

Getting Started in PixInsight

Proving that quick (or easy) is seldom best.

This program is like no other and to some extent divides the astrophotography community. It presents a dilemma, since its powerful image manipulation algorithms are able to deliver the very best results but the user interface is unusual and at odds with the conventions of standard imaging programs, both photographic and astrophotographic. PixInsight (PI) is available for multiple platforms; Windows, OSX, FreeBSD and Linux, though it does require a 64-bit operating system. Most programs have a user guide but presently PixInsight has selected pop-up tools, partial documentation, tutorials and forum posts. Some tools have no documentation. A few experienced users have stepped in and helpfully publish free online tutorials, the most notable of which are Harry's at www.harrysastroshed.com. PI is evolving rapidly and Harry updates his tutorials to keep up with the changes. An Internet search will find other resources, including streamed and DVD media. Personally, I like Harry's down-to-earth style and found his videos and the notes from James Morse on the PI forum helped enormously with my "initiation" into the inner circle.

Many users will try PI without learning about it first but it is not particularly intuitive. For that reason, I have singled out this imaging program for a closer look. Since writing the first edition, Warren Keller published a 380-page reference book. This chapter and the PI content in the following process and practical chapters take things further, with more of a practical application approach:

- help the new user overcome the initial mental hurdles
- understand the general processing workflow
- understand the tools, their purpose and context
- provide insights into alternative techniques
- provide an insight into common tool settings

Although I have a long experience with many programming languages and systems, dating back to CP/M, the user interface took me by surprise and it took a few weeks of head banging before I was able to make some progress. I am by no means an experienced user; there are few that truly can say they are, but this experience allows me to relate to the first-time user.

PI, unlike Photoshop, does not make sole use of a graphical user interface. Its roots are the powerful statistical

and mathematical modeling equations that extract the faintest information from an image, reject unwanted information and remain true to the image intent. There is a philosophy behind PI that avoids manual manipulations of the image or "painting" as they call it, which is more prevalent in the digital imaging programs like Photoshop. PI, like Photoshop, has full color management, something that is lacking in some other astrophotography programs. The visual nature of Photoshop encourages experimentation with immediate feedback of a particular setting or adjustment. This immediacy often hides the fact that there are often many individual steps required to complete a task and the challenge is to remember the sequence rather than the individual actions. For some reason, this makes us more comfortable, but it cannot be called "easy" to interpret a multi-layered Photoshop file with multiple blending modes, masks, and channels. I guess as Photoshop has evolved, our early successes and familiarity with the basic commands help us to cope with the latest advanced tools or simply accept "good enough". At a few points in the first few weeks, I questioned whether the effort of understanding PI was worth the result. My first results confirmed it was and although we live in a world that demands instant results, it is all the more rewarding to accomplish something worthwhile after a little effort. This mantra is mirrored by the PI development team who are rapidly evolving the program and addressing the needs of the amateur community. Regular updates with improved tools, incremental documentation and new scripts are provided free of charge.

As I see it there are three obstacles to using PI effectively; the user interface, the language and the detailed settings of each tool. If you can overcome the first two, you will be able to use PI selectively in combination with other image manipulation programs. As experience grows and you learn from the practical examples on the Internet, further tool settings will become more familiar and find a place in your image processing sequence. Each image and the challenges it presents are unique and one soon learns that for a particular telescope and sensor there may be ball-park settings for a particular tool but the best setting can only be arrived at by experimentation. In some cases, particular tool settings are closely linked to image

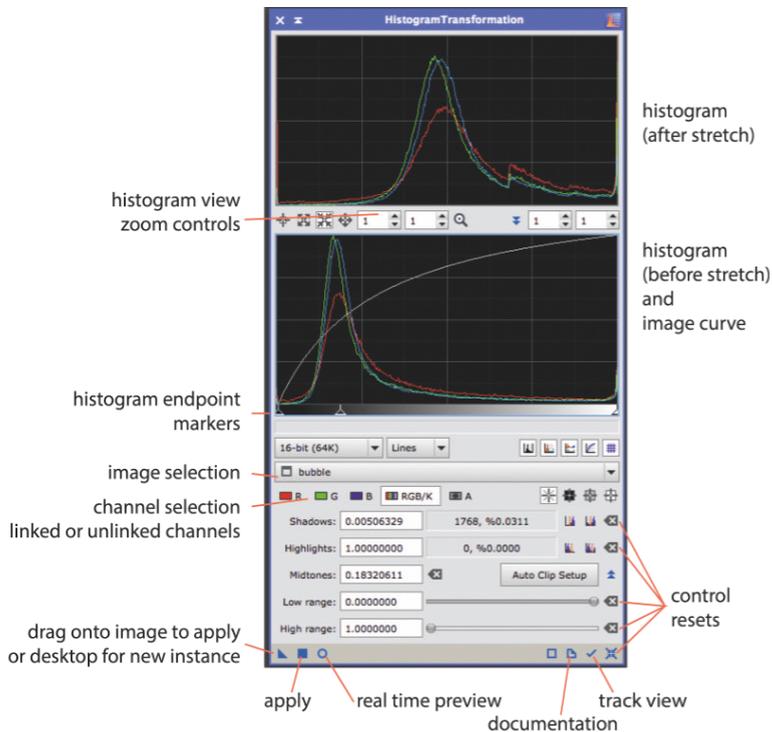


fig.1 A very important function is the HistogramTransformation tool. It is responsible for the non-linear stretches that reveal the detail in deep sky images. It has zoom controls that facilitate precise location of black points and expands its view. It usefully tells you if you are clipping pixels and how many. The blue control icons along the bottom are found in many other process dialog boxes.

statistics, such as background levels, median absolute deviation (MAD) and noise levels.

The User Interface

This is perhaps the most striking immediate characteristic of PI (fig.2). There is a drop-down menu, toolbar, a sidebar and a bottom bar and yet few things are familiar or similar to any other image processing program. This can be quite daunting at first. The foundation of PI is a set of powerful mathematical processes that are applied to an image. They are treated like subroutines in a DOS program. Indeed they can actually be used in a command line interface, similar to the old typed commands in MS-DOS but more conveniently using a scripting language. To make them more useful and to avoid having to type a command and parameter string, the processes (aka tools) have graphical dialog boxes that offer a modern method of selecting files, settings and options. Some of the simpler tools have a real-time preview option; it all depends on the mathematical algorithm being used. (Similarly, other astronomy programs such as Maxim DL do not offer real-time previews on all of their image

processing commands.) The key word here is “real-time”. Some of the algorithms are just too processor intensive for instant results. PI treats the image as sacrosanct and it will not apply or save any image manipulation without an explicit command. For this reason, it uses image previews that facilitate a quick evaluation of a tool setting without affecting the main image. When the setting is just right, the tool and its settings can be applied to one or more images or stored for later use. This is a very useful feature. In this way, it is possible to have multiple tools, say the HistogramTransformation tool, each with a different setting, saved to the desktop as an icon, stored in a different workspace for convenience and saved in a project for later recall. To save time, tools can be combined into a processing sequence to run as a batch and equally, image files can be grouped together and a tool (or tools) applied to all of them. Indeed it is possible to run a batch of processing steps to a batch of files for the ultimate in automation.

When you think about it, this approach allows you to customize your tools to give precisely the effect you want and store them for immediate use or apply later on similar images. The trick is to work out what the tools do, their sequence and what settings give the desired result. This is the part where other users’ recommendations may establish a starting point for further experimentation, to adapt to your image’s quirks.

The dialog boxes have some unusual features too. They have a number of buttons on their margins which have a special purpose. Fig.1 shows the most common of these and what they do. Depending on which button you use, a tool is applied to an image, all open images or a batch of images listed in the tool.

Image file windows also have buttons on their margins for quick access to common zoom tools (fig.2). The tab at the side of the window has special powers too. For instance, if you drag the tab to the desktop a duplicate image is formed. Drag the tab back onto the sidebar of the first image and the image is used as a mask. (The image tab changes color to remind you that there is an active mask on the image.)

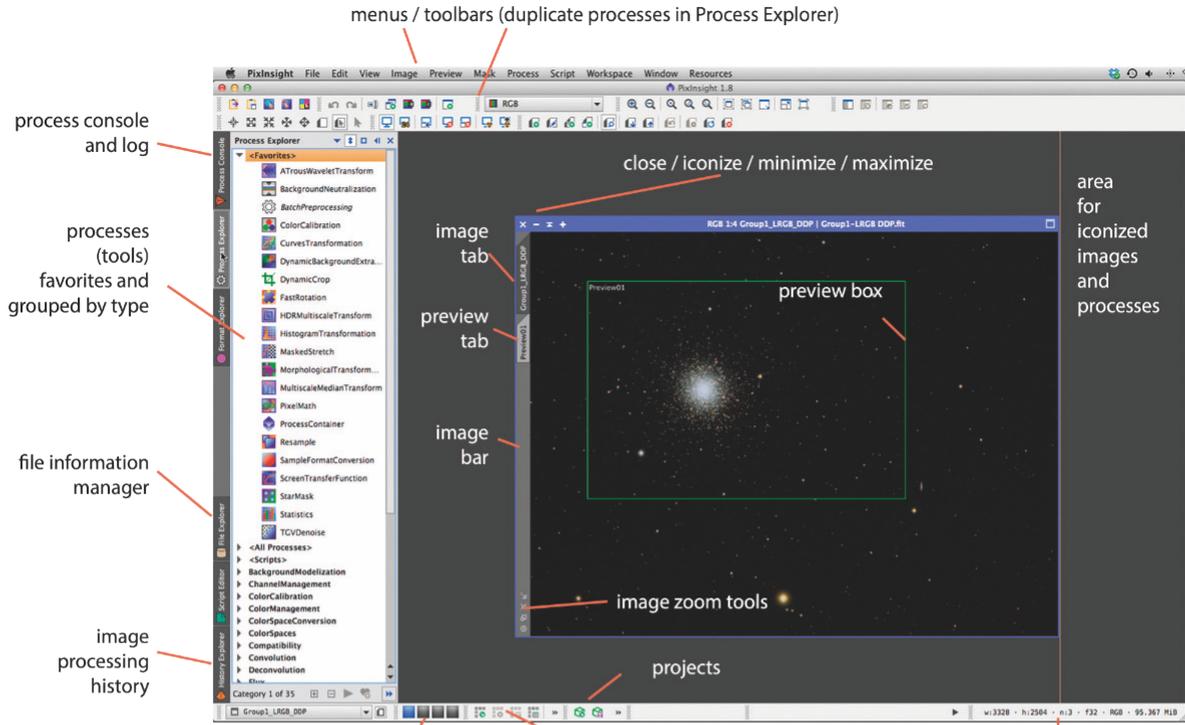


fig.2 The main imaging window of PixInsight has the customary menus and customized toolbars and additional tabs on the left-hand side that give access to process, file and history information and an alternative way of selecting processes (tools).

An image window has some unique controls around its edges to control the zoom and size of the window and a left-hand border that acts as an active area for dropping in images for masking or creating previews.

The four “explorer” tabs on the left-hand side give quick access to the file system, tools, processing history and view information of an image. The PI “File Explorer” is designed for image files and gives unique image analysis information that is not available in a normal OS file browser and includes image analysis statistics. The View Explorer shows detailed information on the selected image file. This comes into its own when you have a lot of images to work with. The History Explorer (fig.3) is a useful extension of the History Palette in Photoshop. PI goes further though; the listed

processes store their settings and it is easy to apply them to another image or batch of images, simply by dragging the line item over to an image, icon, desktop and so on. This may be of particular interest when you process one image of an RGB set and then wish to apply the same tools and settings on the other two images. I guess the same is true when processing a mosaic and you require all the image files to seamlessly match in color and tone.

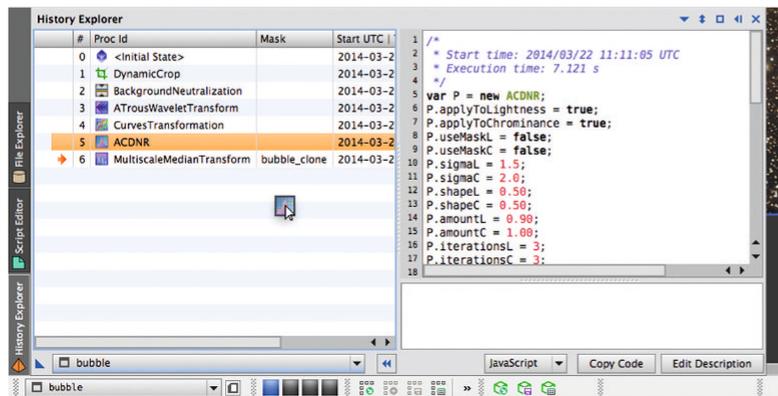


fig.3 The History Explorer uniquely lists all the things you have done to an image and allows you to drag (as seen here) a highlighted process to apply it elsewhere or store it for later use. When you save a project, the history state of the image files are preserved. The process settings are shown on the right in script form, so that they can be copied and edited for sharing and use elsewhere. This is intimidating and powerful at the same time.

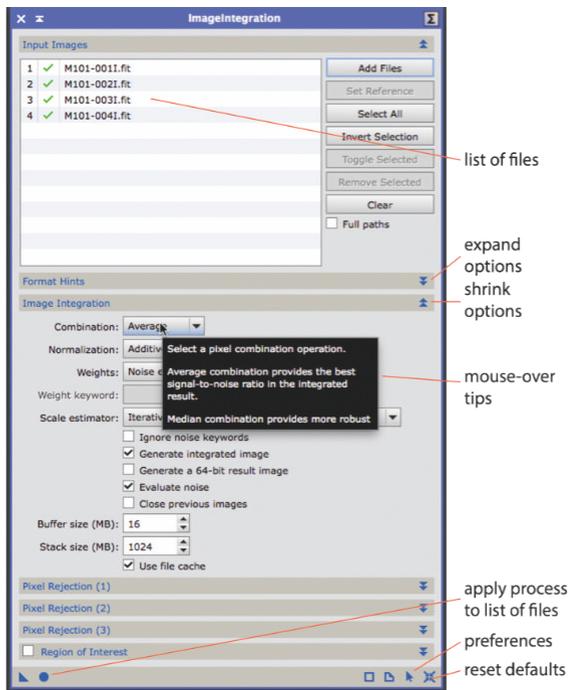


fig.4 A example of a batch process dialog, in this case the statistical combination of files. The file list is at the top and the various controls can be expanded, altered and shrunk back. This is a very powerful tool with many options optimized for darks, bias, flat and light exposures. It can use separate parameters for normalization, rejecting pixels and also for statistically combining the remaining ones.

The Language

I have some issues with the unnecessary academic nomenclature. I am not averse to scientific language but in this case I think it is a barrier to useful and logical user experimentation. The most immediate are a number of tool names and especially their acronyms that offer little clue of their function, especially to a new user. Their icons are no better. That is true of many programs but PI seems to revel in being different and academic. In addition, there are other unique peculiarities: there are processes, icons, containers, scripts, workspaces, projects, objects and previews.

Processes, Objects and Scripts

In the simplest terms, processes are equivalent to tools and objects are typically image files. A script is a small program, often provided by a third party that can provide a utility function, combine existing tools into a more usable form and simplifies batch processing. A script editor is also provided in PI to create your own, or modify existing ones. With a little knowledge, you can read a script and

work out what it does (that is if you have a few cloudy nights). As each tool is applied to an image, the “Process Console” tracks progress with a text-like record. This may appear redundant but it can be useful to scroll back to ensure everything worked on every file and that the various actions were performed according to plan.

Icons are minimized files or tools, floating on the desktop. In a program like Photoshop, the tools are anchored to the toolbars as small icons and a more detailed dialog is expanded when required. In PI, a process is selected from a menu and if the process (tool) is minimized to an icon on the desktop, it keeps its settings and can be renamed to something useful. Its icon form occupies very little space on the desktop and it can be recalled or stored for repeated use or arranged with others to remind you of a processing sequence. An image can be minimized to an icon too. It is handy to keep a partially processed image file close at hand for use later on, say for generating masks or alignment. These unique features allow one to build upon early successes.

Containers, Workspaces and Projects

For the more confident user, “Containers” is a posh word for a collection of image or tools. This is a software device to group together image files for batch processing and assemble a sequence of processes (tools) to apply to one or more images.

Workspaces are just a method to de-clutter the desktop and organize those images and tools that belong together and recall them as a group with a touch of a button. Files and tools are just dragged to the workspace buttons as required.

Projects are a lifesaver. They store all your work, tool settings, work in progress (and each image’s processing history) for a particular session. When you next turn on the computer, run PI and load the project, these are recalled and allow one to carry on with transparent continuity. The workspaces are preserved, and if you organize your work, implicitly record tool settings, workflow and image states. This is not only extremely powerful, a project also serves as a reference for subsequent images. Projects are unbelievably useful, as image processing is often an ongoing activity spanning several sessions.

Previews

Previews may seem obvious enough but PI has a unique slant. One can click on a make preview button and a small outline will appear on your image file. Pulling the edges re-sizes and moves it about. It has its own tab in the image too. You can apply a tool to the preview and get a good idea of what it is going to do. One first trick is to create

multiple previews and quickly compare settings. The second is to combine several previews as a single entity, say, tiled around a central galaxy or nebula. This selection facilitates background or color calibration without being distracted by the higher luminance and color of the subject. Third, the preview has its own processing history. These settings can be copied for later use.

None of the above seem too strange or illogical in their own right. It is just, for a new user, little appears to be familiar or obvious. After a week or two, you will have it figured out and be on the learning curve of how to use the tools themselves. This is when the fun really starts.

PixInsight Workflows

A workflow is a defined processing sequence. You may have come across the term to describe the camera-to-print color management process in digital imaging. In astrophotography, it is theoretically possible to adjust color, noise, sharpness and tonal distribution in any order. In practice, there is a preferred sequence that yields the best quality or put another way, there are certain processing sequences that create issues that are difficult to fix later on. Here is the problem: The initial images start off as a few bright points of light for the brightest stars and everything else appears featureless, or simply black on the screen. All the interesting stuff occupies a tiny tonal range in the shadows. There is nothing else to see in the image, or more importantly, judge processing options with. This detail magically appears with an applied curve, Histogram Transformation or some kind of image stretch. With PI or any other program, do not permanently stretch an image to see the faint details before correcting the basic faults in the image. It is OK to stretch an image preview or view with a temporary screen stretch, so long as the image data is unaltered. A stretched image magnifies all the good and bad points at the same time. Stretching functions change the tone distribution of an image, typified by a curved transfer function between the input and output pixel values, that boost the shadow contrast and compress the highlights. (Altering the endpoints, with a linear characteristic in-between still has a linear transfer function.) Images that have not been stretched are called linear and after stretching, non-linear. Certain processes are optimized to work on linear images and once an image is non-linear, there is no return. The reason for this is simple; the mathematical algorithms behind these processes are designed for the typical parameters of a linear image in terms of signal level, contrast and noise. For instance, most noise reduction algorithms are optimized for linear or non-linear images as they work

on statistical variations that are affected by the image stretching process.

It is possible to classify many PI processes into linear and non-linear camps and in simple terms, the linear processes are applied first, before image stretching. The equivalent processes in the other specialist astrophotography programs generally exhibit similar behavior. Normal photo editing programs are less specific, as they are purposed for general imaging in a non-linear gamma space to improve noise, sharpness and the like. A typical PI workflow is shown in fig.5.

PixInsight Processes

The following chapters step through a typical workflow and draw upon the alternative methods in Photoshop, Maxim DL, Nebulosity and PI. A table of the more common PI processes with a description of what they do and their practical application is shown in fig.6.

When you start using PixInsight, the individual settings in each tool will be just that, individual. Some tools have many controls, so many in fact that it is daunting. My approach is, if in doubt, leave them at their default setting. There are many Internet resources that offer practical examples and good starting points. Further experimentation will give you a feel for what works for your particular image conditions. The mouse-over help windows provide useful information for many tool settings.

After a while, especially with a common camera system, quite a few of the settings become second nature. PI still requires considerable catch-up on documentation, especially considering it is a commercial program. This quick appreciation gives an overview of the kind of tools that are at your disposal and unravels some of the acronyms. Enterprising users continually find new ways to use them for clever effect and in addition, new improved tools replace established ones from time to time. An active and friendly community continues to push the boundaries and shares their insights.

There is no denying PI rewards patience and perseverance and it is not to everyone's taste. It does offer some unique processing tools that make the most of painstakingly acquired images, including high bit-depth processing, multiscale processing to enhance and suppress detail and optimized mask generation. I only started to use PI after the original book's market research identified the need for PI content. It was my intention to use PI to process a few images in the first light section, along with Nebulosity and Maxim DL. After trying it on a few images, however, I quickly noticed a significant jump in image quality, as a result of the complex mathematics that lies at its heart and I have to force myself to use anything else.

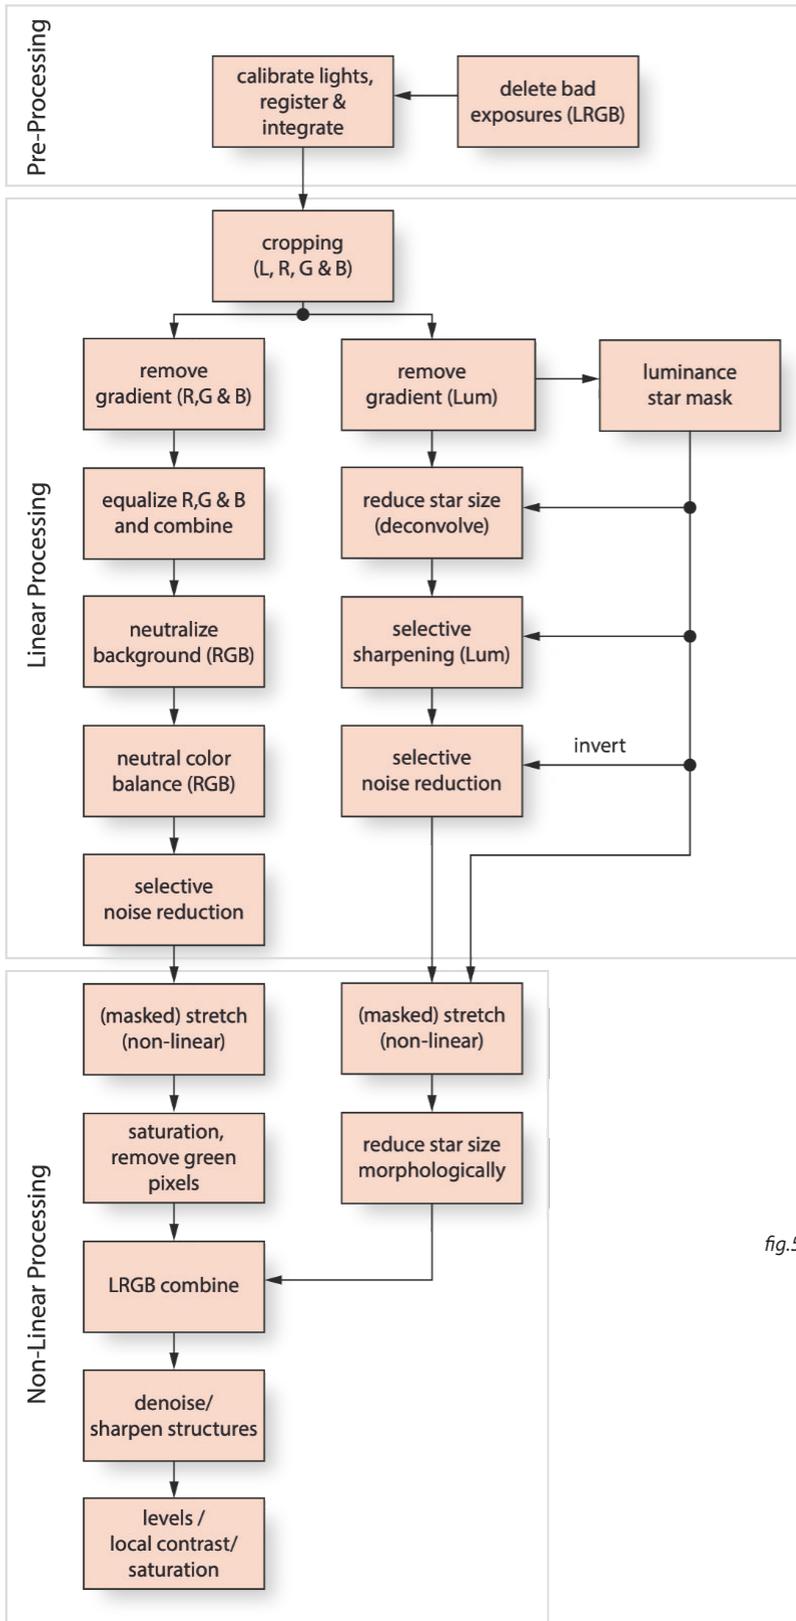


fig.5 A simplified PixInsight LRGB workflow from separate monochrome exposures through to the final image. There are many variations according to taste, for narrow band imaging, one-shot color imaging and also the nature of the deep sky object itself. The processes are generic and in many cases this workflow applies to other imaging software too. In PI, there is often a choice of tools to accomplish a similar function, each optimized for a certain situation. It is finding out what for, how much and when to, that is the fun part.

image calibration tools**Blink and SubFrameSelector utilities (batch process)**

These two utilities are useful for selecting, ranking and discarding image exposures before stacking (integration). These tools allow you to rank and select frames based on a range of image quality parameters.

ImageIntegration (batch process)

This tool stacks images and can be applied to bias, dark and flat calibration frames and calibrated and registered image frames. It combines images and conditionally excludes statistical outliers. It has many powerful options that allow it to optimize the process for calibration master frames and image integration. Some options allow it to normalize frames before averaging them and can also be used to generate a luminance frame from RGB files.

ImageCalibration (batch process)

Processes dark and flat frames from dark, bias and flat exposures as well as calibrating light frames. It works on monochrome images and RGB one shot color images, also called colored filter array (CFA)

StarAlignment (batch process)

After choosing a reference image with the smallest FWHM star value, this generates a new set of registered image files. This also supports aligning files into a mosaic with minimal overlap.

BatchPreProcessing (script / batch process)

The above tools can be used individually to integrate, reject, calibrate and combine registered image pixels. It becomes quite laborious. This script automates the task and creates master calibration files and calibrates, registers image files. It can also integrate the calibrated and registered image files for quick evaluation. Many do this final step using the ImageIntegration tool for finer control.

CosmeticCorrection

This tool is a handy way to remove hot pixels from an image. It allows you to use a master dark file to identify which pixels to remove and replace with interpolated values of the surrounding image pixels.

sharpening tools**HDRMT (High Dynamic Range Multiscale Transformation)**

used on non-linear images

This tool enhances contrast and structure of the image. We perceive it as more detail and dynamic range. The effect can be controlled and directed at objects of different scale. It is normally used with an inverted luminance mask to protect the background.

Morphological Transform

used on linear / non-linear images

This tool is useful for reducing excessive star sizes. It can also make slightly elongated stars round again. It in effect distorts a local area. One useful byproduct is that it reduces star peripheral intensity.

MLT (Multiscale Linear Transformation)

used on linear / non-linear images

This splits the image into its wavelet (size) layers and allows you to process separately. The tool can be used for analysis but also selective noise reduction or detail enhancement at different scales depending on the bias setting for the wavelet size. Effectively supersedes ATWT.

stretch preparation tools**DynamicCrop**

It is important to crop the image to the extent of good data before doing image combination or manipulation. This tool allows you to crop several registered images to the same precise image frame. The best way to see the ragged edges is to apply a screen stretch. The history state can be applied to further images.

LinearFit

normally used on linear images

This tool is useful to linearize the fit between RGB&L channels. It effectively removes the extreme (often red) color cast in the background of a RGB image and achieves approximate balance between the channels. It takes a reference image and linearly scales other images so that their pixels (in a certain range) match background and signal levels. The images must be aligned for this to work.

BN (Background Neutralization)

normally used on linear images

After removing background gradients and the red cast, this tool makes the background a neutral tone. It does this by sampling background image pixels from the entire image or an image preview, within a specified value certain range.

Deconvolution

used on linear images

This tool models the effect of seeing conditions and applies the opposite effect to stars and structures, making them less blurred and more distinct. Uses a Point Spreading Function (PSF) to define how much blur to counteract. The PSF can be generated by measuring stars in the image with the PSF tool. The settings on this tool are very sensitive and it takes many attempts to find one that does not make things worse!

StarMask

It is a good idea to create a star mask before stretching an image, for later use. This tool creates a star mask from an image, using luminance and scale information. It can be adjusted to select big/small stars, bright /dim and extend or blur the mask boundaries. It can be combined with a range mask using PixelMath for use with noise reductions tools.

ColorCalibration

normally used on linear images

Whilst not really a background tool, following background neutralization, the color calibration tool balances the image colors by integrating stars, galaxies or the entire image to obtain a white reference. This allows for maximum color separation but is not technically accurate.

background tools**ABE (Automatic Background Extraction)**

used on linear images

Automatically (or manually) arranges a grid of sample points over image to remove background. Removes background gradients and creates background image to confirm you are not over correcting for deep sky objects.

DBE (Dynamic Background Extraction)

used on linear images

This tool works similarly to ABE but allows the user to place their own sample points to avoid bright stars, nebulosity or faint details.

fig.6 Common PixInsight tool descriptions (continued on next page).

noise reduction tools**ACDNR (Adaptive Contrast Driven Noise Reduction)**

used on non-linear images

Noise reduction tool with a low pass filter and a built in luminance mask to protect lighter areas. It can apply separate levels of noise reduction to color and luminance channels.

TGV Denoise (Total Generalized Variation Denoise)

used on linear / non-linear images

This tool basically replaces the older tool ACDNR. Use with a stretched luminance mask and experiment with strength and edge protection to get the right look.

ATWT (ATrous Wavelet Transformation) (also see MLT)

used on linear / non-linear images

This multi scale tool can be used for noise reduction on linear images, especially at small-scale. When used on linear images, some users only remove noise in the first wavelet layer and use a luminance mask to protect other areas. This sometimes works better on RGBs than TGVDenoise.

MMT (Multiscale Median Transformation)

used on linear / non-linear images

This noise reduction tool is often used with a mask to protect high SNR areas. TGVDenoise is often preferred. Recent enhancements, include differentiation of high contrast structures and a linear mask feature that greatly reduces noise when applied to linear (RAW) images.

SCNR (Selective Color Noise Reduction)

used on linear / non-linear images

This is the preferred method to remove green pixels. Green is not a native color in the sky and the green pixels can be eliminated by this tool and neutralized.

color tools**ChannelCombination**

This tool can be used to replace the color components or channels of an existing image, or to generate a new image from existing channels or components. It supports several color spaces and for instance it can be used to replace the luminance channel in a LAB color space

ChannelExtraction

This tool is the opposite of the one above. This can be useful when assembling composite channels from wide and narrow band filtered exposures and for generating separate masks for RGB images.

LRGBCombination

This tool allows an LRGB image to be assembled out of its constituents, either as separate files or RGB with L. It allows for channel weightings and has built in tools to reduce color noise and boost saturation.

ColorSaturation

This tool allows you to selectively change the saturation and hue of an image. It has a live preview function that helps you fine tune the results. The data input uses a curve, the points of which can be edited and deleted.

ColorCalibration

Described in the stretch preparation tools section.

image stretching tools**STF (Screen Transfer Function)**

This tool gives the appearance of a stretch function but it only does this to the screen image. The auto setting is very useful for a quick appreciation of the image and individual color channels can be tweaked. The settings can be copied over to the histogram tool to apply to an image.

HT (Histogram Transformation)

This tool can alter endpoints but more importantly stretches the tonal range using a gamma-like slider. It appears to be similar to the levels dialog in Photoshop but offers more precise control. This is the primary tool to convert a linear image into a non-linear one. There is no going back! This tool has a live preview function and some functions and read-outs to identify clipped pixels. It is often a good idea to do the stretch in two stages, coarse and fine. This is often used with a mask to protect star bloat or alternatively some use the MaskedStretch tool.

AdaptiveStretch

This is a general contrast and brightness manipulation tool. Some prefer to use this to perform the initial stretch to a linear image. It measures the noise level in an image and adjusts the transfer function to maximize contrast but not intensify image noise.

MaskedStretch (Process and Script)

As the name implies, this applies a stretch function and masks off structures. It does this iteratively to achieve a better result. This prevents star bloat and desaturation. With the right settings, this gives a better result than a HT with a star mask.

CurvesTransformation

Often used with a luminance mask, this tool not only allows subtle control over the tonality of an image (like the curves tool in Photoshop) but also has a saturation and hue modes that allow adjustment based on color. This is very useful for tuning narrowband image color.

Rescale

The rescale function resets the black and white points of an image without clipping pixels. This is done linearly and does not stretch the image with any curve. It can be useful to quickly set endpoints and also clean up after some processes which generate out of limit data.

miscellany**FastRotation**

A tool to rotate images in 90° increments and or mirror transformations. Unlike any other rotation value, these occur without interpolation.

PixelMath

This tool can create an image from a mathematical derivation on one or more images. It can be very simple; adding an offset to all pixels, combining images together or emulating Photoshop blending modes.

ProcessContainer

This enables one to create a sequence of processes, assembled by dragging process icons from the desktop or from the history state of another image.

HDRComposition

This tool blends Linear data of different intensity to extend the dynamic range of the image. It also can be used to replace the centers of burned out stars using a shorter exposure starfield image.