)*
- Updated 'Recommended ASTAP settings' in 'Appendix' *(polslinux)*
- Replaced step 7 'Decon' by 'SVDecon' in Tutorials *(Guy/J Scharmann)*
- Added ' NB Accent' in Tutorials *(J Scharmann)*
- Updated HDR in *'Modules' (Ivo)*
- Updated HDR in *'Tutorials' (Guy)*
- Updated HDR in *'Background Information' (Guy)*
- Updated HDR in *'Modules' (Guy)*
- Added Simulating RGB image using narrowband data to 'Special Techniques' *(Ivo)*
- Updated Compose in *'Modules' (Ivo)*
- Updated Crop in *'Modules' (Ivo)*
- Updated SVDecon in *'Modules' (Ivo)*
- Added 'Using StarNet++ with StarTools on linear data' to 'Special Techniques' *(Ivo)*

V1.8 - Final for ST1.8

- Added 'Mimicking HDR results of V1.7' to 'Special Techniques' *(Ivo)*
- Added 'Reveal DSO Core preset of V1.7' *(Ivo/Mike in Rancho)*
- Added 'MacOS Hints to 'Appendix' *(Ivo/hixx)*
- Added 'Windows & Linux Hints to 'Appendix' *(Ivo)*
- Udates for NBA wthin 'Tutorials' and 'Modules' *(Guy)*
- Formatting & Spelling
- Updated 'SVDecon' in 'Background Information' *(Guy)*

V1.9.0 - Draft for ST1.9

- Added 'FilmDev' to 'Special Techniques' *(Mike in Rancho)*
- Updated 'Using StarNet++ with StarTools on linear data' to 'Special Techniques' *(Guy)*
- Updated 'Loading Narrow Band data using the Hubble Palette in 'special techniques / Compose' *(Guy)*
- *Added* 'Alt stars preset' to 'Modules /Mask' (Ivo)
- *Added* 'Creating solar image in false Color' to 'Special Techniques' (Ivo)
- Added 'Hands On IC1396 - Elephant's Trunk' to 'Tutorials'*(Ivo /Mike in Rancho)*
- Added Color Mapping and Matrices' to Background Information / Color' *(Ivo)*
- Added 'Deconvolution and Binning' to 'Background Information/SVDecon' *(Ivo)*
- Added 'Apod Mask Color Indicators' to 'Background Information/SVDecon' *(Ivo)*
- Added 'Selecting Sample Stars in Apod mask' to 'Tutorials / SVDecon' (Ivo / Mike in Rancho