RawPedia Book

Snapshot from December 16, 2018

General Information

Getting Started

Welcome

RawTherapee is a cross-platform raw image processing program, released under the GNU General Public License Version 3. It was originally written by Gábor Horváth of Budapest, and development was taken over in 2010 by a team of people from around the world. Rather than being a raster graphics editor such as Photoshop or GIMP or a digital asset management program like digiKam, it is specifically aimed at raw photo post-production. And it does it very well - at a minimum, RawTherapee is one of the most powerful raw processing programs available.

Installing RawTherapee

Users can just download a RawTherapee installer from http://rawtherapee.com/downloads (http://rawtherapee.com/downloads) or from their package manager. However it is also possible to compile it yourself, should you want or need to. The RawPedia main page has links to instructions on how to do this.

Many versions are available for download, and this paragraph will attempt to explain the difference between them to people who are unfamiliar with how a rolling-release versioning system works. We make new "development" versions almost daily, and once or twice a year we release a new "stable" version, which is nicely packaged with all known important bugs fixed. Any bugs found in the "stable" version will be subsequently fixed in the newer development versions, and these will accumulate until the next "stable" release several months later, and so on. These "development" versions are also where we improve existing tools and add new ones, though it takes time to polish them and to make sure they work well out of the box. On the one hand, "development" versions always have the highest number of bugs fixed, but on the other hand the new tools in these versions may be rough and unpolished and new bugs will appear. If you want to try out new features then get the latest "development" version - you get to take advantage of all the latest bug fixes and you get to test new tools and report problems and ideas back to us at the cost of discovering new bugs. For general use we recommend the latest "stable" version which gives you a generally more polished experience.

Start RawTherapee

The first time you start RawTherapee, you will see the File Browser tab, and it might be empty. You need to point RawTherapee to where your raw photos are stored. Use the directory tree browser on the left of the *File Browser* tab to navigate to your raw photo repository and double-click on the folder to open it. Then double-click on a raw photo to start editing it.

Edit your first image

Once you've opened a raw photo for editing, you will notice that the preview does not look the same as your out-of-camera JPEG did. The article "Eek! My Raw Photo Looks Different than the Camera JPEG" explains why.

Editing is done in the Image Editor tab. This is where you work with RawTherapee to create stunning works of art - or perhaps just apply first aid to your snapshots.

Take a moment to look around this main Editor tab. Notice that there are tabs within this tab - on the right of screen towards the top. These tabs and the controls under them are the Toolbox. You probably have the first tab open and, if you hover your mouse over it, you'll find that it's called the Exposure tab. Below the choice of tabs are the tools the chosen tab contains – Exposure, Shadows/Highlights, Tone Mapping etc. If you click on one of them it will expand so that you can see its contents. Click again and it will collapse. Right-click on one and that one will expand while all others will collapse - a time-saving shortcut. To the left of each tool's label is a power button b which lets you turn it on or off, or in some cases instead of a power button there is a triangular expander. Read the Tools section of the General Comments



RawTherapee in Single Editor Tab Mode Vertical Tabs, showing: 1- Main sections: File Browser (currently opened), Oueue, Editor and Preferences. 2- Panels used for navigating to files and folders. 3- Thumbnails of the currently opened folder. 4- Filters to limit the thumbnails shown to only those which match some metadata or state, 5- Thumbnail zooming and info, 6- Ouick image operations. 7- Sub-tabs of the File Browser: Filter (currently opened), Inspect (to see a fullsized embedded JPEG preview), Batch Edit (to apply some setting to all selected images) and Fast Export (low quality and bypasses some tools but fast saving - don't use this for typical saving!). 8- Right-click context menu (you will typically use this to apply some processing profile to all selected files).



The Image Editor tab.

About Some Toolbox Widgets article for a detailed explanation. Browse through the tabs and panels until you feel totally overwhelmed by all that's available.

Before you start working on an image, here is some important advice – **Don't Panic!** You are in no danger of destroying any of your prized images if you make a mistake. RawTherapee has some features which help you protect your images:

- RawTherapee does non-destructive editing of your raw files. This means that RawTherapee will never, ever change the raw file itself. All changes are stored in sidecar files. You can find out more about them in the Sidecar Files Processing Profiles article.
- When using the Image Editor, you'll see the History panel on the left. This panel shows a history stack of every change you have made to your image. To go back to any step (including when the image was first loaded), just click on the relevant line in the History panel.
- Under the History panel you'll see a Snapshots panel. You can skip it for now, but you'll find it handy when you gain experience with RawTherapee. This panel stores the state of all the tools as a "snapshot". This allows you to easily, for example, tweak your photo to a nice and colorful look and take a snapshot, then tweak it again to a lovely black-and-white look and take a snapshot, and then compare the two just by clicking on either snapshot. (Note: RawTherapee does not save snapshots to the PP3 file yet, it will do so in the future. If you have three snapshots which you want to retain, you will need to click through them and save a PP3 file each time under a unique name).
- As you might expect, Control-z will undo the previous change. (OK, it's not rocket science but it's still handy!)

Basics

- 1. Start off by clicking on the Color tab and expanding the White Balance tool by right-clicking on it. RawTherapee will start with the white balance used by your camera. Most white balance adjustments involve moving the Temperature and Tint sliders, or using the Spot White-Balance Picker on a colorless (neutral gray) patch. Adjust to taste.
- 2. Next, fix the exposure by going to the \(\subseteq \) Exposure tab, expanding the Exposure tool and adjusting it to taste. For now, just use the Exposure Compensation and Saturation sliders.
- 3. If your image is noisy, switch to the Detail tab, zoom to 100% either using the File:Gtk-zoom-100.png button or using the "z" keyboard shortcut key, because the effects of the tools in this tab are only visible in the zoomed-to-100% preview (and of course in the saved image), and enable the Noise Reduction tool by clicking on the power button obleaving the settings at their default values for now. RawTherapee has automatically removed color (chrominance) noise. Luminance noise is removed manually, though leave it for now as luminance noise generally lends a pleasing, grainy, film-like look. As a general rule, when using noise reduction don't use sharpening. Zoom back out to see the whole image either using the button or using the "f" keyboard shortcut key.
- 4. Now you decided you want to fix the geometry and composition of your photo.
 - First make the horizon level, or correct the things which should be vertical such as street lamps or building edges. To easily do this, press the "s" key on your keyboard (the same as clicking the ₺ button), and click-and-drag a line along the horizon or along the edge of a building over the preview. Your image will rotate accordingly and you will automatically be taken into the ₺ Transform tab.
 - To crop the photo, press the "c" shortcut key on your keyboard (or use the the button) and click-and-drag a crop over the preview; you will notice that the Crop tool becomes automatically enabled. There is no need to "apply" a crop it takes effect the moment you draw it. You may want to set the Crop "Guide type" to "none" if it's a problem.
 - Finally, you want to downscale the photo, because who wants to upload a 10MB JPEG to your social network. Enable the Resize tool and leave it at the default settings. Notice that the resizing effect is only applied to the saved image, not to the preview.
- 5. You're all set, let's save it straight away. Click the Save Current Image button, or use the keyboard shortcut Ctrl+s. Save it as a JPG file, quality at "92", subsampling at "balanced". These are good allround settings. Choose a folder where you want it saved to, and after a few seconds your file will be ready in the folder you selected. If you close RawTherapee, the settings you used will be stored in a PP3 sidecar file next to the raw file, so that you can re-open the raw photo in the future and retain the tool settings you used.

Now that you went through basic photo adjustment and are familiar with the steps, let's recap the steps but with more advanced details.

Advanced

Always read each tool's article here on RawPedia before using it, to get a firm understanding of what it does. The articles explain how the tools work in RawTherapee, while the general concepts unspecific to RawTherapee are left to the user to find on Wikipedia or elsewhere.

Be sure to see the Keyboard Shortcuts.

The order of the tools inside RawTherapee's engine pipeline is hard-coded, so from that point of view it does not matter when you enable or disable a tool. However some tools can make a large impact on other tools, e.g. changing exposure may require you to re-adjust color toning, and some tools may require plenty of CPU power to calculate the preview making updates of the preview from then on slow, so it is for this reason we suggest you stick to this general order of operations:

- 1. Start off by making sure that RawTherapee's environment is set up correctly, meaning:
 - Make sure that RawTherapee is using your monitor's color profile if you use a color-managed workflow. Check Preferences > Color Management. You may also need to load the appropriate calibration curves into your graphics card if you built your monitor color profile on top of them, though how you do that is outside the scope of RawTherapee.
 - Make sure that the Color Management tool is configured correctly. Usually the defaults are best. Read the Color Management and Color Management addon articles. If instead of using the color matrix or DCP or ICC profiles shipped with RawTherapee you decide to use an external one, for example a self-made DCP or one from Adobe, load it as the first thing you do, otherwise you may need to re-adjust some of the color tools. Always use an output profile in most cases the default one, RT_sRGB. If you think you're being smart by selecting "No ICM: sRGB Output", you're mistaken.
- 2. If you want to use a Flat-Field and/or Dark-Frame image, do so now, to avoid re-adjustment.
- 3. Now set the correct White Balance. You may fix the exposure first if the image is too dark (or too bright) to see white balance changes.
- 4. Next, adjust the Exposure, using the Exposure Compensation and Black sliders to get the image into the right ballpark. Once in the right ballpark, continue with using both tone curves. Be sure to read the Tone Curve section in the Exposure article to learn why there are two of them and how best to use them they are a very powerful tool!
- 5. In the Basics section above we suggested that you use the Saturation slider (in the Exposure tool). Now that you've learned the basics and are exploring more advanced techniques, we suggest you not use the Saturation slider anymore, and instead use the more powerful CC curve in the Lab Adjustments tool, as it gives you finer control.
- 6. The order of the rest gets fuzzy. Some tools will unavoidably influence others. Carry on with the Lab Adjustments tool and then the rest of the tools in the Exposure tab.
- 7. Then use the Wavelet tool in the A Wavelet tab.
- 8. Then use the tools in the Color tab. The Color Toning tool specifically is very sensitive to exposure changes, so leave it for last.
- 9. Then zoom to 100% and use the tools in the ② Detail tab. Generally, don't sharpen if you're using noise reduction.
- 10. Finally, zoom out again and use the tools in the Transform tab. The reason you left these for last is that they may make the preview image appear blurry, because in order for the preview to be responsive, RawTherapee uses that very preview image you see at the very resolution you see small to show what the tools do, and when you rotate or otherwise change the geometry of a small image, there is a clear softening. This is not a problem when saving as by that point RawTherapee does its processing on the full-sized image, which is slow but of high quality.
- 11. Save, either directly when you want to save a single photo, or via the Batch Queue when you want to process many photos.

You can edit metadata in the Meta tab at any time. For the changes you make in the Meta tab to take effect in the saved image, make sure "Preferences > Image Processing > Copy Exif/IPTC/XMP unchanged to output file" is unchecked. If it is checked, the changes you make in the Meta tab will be ignored in the saved image.

Features

General Features

- All the standard features you would expect from a raw developer and much more,
- A batch processing queue to make tweaking a photo as fast as possible while leaving the CPU grunt work queued up for later,
- Floating point (https://en.wikipedia.org/wiki/Floating_point) engine the only raw developer on the market which does all calculations in precise floating point notation so nothing gets rounded off and lost,
- SSE (https://en.wikipedia.org/wiki/Streaming_SIMD_Extensions) optimizations for better performance on modern CPUs,
- Color management (https://en.wikipedia.org/wiki/Color_management) using the LittleCMS (https://en.wikipedia.org/wiki/LittleCMS) v2 color management system for more precise color handling, providing control over working and output color space,
- Support for reading and tone-mapping 16-bit, 24-bit and 32-bit floating-point HDR images in the DNG format,
- File browser featuring colored tagging, searching (by text matched to the filename), metadata filtering (by file type, camera model, lens model, photo parameters),
- Support for DCP and ICC (https://en.wikipedia.org/wiki/ICC_profile) color profiles, for precise colors or for replicating the "camera look" to match out-of-camera JPEG images,
- A history panel to easily see what changes you made and jump back to a specific point,
- A snapshot panel to work with multiple versions of changes to a photo,
- A flexible UI where panels and some individual elements can be adapted or hidden,
- Easily pan around photos much larger than your screen thanks to pan rate amplification, eliminating the need for numerous and fidgety mouse movements,
- Scroll the tool panels using your mouse scroll wheel without worrying about accidentally misadjusting
 any tools, or hold the Shift key while using the mouse scroll wheel to manipulate the adjuster the
 cursor is hovering over,
- Maximize screen-space by right-clicking on a tool to keep it visible while automatically collapsing the other ones,
- A Before After view to compare your latest change to any previous one,
- Support for PP3 processing profiles (sidecar files (https://en.wikipedia.org/wiki/Sidecar_file)), whole and partial,
- Show various channels in the preview: red, green, blue, luminosity and a focus mask,
- Show various channels in the histogram: red, green, blue, CIELAB (https://en.wikipedia.org/wiki/Lab_color_space) luminance, chromaticity and raw,
- Color clipping (https://en.wikipedia.org/wiki/Clipping_(photography)) indicators in the preview,
- Export Panel with Fast Export Options,
- Keyboard shortcuts to speed up work,
- Command line support to automate RawTherapee using scripts or call it from other programs,
- Support for most cameras,
- Support new raw formats by simply editing the camconst.json file in a text editor,
- Audio feedback to inform you when a CPU-intensive task completes e.g. when the queue is done being processed,
- Preserve IPTC (https://en.wikipedia.org/wiki/IPTC) and XMP (https://en.wikipedia.org/wiki/Extensible_Metadata_Pla tform) of pre-tagged files,
- Adapt RawTherapee to use your system interface's color scheme and widgets or try the provided custom color schemes,
- Localization in almost 30 languages.

Exposure and Color Features

- The Auto Levels tool tweaks your photos to provide a good starting point,
- Various powerful methods of shadow and highlight recovery and reconstruction,
- Post-crop Vignette Filter,
- Graduated Filter (GND),
- Pipette tool, which, when activated for a specific curve, lets you pick a point on the image preview and then places a corresponding adjustment point on that curve,
- Two RGB tone curves, each with four methods of control, for unprecedented control of colors and exposure,
- Hue, saturation and value (HSV) and red, green and blue (RGB) curve adjustments,
- Wealth of Lab adjustments for separate control of colors and luminance:
 - L* curve for control of lightness,
 - a* curve to control the position of a color between red/magenta and green,
 - b* curve to control the position of a color between yellow and blue,
 - LH curve to control luminance as a function of hue,
 - CH curve to control chromaticity as a function of hue,
 - HH curve to control hue as a function of hue.
 - CC curve to control chromaticity as a function of chromaticity,
 - LC curve to control luminance as a function of chromaticity,
 - CL curve to control chromaticity as a function of luminance.
- Avoid color shift by using Munsell correction (https://en.wikipedia.org/wiki/Munsell_color_system),
- Vibrance control,
- Preservation of natural skin tones,
- Tone mapping (https://en.wikipedia.org/wiki/Tone_mapping) based on edge-preserving decomposition (htt p://www.cs.huji.ac.il/~danix/epd/) for a natural look,
- White balance (https://en.wikipedia.org/wiki/White_balance) automatic, manual or one of a wide variety of predefined sources,
- Channel Mixer,
- Black-and-White conversion,
- Several methods of color toning,
- Support for monochrome cameras
- Adaptation of the CIECAMo2 (https://en.wikipedia.org/wiki/CIECAMo2) color appearance model ratified by the International Commission on Illumination (CIE) to maintain accurate colors and to, given a set of initial viewing condition parameters, convert the image so that it will look the same under the target viewing conditions. Image processing using CIECAMo2 is enabled via a number of methods, using curves and sliders. A multitude of tools are adapted to auto-switch to CIECAMo2 mode when in use, including Tone Mapping, Sharpening, Defringe, etc.
- Color management.

Detail Features

- Various sharpening methods:
 - Unsharp mask (https://en.wikipedia.org/wiki/Unsharp_mask) featuring a unique and powerful threshold slider to bring out detail while avoiding halos,
 - RL deconvolution (https://en.wikipedia.org/wiki/Richardson%E2%80%93Lucy_deconvolution) to undo the blur,
 - Edges (https://web.archive.org/web/20110625093654/http://www.rawness.es/sharpening/?lang=en) true sharpening to enhance edges,
 - Microcontrast to enhance texture,
 - Contrast by Detail Levels using wavelet decomposition into five levels of detail.
- Very powerful wavelet-based noise reduction in the RGB and Lab color spaces,
- Impulse noise reduction to eliminate salt-and-pepper type noise.
- Defringe tool for elimination of purple fringing (or other color).

Transformation Features

- Perspective correction,
- Adobe Lens Correction Profile support for automatic correction of distortion, vignetting and chromatic aberration,
- Distortion correction,
- Post-demosaic chromatic aberration correction,
- Pre-crop vignetting correction.

Raw Pre-Demosaicing Features

- Various demosaicing methods to start off by squeezing out the most detail possible from your raw photo:
 - AMaZE,
 - IGV and LMMSE for use with noise reduction to prevent maze patterns,
 - EAHD,
 - HPHD,
 - VNG4,
 - DCB,
 - AHD,
 - Mono,
 - Fast.
- Line noise filter,
- White and black point adjustment,
- Dark frame subtraction to eliminate some forms of noise,
- Flat field correction to easily correct vignetting, lens color cast and sensor dust,
- Manual and automatic chromatic aberration correction.

The Floating Point Engine

RawTherapee 4 does all calculations in precise 32-bit floating point (https://en.wikipedia.org/wiki/Floating_point) notation (in contrast to 16-bit integer as used in many other converters like dcraw (https://en.wikipedia.org/wiki/Dcraw) and also in RawTherapee up to version 3.0) so nothing gets rounded off and lost.

Classical converters work with 16-bit integer numbers. A pixel channel has values ranging from 0-65535 in 16-bit (to increase precision converters usually multiply the 12-14 bit camera values to fill the 16-bit range). The numbers have no fractions, so for example there is no value between 102 and 103. In contrast, floating point numbers store a value at a far wider range with a precision of 6-7 significant digits. This helps especially in the highlights, where higher ranges can be recovered. It allows intermediate results in the processing chain to over- or undershoot temporarily without losing information. The fraction values possible also help to smooth color transitions to prevent color banding. And not the least, floating point makes the life easier for developers which don't need to worry as much about rounding errors or clipping when developing image algorithms for RawTherapee.

The downside is the RAM space floating point numbers require, which is exactly twice that of 16-bit integer. Together with the ever-increasing megapixel count of digital cameras a 32-bit operating system can quite easily run out of memory and cause RawTherapee to crash. Therefore a 64-bit operating system is highly advised for stability. If you have problems running RT on a 32-bit system, try the following:

- As a general rule, you should avoid having folders with too many raw photos in them as each photo takes up memory when displayed in RawTherapee's File Browser tab. Try not to have more than 100 photos per folder.
- RawTherapee uses more RAM while you are using the File Browser tab, so avoid opening that tab
 while you are processing photos.
- Use 4-Gigabyte Tuning in Windows. See "4-Gigabyte Tuning: BCDEdit and Boot.ini (http://msdn.mic rosoft.com/en-us/library/bb613473%28VS.85%29.aspx)" for an explanation of what it is, and find out how to do it by reading the guide "How to set the /3GB Startup Switch in Windows XP and Vista (http://avatechsupport.blogspot.se/2008/03/how-to-set-3gb-startup-switch-in.html)".
- Close other programs while working in RawTherapee.
- Close the Image Editor tab when you're done editing to free up memory.
- Turn off "auto-start" in the batch queue. Only add photos to the batch queue once you are done editing all of them, and then start it. Use the batch queue, do not use the immediate save button.
- Change to a directory with few or no photos in it before starting the batch queue.
- You can free some memory by deleting (or moving to a different folder, or renaming to something like .unusedicc) all the .icc profiles in iccprofiles\input for the cameras you do not use.
- Assure that the Dark-frames Directory and the Flat-fields Directory in Preferences do not point to folders containing raw-files if you do not use Dark-frames or Flat-fields.
- The most memory-intensive tools are Tone Mapping, Contrast by Detail Levels and Highlight Reconstruction using "Color Propagation", so you might need to avoid using them if your machine and operating system are not up to standard.

8-bit and 16-bit

"8-bit" when referring to image formats typically means that the program assigns 8 bits (https://en.wikipedia.org/wiki/Bit) (eight "1" or "0" values, which together make one byte (https://en.wikipedia.org/wiki/Byte), capable of representing an integer (https://en.wikipedia.org/wiki/Integer) value from o [00000000] to 255 [11111111]) to each pixel's color channel, and each pixel in the files RawTherapee saves has three color channels - red, green and blue.

Most modern raw-capable DSLR (https://en.wikipedia.org/wiki/DSLR) cameras use 12- or 14-bit analog-to-digital converters to record the sensor data. This means that when choosing an 8-bit-per-channel output format in your camera, such as JPEG, one loses some information. It's not as simple as it may seem though, cameras record data linearly (due to limitations in hardware design) while JPEG, TIFF and PNG gamma encode (https://en.wikipedia.org/wiki/Gamma_correction) their data meaning that they distribute more values in the shadow range and less in the highlights which better matches the eye's sensitivity. This means that an 8 bit JPEG can display as much as log2((1/2^8)^2.2) = 17.6 stops of dynamic range which indeed exceeds the 14 stops of the current best cameras, which explains why you sometimes can see a camera's shadow noise even in an 8 bit JPEG. However due to the fewer values in highlights we lose precision there compared to the camera. Practically this is not a problem when the output file is the definitive one and will not be processed anymore, however a photo can be vastly improved when saved as raw data and processed using a state of the art raw processing program, such as your's truly - RawTherapee.

Once you have processed a photo in RawTherapee, you are faced with the same choice - to save the image with a color resolution of 8 bits per channel, or 16 bits per channel (only TIFF (https://en.wikipedia.org/wiki/TI FF) and PNG (https://en.wikipedia.org/wiki/Portable_Network_Graphics), not JPEG (https://en.wikipedia.org/wiki/JPEG)). If you plan to post-process your photos after RawTherapee in a 16-bit-capable image editing program, it is better to save them in a lossless 16-bit format. Uncompressed TIFF is suggested as an intermediate format, as it is quick to save and stores all the metadata (Exif (https://en.wikipedia.org/wiki/Exif), IPTC (https://en.wikipedia.org/wiki/IPTC_Information_Interchange_Model), XMP (https://en.wikipedia.org/wiki/Extensible_Metadata_Platfo rm)) of the original file (PNG generally discards metadata!).

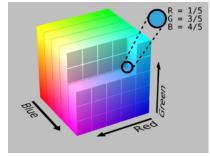
There is some confusion over the naming of 8, 16, 24 and 32-bit files. Here is a clarification, but it gets confusing, so put your tin foil hat on. You do not actually need to read this to use RawTherapee, it is just background knowledge. Each of the red, green and blue channels stored in a JPEG, PNG or TIFF file is actually a colorless image, but when you combine these three colorless images together, you get a color image. This is how all digital representation of images works - color images are always decomposed into their components in one way or another. Of the file formats that RawTherapee can save to (JPG, PNG and TIFF), each pixel contains information for three color channels - red, green and blue. We say "8 bits per channel" to make it clear that these 8 bits apply to one color channel only. The reason is that you might encounter references to "8-bit images", and here it gets confusing, because the person who wrote that may have been referring to a grayscale format which stores only one channel, or to a color format that stores three channels, with 8-bit precision each. Another notation for the very same "8-bit" images that RawTherapee saves is "24-bit". Woo, confusing. Or is it? Each pixel is made of 3 channels, and each channel stores 8 bits of data, so we have a total of 24 bits of data per pixel. It gets worse. Image editing programs can also store a fourth channel, called "alpha". To put it simply, alpha describes how transparent a pixel is. These alpha channels also have a "color resolution" of 8 bits. Both PNG and TIFF files can handle alpha, JPEG can't. If you have an 8 bits per channel image with an alpha channel, it can also be described as a 32-bit image; R (8) + G (8) + B (8) + alpha (8) = 32. The ultimate problem is that you can also have an image that assigns as many as 32 bits per color channel. These images can be

described as "32-bit" images as well as "96-bit images" (because R (32) + G (32) + B (32) = 96). All real HDR files are stored in image formats that assign at least 16-bit floating-point numbers for each pixel per color channel, such as the EXR format, or 32-bit ones, such as the RGBE format. To summarize, an "8-bit per channel" image with three channels (RGB) can also be called a "24-bit per pixel" image, and a "16-bit per channel" image with three channels can also be called a "48-bit per pixel" image. In both cases use the former (the full "x bit per channel" description, don't just say "x bit"!), it's more clear what you mean.

RGB and Lab

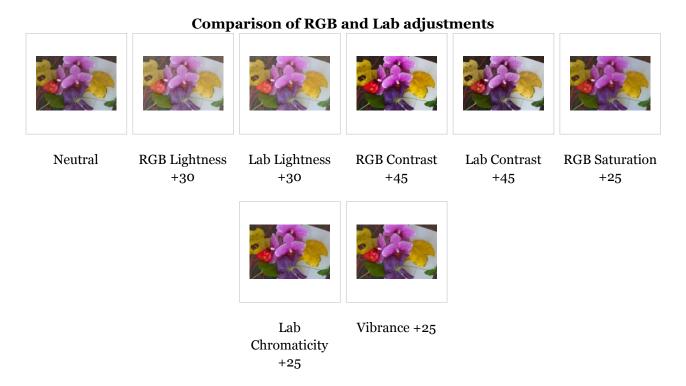
RGB (https://en.wikipedia.org/wiki/RGB_color_space) and $CIE\ L^*a^*b^*$ (http s://en.wikipedia.org/wiki/Lab_color_space) (or just "Lab") are two different color spaces (https://en.wikipedia.org/wiki/Color_space), or ways of describing colors.

Many people wonder what the differences are between adjusting lightness, contrast and saturation in the RGB color space, or lightness, contrast and chromaticity in the Lab color space. RGB operates on three channels: red, green and blue. Lab is a conversion of the same information to a lightness component L*, and two color components - a* and b*. Lightness is kept separate from color, so that you can adjust one without affecting the other. "Lightness" is



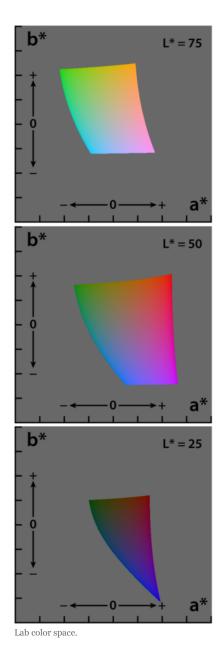
RGB cube.

designed to approximate human vision, which is very sensitive to green but less to blue. If you brighten in Lab space, the result will often look more correct to the eye, color-wise. In general we can say that when using positive values for the saturation slider in Lab space, the colors come out more 'fresh', while using the same amount of saturation in RGB makes colors look 'warmer'.



The difference between the *Lightness* slider in the *Exposure* section (in RGB space) and the *Lightness* slider in the *Lab* section is subtle. A RGB *Lightness* setting of +30 produces an image that is overall a bit brighter than when using a Lab *Lightness* setting of +30. The colors in Lab *Lightness* are somewhat more saturated. The contrary is true for the Contrast sliders; when using a RGB *Contrast* of +45 the colors will be clearly warmer than when using a Lab *Contrast* of +45. The contrast itself is about the same with the two settings. Do not hesitate to use both sliders to adjust saturation and/or contrast. As for the *Saturation/Chromaticity* sliders, setting the RGB *Saturation* slider to -100 renders a black and white image which appears to have a red filter applied, while the Lab *Chromaticity* slider renders a more neutral

black and white image. Positive RGB *Saturation* values will lead to hue shifts (the larger the value, the more visible the shift), while positive Lab *Chromaticity* values will boost colors while keeping their hues correct, rendering a crisp and clean result. Lab chromaticity (via the "*Chromaticity*" slider or "*CC*" curve) is the recommended method for boosting colors.



Making a Portable Installation

RawTherapee and the cache folder can be stored "self-contained" on a USB flash drive or any other mass-storage device.

For Windows

Get the latest build of RawTherapee. Since we want it portable, we don't want the installer, just the bare, zipped program. If the latest version on our website is in simple zipped form without an installer, you can skip this step. However, if it is an installer, you need to first extract the RawTherapee files.

- If it is an Inno Setup installer (.exe extension, all recent Windows installers are Inno Setup ones at the time of writing, summer 2014), get innounp (http://innounp.sourceforge.net/) or innoextract (http://constexpr.org/innoextract/) to unpack it.
- If it is an MSI installer (no recent Windows builds use this at the time of writing), fire up a command prompt and type:

```
msiexec /a RawTherapee.msi TARGETDIR="C:\TargetDir" /qb
```

Replace the name of the MSI installer and the target directory as appropriate. Spaces in the TargetDir path are allowed, as the path is enclosed in quotes.

Let's assume that you've unzipped your archive into E:\RawTherapee, where E:\ is the drive letter of your USB flash drive. Open the E:\RawTherapee\options file, and set the MultiUser option to false. Now when you run RawTherapee, it will store the cache and your settings in subfolders relative to the executable, named mycache and mysettings, respectively, so in E:\RawTherapee\mycache and E:\RawTherapee\mysettings.

See also the File paths page on how to set a different location for these two folders.

When updating RawTherapee, it is recommended to unzip the new version to a new folder and simply move mycache and mysettings into it.

For Linux

Getting RawTherapee to run off a portable medium such as a USB flash drive on various Linux systems is not straightforward due to the nature of Linux systems. While the Windows version of RawTherapee comes bundled with all required libraries to run on any Windows version, Linux distributions differ significantly from each other and as a result a version of RawTherapee built for one distribution is unlikely to run under a different distribution. One way around this is by using an AppImage.

A RawTherapee AppImage is a single file which contains a RawTherapee executable along with all the required files needed for it to run on any Linux distribution. Download it, make it executable, and run it. We are currently in the testing phase regarding AppImages. They are not yet available from our Downloads page, but you can find them on our "development builds" page in the forum (https://discuss.pixls.us/t/download-rawtherapee-development-builds/2924?u=morgan_hardwood).

Regardless whether you use the AppImage or a "proper" RawTherapee build from the distribution's package manager, you will want to be able to hang on to your RawTherapee configuration and processing profiles.

In order to backup your configuration you will want to copy RawTherapee's *config* folder onto your USB stick. Specifically, you want the "*options*" file, your custom "*camconst.json*" if you made one, and any custom PP3, ICC, DCP and LCP profiles. The File Paths article describes where to find these.

The File Browser Tab

The File Browser tab is where you review your photos, select photos for editing, or perform batch-editing operations. It consists of the following parts:

• The left panel

- The "Places" panel on the top links to your home folder, USB card readers, the system's default "photos" folder, or custom folders.
- Below this is a standard tree-type file browser that you can
 use to navigate to folders containing your photos.
 RawTherapee does not complicate things by requiring you to
 import photos into databases as some other software do.

The right panel

- The "Filter" tab lets you show only photos which match the parameters you specify.
- The "Inspect" tab shows a preview at a fixed scale of 100% of the image your mouse cursor is hovering over, which is either the largest JPEG image embedded in the raw file, or the image itself when hovering over non-raw images.
- The "Batch Edit" tab allows you to apply tool settings to the selected image or images. This allows you to quickly enable some tool in many photos at once.
- The "Fast Export" tab lets you quickly process the selected images by bypassing certain tools even if they are enabled in the processing profiles of those images, so that you can get a quick preview of the raw files for example to delete the shots which are blurry or out of focus.
- The central panel shows thumbnails of the folder currently selected.

You can hide the individual panels using the "Show/Hide the left panel +" and "Show/Hide the right panel +" buttons - see the Keyboard Shortcuts page.

When you open a folder, RawTherapee will generate thumbnails of the photos in that folder in the central panel. The first time you open a folder full of raw photo files, RawTherapee will read each file and create a thumbnail based on the embedded JPEG image (every raw photo has an embedded JPEG image, sometimes even a few of various sizes). This can take some time on folders with hundreds of photos, but it only happens the first time you open that folder. All subsequent times you go to a previously opened folder, RawTherapee will read the thumbnails from its cache if they exist, and this will be much faster than the first time you opened that folder.

The JPEG image embedded in each raw photo is identical to the out-of-camera JPEG image you would get if you shot in JPEG mode (or in "RAW+JPEG" mode). This JPEG is not representative of the actual raw data in that photo, because your camera applies all kinds of tweaks to the JPEG image, such as increasing the exposure a bit, increasing saturation, contrast, sharpening, etc.

After you start editing a photo, its thumbnail in the *File Browser* tab is replaced with what you see in the preview in the *Editor* tab, and every tweak you make is reflected in the thumbnail. The thumbnails are stored in the cache for quick future access. If you want to revert to the embedded JPEG image as the thumbnail, then right-click on the thumbnail (or selection of thumbnails) and select "*Processing Profile Operations > Clear*".



RawTherapee in Single Editor Tab Mode Vertical Tabs, showing: 1- Main sections: File Browser (currently opened), Queue, Editor and Preferences, 2- Panels used for navigating to files and folders. 3- Thumbnails of the currently opened folder. 4- Filters to limit the thumbnails shown to only those which match some metadata or state. 5- Thumbnail zooming and info. 6- Quick image operations, 7- Sub-tabs of the File Browser: Filter (currently opened), Inspect (to see a fullsized embedded JPEG preview), Batch Edit (to apply some setting to all selected images) and Fast Export (low quality and bypasses some tools but fast saving - don't use this for typical saving!). 8- Right-click context menu (you will typically use this to apply some processing profile to all selected files).

Use the zoom icons in the File Browser's top toolbar to make the thumbnails smaller or larger. Each thumbnail uses some memory (RAM), so it is advisable not to set the thumbnail size too high ("Preferences > File Browser > Maximal Thumbnail Height").

You can filter the visible photos by using the buttons in the *File Browser's* or *Filmstrip's* top toolbar, as well as by using the "*Find*" box or the "*Filter*" tab. Possible uses:

- Show only unedited photos,
- Show only photos bracketed +2EV,
- Show only photos ranked as 5 star,
- Show only photos with a specific ISO range,
- Show only photos with a NEF extension.

If your screen's resolution is too low to fit the whole toolbar, some of the toolbar's contents (buttons, drop-downs, etc.) may become hidden. To see them, simply hover the cursor over the toolbar and use the mouse scroll-wheel to scroll the contents left and right.

Batch Adjustments - Sync

RawTherapee lets you batch-adjust, or sync, the processing settings in many photos at the same time in generally two ways. It lets you copy and paste a processing profile (a collection of tool settings), in parts or in full, to any number of images. It also lets you select any number of images and adjust any tool in all of them at once (sync), and it lets you do this in two ways. Let's take a closer look.

Both ways involve making a selection of photos you want the processing profile or adjustments applied to. Selections are made using standard key combinations: Shift+click to select a range, Ctrl+click to select individual images, or Ctrl+A to select everything. Both ways are performed from the File Browser tab. The "copy & paste" method can also be done via the Filmstrip.

Copy & Paste

Copying and pasting a processing profile to a selection of images is a very common task. Assume you took a series of photos - for example studio shots, wedding portraits or focus-bracketed macro photos. All images in each series are going to be very similar; they will probably use the same lens, the same ISO, the same white balance, and end up being used for the same purpose. This means that they will all probably require the same processing settings - the same noise reduction, the same sharpening and lens distortion correction, and so forth.

To process the lot, what you would usually do is open any one image from the whole series in the Editor tab and tweak it to your liking. Once you have finished tweaking it, you will apply this image's processing profile to all other images in the same series. To do that, go to the File Browser tab, right-click on this photo and select "*Processing Profile Operations > Copy*", then select the images you want to apply this profile to, right-click on any one of them (it doesn't matter which) and select "*Processing Profile Operations > Paste*". In one quick operation you have replicated the same tool settings in the whole series of images.

Additionally, RawTherapee lets you apply only a part of the copied processing profile, for example only the "Resizing" tool. To do this, use the "Processing Profile Operations > Paste Partial" option instead of the "Paste" option.

Sync

RawTherapee lets you instantly apply tool adjustments to a selection of images. Similar functionality in other software is called "sync". This method is useful for when you don't need to see an accurate preview of your changes, for example when you only want to enable the "Resizing" tool in a selection of photos, because when working in the File Browser tab your only preview are the small and inaccurate thumbnails. This method can only be performed from the File Browser tab because you need access to that tab's batch tools (the panel on the right).

When you're in the File Browser tab, select the images you want to batch-adjust (sync), then use the tool panel on the right to make adjustments. Your tweaks can either replace the existing ones ("Set" mode), or be added to them ("Add" mode). For example if you select two photos, one of which has previously been tweaked with +1EV Exposure Compensation and one which has not, and you set Exposure Compensation to +0.6EV, then the previously-tweaked photo would end up having +1.6EV Exposure Compensation in "Add" mode and just +0.6EV in "Set" mode. The photo which was not previously tweaked would have +0.6EV in both modes. You can decide which tools should work in which mode from the Batch Processing tab in Preferences.

Deleting Files

As RawTherapee is a cross-platform program, it has its own trash bin, independent from your system one if you have a system one.

Using the Trash Bin

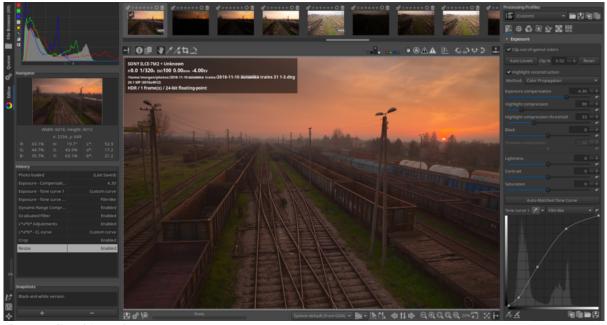
To move files to the trash bin, either use the "Move to trash" button $\widehat{\mathbf{m}}$ in the top-right corner of each thumbnail, or right-click on a selection of files and choose "File operations > Move to trash". These files are then marked as being in the trash bin, but they are not deleted from your hard drive.

- To hide all files which are marked as being in the trash bin, click the "Show only non-deleted images" button not in the top toolbar.
- To see the contents of the trash bin, click the "Show contents of trash" button 📆.
- While you are viewing the contents of the trash bin a new "Permanently delete the files from trash" button m appears to the left of the thumbnails use it to delete all trashed files from your hard drive.
- Click on the "Clear all filters" button ▼ to return to the default view.

Deleting From the Hard Drive

To delete files from your hard drive without using the trash bin, just right-click on a file or on a selection of files and choose "File operations > Delete" or "Delete with output from queue". Both options delete the selected photo and its sidecar file from your hard drive, but "Delete with output from queue" also deletes the saved image whose filename matches the template which you currently have set in the Queue tab, in the "Use template:" field.

The Image Editor Tab



The Image Editor tab in RawTherapee 5.5.

Introduction

The Image Editor tab is where you tweak your photos. By default RawTherapee is in "Single Editor Tab Mode, Vertical Tabs" (SETM) which is more memory-efficient and lets you use the Filmstrip (described below). You can switch to "Multiple Editor Tabs Mode" (METM) by going to "Preferences > General > Layout", however each Editor tab will require a specific amount of RAM relative to the image size and the tools you use, and also the Filmstrip is hidden in this mode, so we recommend you first give SETM a try.

The Preview Panel

The central panel holds a preview of your photo. This preview is generated from the actual raw data by processing it according to the settings either you manually set, or those that are stored in the processing profile used when opening that photo, as specified in "*Preferences > Image Processing > Default Processing Profile*". The preview will show you the effect of all the adjustments you make. Note that the effects of some tools are only accurately visible when you are zoomed in to 1:1 (100%) or more. These tools are marked in the interface with a "1:1" icon 1:1 next to the tool's name.

The image you see in the preview is taken from the working profile's color space and converted into the monitor profile's color space, if a monitor profile is loaded, or into sRGB if one is not. It does not take into account the "Output Profile" section of the "Color Management" tool.

Eek! My Raw Photo Looks Different than the Camera JPEG

After opening a raw photo you will notice that it looks different, often worse - darker, less sharp, more dull, lacking contrast, more noisy - than your camera's JPEG, or than the same raw photo when viewed in other software. What gives? Witches, aliens, possums, or by design?

There are three things you must know first to understand what is happening here:

- 1. Your camera does not show you the real raw data when you shoot raw photos. It processes the raw image in many ways before presenting you with the histogram and the preview on your camera's digital display. Even if you set all the processing features your camera's firmware allows you to tweak to their neutral, "o" positions, what you see is still not an unprocessed image. Exactly what gets applied depends on the choices your camera's engineers and management made, but usually this includes a custom tone curve, saturation boost, sharpening and noise reduction. Some cameras, particularly low-end ones and Micro Four-Thirds system (http://en.wikipedia.org/wiki/Micro Four Thirds_system), may also apply lens distortion correction to not only fix barrel and pincushion distortion (http://en.wikipedia.org/wiki/Distortion (optics)#Radial distortion) but also to hide severe vignetting (http://en.wikipedia.org/wiki/Vignetting) problems. Most cameras also underexpose every photo you take by anywhere from -0.3EV to even -1.3EV or more, in order to gain headroom in the highlights. When your camera (or other software) processes the raw file it increases exposure compensation by the same amount, making the brightness appear correct and hoping to recover some highlights in the process. RawTherapee shows you the real raw data which may mean your photos appear dark, so it is up to you whether you apply the required exposure increase and how you go about doing so, whether by using the Exposure Compensation slider or one of the various tone curves. Increasing exposure compensation makes noise more apparent regardless whether it is your camera or RawTherapee which does it, but other than this [b]RawTherapee does not "add noise"![/b] Many cameras apply noise reduction to the JPEGs (behind your back) to lower the noise level after increasing the exposure compensation, so you should expect there to be a difference between your out-of-camera JPEG and RawTherapee's image if noise reduction in RawTherapee is not enabled.
- 2. Every DSLR raw file contains a processed JPEG image. Most raw files contain a JPEG image of the same full resolution as your camera can shoot, and some raw files contain as many as three JPEG images differing only in resolution. When you open raw files in other software, what you are usually seeing is **not** the raw data, but the embedded, processed JPEG image. Examples of software which are either incapable of or which in their default settings do not show you the real raw data: IrfanView (htt p://en.wikipedia.org/wiki/IrfanView), XnView (http://en.wikipedia.org/wiki/XnView), Gwenview (http://en.wikipedia.org/wiki/Gwenview), Geeqie (http://en.wikipedia.org/wiki/Geeqie), Eye of GNOME (http://en.wikipedia.org/wiki/Eye_of_GNOME), F-Spot (http://en.wikipedia.org/wiki/F-S pot), Shotwell (http://en.wikipedia.org/wiki/Shotwell_(software)), gThumb (http://en.wikipedia.org/wiki/GThumb), etc. It is worth mentioning at this point that if you shoot in "RAW+JPEG" mode, you are in fact wasting disk space and gaining nothing for it, as your raw files already contain the embedded JPEG files which you can view using the listed programs. The embedded JPEG may differ from an 'external' one as saved using "RAW+JPEG" mode in compression.
- 3. Most raw development programs (programs which do read the real raw data instead of just reading the embedded JPEG) apply some processing to it, such as a base tone curve, even at their most neutral settings, thereby making it impossible for users to see the real, untouched contents of their raw photos. Adobe Lightroom is an example. Comparing RawTherapee's real neutral image to a quasi-neutral one from these other programs will expose the differences.

RawTherapee, on the other hand, is designed to show you the real raw image in the main preview, leaving the way you want this data processed up to you. When you use the "Neutral" processing profile you will see the demosaiced image with camera white balance in your working color space with no other modifications. You can even see the non-demosaiced image by setting the demosaicing option to "None". To provide you with a more aesthetically pleasing starting point, we do ship a collection of processing

profiles with RawTherapee. After installing RawTherapee, the default profile for processing raw photos is eponymously called "Default". We also ship the "Default ISO Medium" and "Default ISO High" profiles which are designed to give a good starting point to moderately noisy and very noisy images, respectively.

None of the shipped profiles (at least none of the ones shipped in RawTherapee 5.0) are designed to imitate your camera's look. Why not? Every camera is different. My camera's image quality at ISO1600 could be far noisier than your camera's. My camera's response to colors differs from yours. Even the same camera can behave differently at various settings. To provide such profiles, we would need access to raw files for every supported camera model, often multiple raw files in various shooting modes for a single camera, and countless person-hours (http://en.wikipedia.org/wiki/Man-hour). This may be possible as a community effort, but it is not a job for a small team. Even then, of what purpose would RawTherapee be if you ended up with a camera JPEG look?

It is far more reasonable that you learn how to use the powerful tools that RawTherapee provides to get the most out of your raws, to surpass the camera look.

As of September 2015 we are starting to ship DCP input profiles made using DCamProf (http://www.ludd.ltu.se/~torger/dcamprof.html) which include an optional tone curve (http://www.ludd.ltu.se/~torger/dcamprof.html#dcp_ton e). This curve is modeled after Adobe Camera Raw's default film curve and renders a result similar to your "camera look". The reason we include the curve in new DCP profiles is because it makes for a good vibrant starting point (as opposed to the flat look of using the "Neutral" processing profile) without having to use Auto Levels and without having to touch any of the other tools, and it is entirely optional. Do read the article on input profiles. If we ship a DCP for your camera model which includes the tone curve, the "Tone curve" checkbox in Color Management > Input Profile > DCP will be clickable. Applying the (Neutral) processing profile will disable the tone curve. While the input color profile is applied at the first stages of the toolchain pipeline, the DCP tone curve is applied later in the pipeline at some point after the Exposure tool.

You can create a processing profile ideally tailored to your camera and lens combination, and set RawTherapee to use it by default on your raw photos. See the Creating processing profiles for general use article to learn how.

Scrollable Toolbars

The toolbars above and below the main preview hold a certain number of buttons and other widgets which might not fit on lower resolution screens. If your screen's resolution is too low to fit the whole toolbar, some of the toolbar's contents (buttons, drop-downs, etc.) may become hidden. To see them, simply hover the cursor over the toolbar and use the mouse scroll-wheel to scroll the contents left and right.

Preview Modes

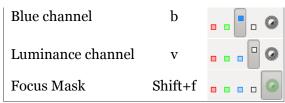
In addition to the normal preview, RawTherapee supports a number of other preview modes to help you tweak your photos. Preview modes are controlled via buttons in the *Editor* toolbar or via keyboard shortcuts. Only one preview mode can be engaged at a time.

The following preview modes are currently supported:

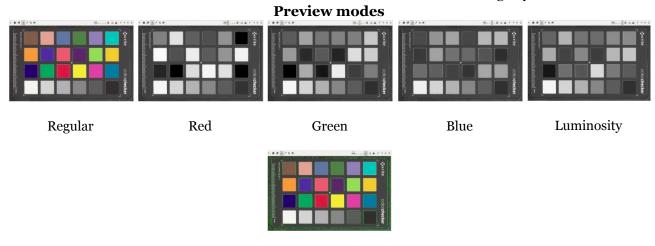
- Red channel.
- Green channel,
- Blue channel,
- Luminosity, which is calculated as 0.299*R + 0.587*G + 0.114*B,

Preview Mode	Button					
Regular*						
Red channel	r					•
Green channel	g					•

• Focus mask, to see which areas are in focus



* The preview is returned to normal by deselecting any other mode.



Focus Mask

Red, Green, Blue and Luminosity Preview Modes

The clipped shadow \triangle and \triangle highlight indicators in the Editor allow you to easily see which areas of the image are too dark or too bright. Highlighted areas are shaded according to the much they transgress the thresholds.

The thresholds for these indicators are defined in Preferences > General.

The clipped shadow indicator will highlight areas where all three channels fall at or below the specified shadow threshold. The clipped highlight indicator will highlight areas where at least one channel lies at or above the specified highlight threshold.

Clipping is calculated using data which depends on the state of the gamut button \triangleright which you can toggle above the main preview in the Editor tab. When the gamut button is enabled the working profile is used, otherwise the gamma-corrected output profile is used.

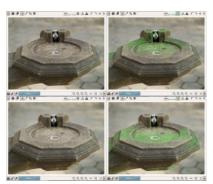
Preview of individual channels may be helpful when editing RGB curves, planning black/white conversion using the channel mixer, evaluating image noise, etc. Luminosity preview is helpful to instantly view the image in black and white without altering development parameters, to see which channel might be clipping or for aesthetic reasons.

Focus Mask

The focus mask is designed to highlight areas of the image which are in focus. Naturally, focused areas are sharper, so the sharp areas are being highlighted. The focus mask is more accurate on images with a shallow depth of field, low noise and at higher zoom levels. To improve detection accuracy for noisy images

evaluate at smaller zoom, around the 10-30% range. Note that the preview is rendered more slowly when the focus mask is enabled.

The current implementation analyzes the preview image which is rescaled from the original captured size. This process of rescaling reduces the noise and is helpful to identify truly sharper details rather than noise itself which may also contain micro texture. At the same time, rescaling of the original image to the preview size compresses larger scale details into a smaller size, and it may introduce aliasing artifacts, both of which could lead to false positives. You can increase your confidence by viewing the mask at various zoom levels. It is not always fault proof, but can be helpful in many cases.



Focus mask indicating the focusing plane

Warning: Be sure to double-check your images if you decide to delete them based on the focus mask.

Background color of the preview

The background color of the preview panel surrounding the image area may be changed to ease image preview during editing and to better visualize image cropping. A vertical stack of three thin buttons in the preview modes toolbar above the image preview panel allows to set the background color of the area around the photo preview.

Preview Background	Shortcut	Button	Preview Background and Crop Visualization	Description
Theme-based	8			The cropped area of the image is masked with a theme-based color. The cropped area visibility is based on the crop mask color and transparency as set in "Preferences > Default Theme > Crop mask color/transparency".
Black	9			The cropped area of the image is masked with black.
White	0			The cropped area of the image is masked with white.

Detail Window

The "New detail window" button [1], situated below the main preview next to the zoom buttons, opens a new viewport over the main preview of an adjustable size and of adjustable zoom. This lets you work on the photo zoomed-to-fit while examining several areas of interest at a 100% zoom (or even more). The benefit of using this feature is particularly important to users with slower machines, though not only them, as the zoomed-out main preview takes a shorter amount of time to update than if you were to zoom it to 100% because working at a zoom level less than 100% excludes certain slow tools, such as Noise Reduction, while the little detail windows zoomed to 100% do include all tools and are fast to update because of their small size. This allows you can use the main preview for your general exposure tweaks where it is necessary to see the whole image, and one or more detail windows to get sharpening and/or noise reduction just right.

Preview refresh delay

Changing any tool's parameters sends a signal for the preview image to be updated accordingly. Imagine what would happen if there was no "delay period", and you dragged, for example, the exposure compensation slider from 0.00 to +0.60. A signal would be sent to update the preview for every single change of that value - for +0.01, +0.02, ... +0.59, +0.60. Updating the preview 60 times would be completely unnecessary and actually take longer than it takes you to move the slider. This is especially true for more complicated tools, such as noise reduction, where a preview update can take even a second (depending on your CPU and preview size). The solution is for RawTherapee to wait for a very short period from the moment you stop moving a slider (you don't have to let go of it, pausing movement is enough) until the moment it sends a signal for the preview to be refreshed.

We have introduced two parameters which control the length of this waiting period:

AdjusterMinDelay

Default value = 100ms.

This is used for tools with a very fast response time, for example the exposure compensation slider.

AdjusterMaxDelay

Default value = 200ms.

This is used for tools with a slow response time, for example the CIECAMo2 sliders.

You can adjust both of these values in the options file in the config folder.

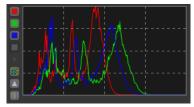
The Left Panel

To the left is a panel which optionally shows the main histogram ("Preferences > General > Layout > Histogram in left panel"), and always shows the Navigator, History and Snapshots. You can hide this panel using the + hide icon, or its keyboard shortcut.

Main Histogram

The main histogram can show the histograms of the red , green , blue , CIELab Luminance and Chromaticity (http://en.wikipedia.org/wiki/Chromaticity) channels of the photo as it would look if you saved it. Regardless of the photo's bit depth, the histogram itself has a precision of 256 points.

The values the main histogram and Navigator panel show are those of the working profile when the gamut button \triangleright is enabled, or those of the gamma-corrected output profile when the button is disabled. The clipped shadow \blacktriangle and \blacktriangle highlight indicators depend on this gamut setting.



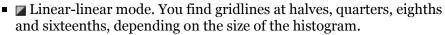
Histogram of the input raw image. We can see there is no clipping in this image and that it is underexposed by a stop.

It can also show the histogram of the raw data as it is in the raw file, before any transformations such as demosaicing are applied to it. Use this information to see whether the raw image is exposed correctly and there it has any clipping. Clipped shadows and highlights cannot be recovered - if it is clipped, the data is simply not there. However, some clipped highlights can be reconstructed using the Color Propagation method.

You can show or hide the RGB Indicator Bar , which is situated under the histogram and shows you the exact place on the histogram of the R, G, B and L values of the pixel your cursor is currently hovering over in the main preview.

The histogram can be moved to the left/right panel from "*Preferences* > *General* > *Layout* > *Histogram in left panel*".

To help you visualize the data, the histogram (as of RawTherapee 5.5) has three modes which scale the data in the x and y axes differently:



- Linear-log mode. The x-axis is linear, the y-axis and the horizontal gridlines are scaled logarithmically. The position of the gridlines still corresponds to the halves, quarters, etc.
- Log-log mode. Both the x- and y-axes are scaled logarithmically. The gridlines are not scaled logarithmically, but correspond to stops (https://en.wikipedia.org/wiki/Exposure_value) with every gridline the value doubles, so there are lines for the values 1, 3, 7, 15, 31, 63, and 127 (pow (2.0, i) 1)).

When there is a disproportionately bright area relative to the rest of the image, it will show up as a spike in the histogram. If you want to show this on a histogram with a linear y axis, the spike may push the lesser values down the y-axis, making them difficult to see. Switch to one of the log modes to scale the data and help you get a better overview of all values.

Navigator

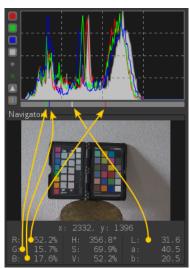
The *Navigator* panel shows a thumbnail of the currently opened image, and RGB, HSV and Lab values of the pixel your cursor is currently hovering over.

The values the main histogram and Navigator panel shows are either those of the working profile, or of the gamma-corrected output profile. You can choose which you prefer in "*Preferences > General > Use working profile for main histogram and Navigator*".

By clicking on the values in the Navigator you can cycle between these three formats:

- **■** [0-255]
- **■** [0-1]
- **[**%]

RawTherapee 5.1 onward can show the real raw photosite values. To see them, set the Navigator to use the [0-255] range, apply the Neutral processing profile, then set the Demosaicing method to "None". The Navigator will show the real raw photosite values after black level subtraction within the range of the original raw data.



The RGB Indicator points out the position in the histogram of the R, G, B and L values of the pixel your cursor is hovering over.

History

Under the *Navigator* it is the *History* panel. While editing a photo, all your actions are recorded in this *History* panel. By clicking on the different entries, you can step back and forth through the different stages of your work.

Snapshots

Under the *History* panel is a panel called *Snapshots*. Its use is in that you can save a snapshot of the photo with all the adjustments up to that point in time, and then proceed to further modify your photo to give it a different appearance, saving new snapshots at every moment you feel you might have reached a version of your photo worth saving. Once you have two or more snapshots, you can just click on them to flip through the different versions and stick with whichever one you like best. In the future, the snapshots will be saved to the PP3 sidecar file. For now, the history and snapshots are lost when you load a new photo in the *Image Editor* or close RawTherapee.

The Right Panel

To the right is a panel which optionally shows the main histogram and *Processing Profiles* selector ("*Preferences > General > Layout > Histogram in left panel*"), and always shows the Toolbox. You can hide this panel using the \vdash hide icon, or its keyboard shortcut.

Processing Profile Selector

The *Processing Profiles* drop-down list lets you apply bundled or custom processing profiles. See the File Paths article for information on where these processing profiles reside on your system.

Pay attention to the "Processing profile fill mode" button!

"Fill" mode **!**≣

When the button is activated and you open a partial profile, the missing values will be replaced with RawTherapee's hard-coded default values.

For instance if you apply a partial profile which contains only sharpening settings, all of the remaining tools (such as Exposure, Tone Mapping, Noise Reduction, Resize, etc) will pop into their default positions.

"Preserve" mode =

If the button is deactivated and you open a partial profile, only those values in the profile will be applied, and the missing ones remain unchanged.

For instance if you apply a partial profile which contains only sharpening settings, only those sharpening settings will be applied, and your other tools remain unchanged.

The state of this button will make no difference if you apply a full profile, but most of the profiles bundled with RawTherapee are partial (for good reason).

Toolbox

The *Toolbox*, in the right panel, contains all the tools you use to tweak your photos. Each tool has its own RawPedia article.

Editor Tab Modes

RawTherapee allows you to work on photos in two modes:

- Single Editor Tab Mode (SETM), where you work only on one photo at a time, and each photo is opened in the same Editor tab. There is a horizontal panel called the Filmstrip at the top of the Editor tab showing the rest of the photos in that folder for easy access. There are Previous Image and Next Image � b buttons in the bottom toolbar (and keyboard shortcuts for them) to switch to the previous/next image.
- *Multiple Editor Tabs Mode* (METM), where each photo is opened in its own *Editor* tab. The *Filmstrip* is hidden in this mode and there are no previous/next buttons. Having multiple photos opened at the same time requires more RAM.

Try both modes and see which one suits you best. To do that, click on the *Preferences* icon in the bottom-left or top-right corner of the RT window, choose "*General* > *Layout*" and set *Editor Layout* to your preferred choice.

Use this *Preferences* window to select a different language for the user interface, to choose a different color theme, change the font size, etc.

It is also possible to start RawTherapee in no-File-Browser-mode (without the *File Browser* tab) by specifying RawTherapee to open an image from your operating system's file browser (in other words, right-click on a photo and select "*Open With > RawTherapee*"), or by using the image filename as an argument when starting RawTherapee from the command line (rawtherapee)

/path/to/some/photo.raw). This mode was introduced for people with little RAM as not having a *File Browser* tab means RawTherapee uses a little less memory, however in practice the amount of memory saved is little and the usability cost outweighs the little benefit, so it is likely to be removed in the future (see issue 2254 (https://code.google.com/p/rawtherapee/issues/detail?id=2254)).

The Filmstrip

If you use



RawTherapee-4.2 showing the Filmstrip with the toolbar visible, which takes up more screen space but lets you easily label, rate and filter the visible thumbnails.



RawTherapee-4.2 showing the Filmstrip with the toolbar hidden, which makes it less high and provides more screen space for the main preview (partially visible at the bottom).

Single Editor Tab Mode ("Preferences > General > Layout") you can display a horizontal panel above the preview, this is called the *Filmstrip*. It contains thumbnails of all images in the currently opened album, and is synchronized with the currently opened image so that you can use keyboard shortcuts or the previous \triangleleft and next \triangleright image buttons to open the previous/next image without needing to go back to the *File Browser* tab.

As of RawTherapee version 4.2.10, you can hide the Filmstrip's toolbar to save screen space. There are two ways of doing this: one way just toggles the toolbar on/off without resizing the filmstrip to the new height, and the other way does the same but also automatically resizes the filmstrip's height. Both are invoked via keyboard shortcuts only. As resizing the filmstrip's height will trigger a refresh of the image preview and this might take a while if using CPU-hungry tools like noise reduction while zoomed in at 100%, the mode that doesn't resize has been implemented for users with slow machines. Users with fast machines will find the auto-resizing mode more helpful.

Monitor Profile and Soft-Proofing

The widgets under the main preview in RawTherapee 5 allow you to apply a monitor color profile to the preview image. This enables users who have calibrated and profiled their monitors to get an instant and accurate preview of their work, whether you're staying in sRGB or working in a wide gamut. Note: OS X users are limited to sRGB and will not get an accurate preview otherwise (see discussion (https://discuss.pixls.us/t/wide-gamut-preview-in-os-x/2481)), while users of Linux and Windows will get a correct wide-gamut preview.

Go to Preferences > Color Management and point the "Directory containing color profiles" to the folder into which you saved your monitor and printer ICC profile. Restart RawTherapee for the changes to take effect. Now you will be able to select your monitor's color profile in the combo-box under the preview. Use the "Relative Colorimetric" rendering intent unless you have a good reason otherwise.

One can also enable soft-proofing of the preview. This will show you what your image will look like once it gets transformed by the printer profile set in Preferences > Color Management. If you want to adjust an image for printing and you have an ICC profile for your printer-paper combination you could set that as your output profile, enable "Black point compensation" in Preferences so that the blackest black in your image will match the blackest black your printer-paper combination is capable of reproducing, then enable soft-proofing. You will see what your image will look like if you print it. This allows you to make adjustments and get an instant preview of the result, saving you time and ink on test prints.

The icon with exclamation mark next to the soft-proofing button will gray out areas that cannot be reproduced by your printer, i.e. areas where you will loose details.

You should have a calibrated and profiled monitor in order for the soft-proofing preview to be accurate.

The items you see in the monitor profile combo-box (under the main preview) and in the printer profile combobox (in Preferences > Color Management) are ICC files located in a folder which you can point RawTherapee to by going to "Preferences > Color Management > Directory containing color profiles".

Batch Adjustments - Sync

RawTherapee lets you batch-adjust, or sync, the processing settings in many photos at the same time in generally two ways. It lets you copy and paste a processing profile (a collection of tool settings), in parts or in full, to any number of images. It also lets you select any number of images and adjust any tool in all of them at once (sync), and it lets you do this in two ways. Let's take a closer look.

Both ways involve making a selection of photos you want the processing profile or adjustments applied to. Selections are made using standard key combinations: Shift+click to select a range, Ctrl+click to select individual images, or Ctrl+A to select everything. Both ways are performed from the File Browser tab. The "copy & paste" method can also be done via the Filmstrip.

Copy & Paste

Copying and pasting a processing profile to a selection of images is a very common task. Assume you took a series of photos - for example studio shots, wedding portraits or focus-bracketed macro photos. All images in each series are going to be very similar; they will probably use the same lens, the same ISO, the same white balance, and end up being used for the same purpose. This means that they will all probably require the same processing settings - the same noise reduction, the same sharpening and lens distortion correction, and so forth.

To process the lot, what you would usually do is open any one image from the whole series in the Editor tab and tweak it to your liking. Once you have finished tweaking it, you will apply this image's processing profile to all other images in the same series. To do that, go to the File Browser tab, right-click on this photo and select "*Processing Profile Operations > Copy*", then select the images you want to apply this profile to, right-click on any one of them (it doesn't matter which) and select "*Processing Profile Operations > Paste*". In one quick operation you have replicated the same tool settings in the whole series of images.

Additionally, RawTherapee lets you apply only a part of the copied processing profile, for example only the "Resizing" tool. To do this, use the "*Processing Profile Operations > Paste Partial*" option instead of the "Paste" option.

Sync

RawTherapee lets you instantly apply tool adjustments to a selection of images. Similar functionality in other software is called "sync". This method is useful for when you don't need to see an accurate preview of your changes, for example when you only want to enable the "Resizing" tool in a selection of photos, because when working in the File Browser tab your only preview are the small and inaccurate thumbnails. This method can only be performed from the File Browser tab because you need access to that tab's batch tools (the panel on the right).

When you're in the File Browser tab, select the images you want to batch-adjust (sync), then use the tool panel on the right to make adjustments. Your tweaks can either replace the existing ones ("Set" mode), or be added to them ("Add" mode). For example if you select two photos, one of which has previously been tweaked with +1EV Exposure Compensation and one which has not, and you set Exposure Compensation

to +0.6EV, then the previously-tweaked photo would end up having +1.6EV Exposure Compensation in "Add" mode and just +0.6EV in "Set" mode. The photo which was not previously tweaked would have +0.6EV in both modes. You can decide which tools should work in which mode from the Batch Processing tab in Preferences.

Saving

Your original raw file will never be altered by RawTherapee.

There are several ways of saving an image from the Image Editor tab:

- Save immediately from the Editor tab,
- Mar Save via the Batch Queue,
- \\ "Edit Current Image in External Editor" (described in its own article).

Save Immediately

In the Editor tab, if you click on the little hard disk icon at the bottom-left of the preview image, or hit the Ctrl+s shortcut, you can "Save immediately". This works as a standard "Save As" dialog. You can select the name and location for the output file (RawTherapee will automatically add the extension based on the chosen format), choose the output file format and bit depth, set the compression level, choose whether you want the processing profile saved alongside the output image, etc. The last option lets you choose whether you want to "Save immediately" or "Put to the head/tail of the processing queue". If you choose to "Save immediately", RawTherapee will be busy saving your photo as soon as you click "OK", so it will be less responsive to any adjustments you might try doing while it's busy saving, and it will also take longer to open other images as long as it's busy saving this one. For this reason it is recommended that you use the queue if you want to tweak other photos right away.

The Save window will by default open the location you saved to the last time you used this window. For your convenience, a shortcut to the folder containing the source image is automatically added to the bookmarks panel on the left side of the Save window. If you want to save the image to the source folder, clicking on this bookmark will save you having to manually click through to navigate to that location.

A shortcut for the *OK* button is **Ctrl+Enter**.

Put to the Head / Tail of the Processing Queue

If you click on the gears icon or in the "Save" window choose "Put to head or tail of the processing queue", your image will be kept in a queue of files to be processed, so RawTherapee can make the most of your CPU and be responsive while you tweak your photos. Once you're done tweaking and adding them to the queue, you can have RawTherapee start processing the queue while you go and enjoy some tea. The benefit of putting it to the queue using the "Save" window is that you can individually change the file format, name and destination of each image, whereas putting images to the queue without using the "Save" window will use the settings from the "Batch Queue" tab.

Naming

If your original raw file was called "photo_1000.raw", the default processed file name will be "photo_1000.jpg" (or .tif or .png). There is an option in the "Save current image" window: "Automatically add a suffix if the file already exists". When checked, you can make different versions of one raw, which will be saved as "photo_1000.jpg", "photo_1000-1.jpg", "photo_1000-2.jpg", etc. The same applies when you send different versions of the same image to the batch queue.

The Batch Queue

Introduction

Saving images from RawTherapee can be done in several ways, the two most common of which are either saving the image immediately if from the Editor tab, or adding it to the batch processing queue which resides in the Queue tab.

Using the "Save immediately" feature will put your CPU immediately to work, and as a result, opening and tweaking other images in the Editor will be somewhat slow while the image is being saved. The queue mechanism allows you to put edited images which are ready to be saved to a virtual queue which you can start processing at a later time. Adding them to the queue is instant, so you can continue editing other images and making the most of your CPU for editing. Once you are done editing and putting images to the queue, you can flip the large "On" switch and go off to brew yourself a coffee while RawTherapee grinds away at all the images in the queue.

The queue is persistent - you can exit RawTherapee and restart it later; the queued images will still be there. The queue can even survive a crash.

Adding Images to the Queue

There are several ways of adding an image to the queue:

- 1. When you are done tweaking an image in the Editor, click the "Put current image to processing queue" button 🔗.
- 2. Also in the Editor tab, click the "Save current image" button 📑 and select "Put the head/tail of the processing queue".
- 3. Right-click on a thumbnail in the File Browser or Filmstrip and select "Put to queue".

Regardless which method you use, when you go to the *Queue* tab you will see your photos lined up, ready for processing (if you had the queue set to "Auto-start", it may have finished processing before you viewed it).

Queue Settings

The Queue has several settings, such as the output file format and destination. These settings take effect in all cases except when you use the "Save current image" button , select "Put the the head/tail of the processing queue" and enable the "Force saving options" checkbox. In this case, the settings seen in the "Save" window will be used, and the ones from the Queue tab ignored. In all other cases, the settings from the Queue tab will be used.

The settings speak for themselves. Two things worth pointing out:

- 1. "Save processing parameters with image" will save a sidecar file alongside the output file, with the same filename as the output image but with a ".pp3" extension. This is useful when you want to save multiple copies of the same photo, each one tweaked a bit differently.
- 2. The destination folder can be set by selecting "Save to folder", but if you need to dynamically customize the destination folder and filename then select "Use template" instead. Hover your mouse over the *Use template* input box and a tooltip with an explanation will pop up:

```
You can use the following formatting strings:
%f, %dl, %d2, ..., %pl, %p2, ..., %r, %s1, %s2, ...
These formatting strings refer to the different parts of the photo's
For example, if the photo being processed has the following pathname:
//home/tom/photos/2010-10-31/dsc0042.nef
the meaning of the formatting strings are:
%d4 = home
%d3 = tom
%d2 = photos
%d1 = 2010-10-31
%f = dsc0042
%p1 = /home/tom/photos/2010-10-31/
%p2 = /home/tom/photos/
%p3 = /home/tom/
%p4 = /home/
%r will be replaced by the rank of the photo. If the photo is unranke
%s1, %s2, etc. will be replaced by a sequence index which is padded t
If you want to save the output image where the original is, write:
%p1/%f
If you want to save the output image in a directory named "converted"
%p1/converted/%d1/%f

If you want to save the output image in a directory named "/home/tom/
%p2/converted/%d1/%f
```


The "Save current image" window. Notice the "Force saving options" checkbox, which is disabled in the screenshot because neither of the "Put to the head/tail of the processing queue" options are enabled.

Running the Queue

In the top-left corner of the Queue tab you will find an "On/Off" switch, and an "Auto-start" checkbox.

- 1. If "Auto-start" is enabled, processing will start as soon as an image is sent to the queue. Usually you will not want this, as this will use your CPU for developing the photos in the queue leaving very little CPU for tweaking other images.
- 2. If "Auto-start" is unchecked, you will have to activate the queue manually by hitting the "On/Off" switch.

You can pause the queue by hitting the "On/Off" switch - RawTherapee will first finish processing the current photo.

Clearing the Queue

You can remove a specific image from the queue by clicking the small "Cancel job" **x** button in the corner of each thumbnail.

You can clear the whole queue right-clicking on a thumbnail and clicking "Select all" and "Cancel job", or by using the ^Ctrl + a keyboard shortcut to select all thumbnails and then hitting the Delete key on the keyboard.

Edit Current Image in External Editor

The "Edit current image in external editor" feature allows you to have RawTherapee fully process the current image and immediately open it in any external program. You can use this feature to easily send the image to an image editor such as GIMP or Photoshop for further processing, or you can set it up to send the image to an image viewer so that with a single click of a button you get to see the final high quality image.

The button to send the image to an external application \infty is located at the bottom-left of the preview panel. When using this feature, RawTherapee processes your image and saves it as a gamma-encoded 16-bit integer TIFF to the temporary folder. These intermediate files, due to being outside of RawTherapee's control, do not get automatically deleted when you close RawTherapee, so you should keep this in mind and clean them out manually.

You should be aware that GIMP-2.8 and below cannot handle 16-bit images, so it will down-sample them to 8-bit. You should also be aware that GIMP-2.8 and below discards all Exif data from TIFF files! This is a GIMP bug, not a RawTherapee one. GIMP 2.9 and above handle high bit depth images (up to 64-bit per channel!), retain all metadata, and are quite stable, so you are advised to get GIMP 2.9 or higher (http://www.gimp.org/downloads/).

You can specify your external editor of choice in "*Preferences > General > External Editor*". Click the link for more information on specifying an external editor.

Command-Line Options

Explanation

<Chevrons> indicate parameters you can change.

[Square brackets] mean the parameter is not mandatory.

The pipe symbol | indicates a choice of one or the other.

The dash symbol – denotes a range of possible values from one to the other.

Since RawTherapee 5.1, two executables are provided.

RawTherapee GUI

Use this application to start the version with graphical user interface.

Usage:

```
rawtherapee <selected dir>
Start File Browser inside folder.
rawtherapee <file>
Start Image Editor with file.
```

-₩

Do not open the Windows console. This option is available in Windows only. If you pass parameters to the RawTherapee executable it spawns a console window so that you can see the verbose output of your processing. Normally Windows closes this console directly after RawTherapee is terminated. To let you see the output we added a prompt which waits for you to hit a key before closing the console. By specifying <code>-w</code> no console will be opened and therefore no key press is needed. Useful if you want to invoke <code>rawtherapee.exe</code> in batch, e.g. from a PowerShell script. Please note that <code>-w</code> will have no effect for "Debug" builds where a console window will be opened unless you're starting RawTherapee from a console window already.

 $-\nabla$

Print the RawTherapee version number and exit.

-R

"Remote" mode, available since RawTherapee 5.2. When opening an image using "Open with" or by passing its filename as an argument, without using the ¬R option RawTherapee will open in "no-File-Browser" mode - that is a mode which lacks the File Browser and Queue tabs as well as the Preferences button. Using the new ¬R mode, RawTherapee will open in a full-fledged instance. Using ¬R also allows you to open an image in an already-running instance of RawTherapee, if that instance was also started using ¬R. The no-File-Browser mode exists for historical reasons when RAM requirements were higher and stability was worse. Now that RawTherapee's memory usage is optimized and it can quickly and reliably open folders with thousands of images, users may prefer using the ¬R mode by default.

-h -?

Display these commands.

RawTherapee CLI

Use this application to start the command line only version. You'll find all command line options to develop your photos without any graphical user interface.

Usage:

```
rawtherapee-cli <options> -c <dir>|<files>
```

Convert files in batch with default parameters if no <options> specified.

-w

Do not open the Windows console. This option is available in Windows only. If you pass parameters to the RawTherapee executable it spawns a console window so that you can see the verbose output of your processing. Normally Windows closes this console directly after RawTherapee is terminated. To let you see the output we added a prompt which waits for you to hit a key before closing the console. By specifying <code>-w</code> no console will be opened and therefore no key press is needed. Useful if you want to invoke <code>rawtherapee-cli.exe</code> in batch, e.g. from a PowerShell script.

Other options used with -c:

```
rawtherapee-cli [-o <output>|-O <output>] [-q] [-a] [-s|-S] [-p <files>] [-d] [-j[1-100] [-js<1-3>]|[-b<8|16>] <[-t[z] | [-n]]] [-Y] [-f] -c <input> -c <files>
```

Specify one or more input files or folders.

When specifying folders, RawTherapee will look for image files which comply with the selected parsed extensions (see the -a option).

The -c option must always be the last one.

```
-o <file>|<dir>
```

Select output file or folder.

Saves output file alongside input file if -o is not specified.

```
-O <file>|<dir>
```

Select output file or folder and copy PP3 file into it. Saves output file alongside input file if -O is not specified.

```
-q<file>|<dir>
```

Quick-start mode. Does not load cached files to speedup start time.

```
-a<file>|<dir>
```

Process all supported image file types when specifying a folder, even those not currently selected in Preferences > File Browser > Parsed Extensions.

-s

Use the existing sidecar file to build the processing parameters, e.g. for photo.raw there should be a photo.raw.pp3 file in the same folder. If the sidecar file does not exist, neutral values will be used.

-S

Like -s but skip if the sidecar file does not exist.

```
-p <file.pp3>
```

Specify processing profile to be used for all conversions. You can specify as many sets of "-p <file.pp3>" options as you like, each will be built on top of the previous one, as explained below.

-d

Use the default raw or non-raw PP3 file as set in "Preferences > Image Processing > Default Processing Profile"

```
-j[1-100]
```

Specify output to be JPEG (default, if -t and -n are not set). Optionally, specify compression 1-100 (default value: 92).

```
-js<1-3>
```

Specify the JPEG chroma subsampling (http://en.wikipedia.org/wiki/Chroma_subsampling) parameter, where:

1 = Best compression: 2x2, 1x1, 1x1 (4:2:0)

Chroma halved vertically and horizontally.

2 = Balanced: 2x1, 1x1, 1x1 (4:2:2)

Chroma halved horizontally.

3 = Best quality: 1x1, 1x1, 1x1 (4:4:4)

No chroma subsampling.

-b<8|16>

Specify bit depth per channel (16 by default). Only applies to TIFF and PNG output, JPEG is always 8.

-t[z]

Specify output to be TIFF (16-bit if -b8 is not set). Uncompressed by default, or ZIP compression with 'z'.

-n

Specify output to be compressed PNG (16-bit if -b8 is not set). Compression is hard-coded to level 6.

-Y

Overwrite output if present.

-f

Use the custom fast-export processing pipeline.

Your PP3 files can be incomplete, RawTherapee will set the values as follows:

- 1. A new processing profile is created using neutral values,
- 2. If the -d option is set, the values are overridden by those found in the default raw or non-raw processing profile,
- 3. If one or more -p options are set, the values are overridden by those found in these processing profiles,
- 4. If the -s or -S options are set, the values are finally overridden by those found in the sidecar files.

The processing profiles are processed in the order specified on the command line.

Redirect Output

To redirect RawTherapee's output to a text file, you have to start it from a console and append the redirection code as follows:

Windows (cmd.exe)

```
rawtherapee.exe > rtlog.txt 2>&1
Linux
rawtherapee &> rtlog.txt
```

Examples

Example 1

In Linux, process a single raw which resides in /tmp and is called "photo.raw", use its sidecar file "photo.raw.pp3" during conversion, save it in the same folder as "foo.tif", and overwrite the file "foo.tif" if it exists:

```
rawtherapee-cli -o /tmp/foo.tif -s -t -Y -c /tmp/photo.raw
```

Example 2

In the next example, we'll assume that you want to quickly process all your raw photos from the /tmp/janeo1 folder to a web sub-folder by using the default profile as a basis, using the sidecar profile if it exist, but with removing some Exif tags (e.g. the camera's serial number) and adding some IPTC tags (e.g. your usual copyright parameters), plus resize and sharpen the image for the web (spread over multiple lines for clarity):

```
rawtherapee-cli -o /tmp/Jane01/web -p ~/profiles/iptc.pp3 -s -p ~/profiles/exif.pp3 -p ~/profiles/web.pp3 -t -Y -d -c /tmp/Jane01/
```

The processing profile will be built as follows:

- 1. A new profile is created using internal default values (hard-coded into RawTherapee),
- 2. then overridden by those from the default raw profile (-d),

- 3. then overridden by those found in iptc.pp3,
- 4. then overridden by those found in the sidecar file (-s) if it exists, so you can force some IPTC tags even if already set by iptc.pp3,
- 5. then overridden by those found in exif.pp3, so you can force the profile to erase some tags,
- 6. then overridden by those found in web.pp3, to resize and sharpen the image, and make sure that the output colorspace is sRGB.

As you can see, the position of the -s switch tells when to load the sidecar profile relative to the other -p parameters. That is not the case for the -d switch.

Example 3

In the third example, we will see how long it takes to process every raw file in a folder, assuming that each raw photo has a corresponding processing profile, and discard each output file:

```
time {
    for f in /home/user/photos/2011-11-11/*.raw; do
    rawtherapee-cli -o /dev/null -S -t -Y -c "$f";
    done
    }
```

Keyboard Shortcuts

Note that some shortcuts, specifically ones that involve one or more modifier keys such as <u>^ Ctrl</u> + <u>î Shift</u>, are known to not work on some non-QWERTY keyboard layouts. This is not a fault of RawTherapee.

- SETM = Single Editor Tab Mode (required for Filmstrip panel visibility)
- METM = Multiple Editor Tabs Mode



For reference, this is the standard United States-NoAltGr QWERTY keyboard layout (http://en.wi kipedia.org/wiki/Keyboard_layout).

Cross-Platform Shortcuts

Where	Shortcut	Action	
Anywhere	^ Ctrl] + F2	Switch to the File Browser tab.	
Anywhere	^ Ctrl] + [F3]	Switch to the Queue tab.	
Anywhere	^ Ctrl] + F4	Switch to the Editor tab.	
Anywhere	$rac{\ \ \ \ \ \ \ \ \ \ \ \ \ \ \ \ \ \ \ $	Copy the processing profile.	
Anywhere	^ Ctrl + v or î Shift + v	Paste the processing profile.	
Anywhere	^ Ctrl + 1 Shift + v	Partial-paste the processing profile. Pay attention to the processing profile fill mode!	
Editor, File Browser	✓ Alt + e	Switch to Exposure tab.	
Editor, File Browser	∠ Alt + d	Switch to Detail tab.	
Editor, File Browser	∠ Alt + c	Switch to Color tab.	
Editor, File Browser	✓ Alt + t	Switch to Transform tab.	
Editor, File Browser	✓ Alt + r	Switch to Raw tab.	
Editor	∠ Alt + m	Switch to Metadata tab.	
Anywhere	F11	Toggle fullscreen.	
Anywhere	^ Ctrl + q	Quit.	
Editor	z	Zoom to 100%.	
Editor	f	Zoom to crop.	
Editor	✓ Alt + f	Zoom to fit whole image.	
Editor	+ or =	Zoom in.	
Editor	- or _	Zoom out.	
Editor	8	Background color of image preview: theme-based.	
Editor	9	Background color of image preview: black.	
Editor	0	Background color of image preview: white.	
Editor	<	Toggle clipped shadows indicator.	
Editor	>	Toggle clipped highlights indicator.	
Editor	î Shift + b	Before/After view.	
Editor	c	Crop tool. © Shift]+[left-click-drag] within the cropped area to move the entire crop border.	
Editor	^ Ctrl + e	Edit current image in external editor.	
Editor	^ Ctrl + s	Save current image.	
Save Window	^ Ctrl + = Enter	"OK" to save.	
P.3:4		Force saving current settings to the processing profile. This happens automatically each time you close the image or RawTherapee, but you can also force it to happen	
Editor	^ Ctrl] + 1 Shift] + s	immediately so you don't lose any work should the program crash.	
Editor	^ Ctrl + z	Undo.	
Editor	^ Ctrl + 1 Shift + z	Redo.	
Editor	h	Hand/cross tool (standard); use this to navigate around a zoomed image or to move the frame that defines where the detail window is.	
Editor	i	Toggle quick info (Exif info overlay).	
Editor	Right-click in the preview area	Return to default cursor (after having used the pipette for white balance or the straightening tool). If the RGB indicator bar is enabled, this (un)freezes its indicators.	
Editor	s	Select straight line (Rotate tool).	
Editor	✓ Alt + s	Add a snapshot.	
Editor	1 Shift + drag crop	Move crop border.	
Editor	w	Spot white balance pipette.	
231101	ے ا	oper mane summer piperter	

Where	Shortcut	Action	
Editor	r	Toggle preview mode - red.	
Editor	g	Toggle preview mode - green.	
Editor	b	Toggle preview mode - blue.	
Editor	v	Toggle preview mode - luminosity.	
Editor	Î Shift + f	Toggle preview mode - focus mask.	
Editor		Rotate open image left.	
Editor]	Rotate open image right.	
Editor	-	Decrease preview size.	
Editor	+	Increase preview size.	
File Browser		Rotate selected thumbnail left.	
File Browser]	Rotate selected thumbnail right.	
Filmstrip (SETM)	✓2 Alt + [Rotate selected thumbnail left.	
Filmstrip (SETM)	✓ Alt +]	Rotate selected thumbnail right.	
Filmstrip (SETM)	✓ Alt + -	Decrease thumbnail size.	
Filmstrip (SETM)	✓ Alt + +	Increase thumbnail size.	
File Browser	-	Decrease thumbnail size.	
File Browser	+	Increase thumbnail size.	
File Browser	ī	Toggle thumbnail information.	
Filmstrip (SETM)	✓2 Alt + i	Toggle thumbnail information.	
Editor (SETM)	F ₃	Navigate to the previous image relative to the thumbnail selected in the File Browser/ Filmstrip.	
Editor (SETM)	1 Shift + F3	Navigate to the previous image relative to the image open in the Editor.	
Editor (SETM)	F4	Navigate to the next image relative to the thumbnail selected in the File Browser/ Filmstrip.	
Editor (SETM)	î Shift + F4	Navigate to the next image relative to the image open in the Editor.	
		Synchronize the Filmstrip with the Editor to reveal and focus the thumbnail of the currently open image and	
Filmstrip (SETM)	[x]	clear active filters (rating, labels, metadata filters).	
Filmstrip (SETM)	У	Synchronize the Filmstrip with the Editor to reveal and focus the thumbnail of the currently open image, but without clearing filters in File Browser/ Filmstrip. (Note that the thumbnail of the open file will not be shown if filtered out).	
File Browser/ Filmstrip (SETM)	⊚ Esc	 Resets Browse Path box to current directory. Clears Find box. Does not defocus. 	
File Browser/ Filmstrip (SETM)	① Shift + ⊙ Esc	 Resets Browse Path box to current directory. Does not clear Find box. Defocuses. 	
File Browser/ Filmstrip (SETM)	^ Ctrl) + [0]	Sets focus to Browse Path box, selects contents. Browse Path shortcuts: Press	
File Browser/ Filmstrip (SETM)	^ Ctrl) + [f]	Sets focus to the Find box, selects contents. Press	

Where	Shortcut	Action	
		Rank selected thumbnails with 1-5 stars:	
File Browser/ Filmstrip (SETM)	1 Shift + 1 through 5	1. ** 2. *** 3. **** 4. **** 5. ****	
File Browser/ Filmstrip (SETM)	Î Shift + O	Un-rank selected thumbnails.	
		Apply a color label to the selected thumbnails:	
File Browser/ Filmstrip (SETM)	1 Shift + ^ Ctrl + 1 through 5	1. Red 2. Yellow 3. Green 4. Blue 5. Purple	
File Browser/ Filmstrip (SETM)	Ît Shift + ^ Ctrl + O	Remove color label from the selected thumbnails.	
El D		Toggle filter for images without a rating.	
File Browser/ Filmstrip (SETM)	0	^ Ctrl allows to apply an additive filter.	
File Browser/		Toggle filter for images with a specific rating.	
Filmstrip (SETM)	1 through 5	^ Ctrl allows to apply an additive filter.	
Eile Provenen/		Toggle filter for not edited images.	
File Browser/ Filmstrip (SETM)	6	^ Ctrl allows to apply an additive filter.	
File Browser/		Toggle filter for edited images.	
Filmstrip (SETM)	7	^ Ctrl allows to apply an additive filter.	
File Browser/		Toggle filter for images without a color label.	
Filmstrip (SETM)	<u>∠</u> Alt + 0	^ Ctrl allows to apply an additive filter.	
File Browser/		Toggle filter for images with a specific color label.	
Filmstrip (SETM)	<u> </u>	^ Ctrl allows to apply an additive filter.	
File Browser/		Toggle filter for images that were not saved.	
Filmstrip (SETM)	✓2 Alt + 6	^ Ctrl allows to apply an additive filter.	
File Browser/		Toggle filter for images that were saved.	
Filmstrip (SETM)	✓2 Alt + 7	<u>^ Ctrl</u> allows to apply an additive filter.	
File Browser/ Filmstrip (SETM)	d	Clear all filters: ratings, color labels, (not)edited, (not)saved, trash.	
File Browser/ Filmstrip (SETM)	F2	Rename file. If multiple thumbnails are selected, the rename dialog will be called for each.	
File Browser/ Filmstrip (SETM)	^ Ctrl + [t]	Show trash contents. Was just t prior version 4.2.10.	
File Browser, Editor	L	Toggle left panel visibility.	
File Browser, Editor	✓ Alt + L	Toggle right panel visibility.	
Filmstrip	Shit + L	Toggle Filmstrip visibility.	
Filmstrip	t	Toggle the visiblity of the Filmstrip's toolbar and automatically resize the height of the filmstrip.	
Filmstrip	1 Shift + t	Toggle the visiblity of the Filmstrip's toolbar but do not automatically resize the height of the filmstrip.	
Editor (SETM)	^ Ctrl + î Shift + L	Toggle left panel and Filmstrip visibility.	
Editor (SETM)	✓ Alt + î Shift + L	Toggle right panel and Filmstrip visibility.	
File Browser, Editor	^ Ctrl + ✓ Alt + L	Toggle left and right panel visibility.	
Editor	$ \begin{bmatrix} $	Toggle left panel, right panel and Filmstrip visibility.	

Where	Shortcut	Action	
File Browser, Editor	m	Toggle all panels' visibility (maximize preview area).	
File Browser, Editor	^ Ctrl + b	Add current image to the Batch Queue.	
File Browser	$\leftarrow \uparrow \rightarrow \downarrow$	Navigate through displayed thumbnails.	
File Browser	1 Shift + ← ↑ → ↓	Change the selected range of thumbnails.	
Tool Panel	^ Ctrl +	The slider is reset to the value it had when the image was loaded in the Editor, or when the image was selected in the File Browser.	
Tool Panel	^ Ctrl + drag a curve point	Makes the point move slower than the mouse cursor, allowing finer adjustments to the curve.	
Tool Panel	f Shift + drag a curve point	The point is snapped to key positions (highlighted in red).	
Tool Panel	Right-click on a tool's name	Unfolds clicked section and folds all others.	
Queue	^ Ctrl + s	Start/stop the queue.	

Windows-Only Shortcuts

Where	Shortcut	Action	
File Browser, Editor	F5	Open resulting image in the default image viewer (must have been processed before). Uses current queue settings to determine output file path.	
File Browser, Editor	^ Ctrl + F5	Open Explorer in the raw's directory and select current image.	
File Browser, Editor	1 Shift + F5	Open Explorer in the Batch Queue's output directory and selects the current image if already converted.	
Editor	Right-Alt	Temporarily enables both shadow and highlight clipping indicators while the preview is being re-drawn (e.g. while adjusting exposure compensation).	

Adding Support for New Raw Formats

Introduction

Supporting a raw format requires generally two things:

- Being able to decode the raw file, so that the program has access to the image data and metadata stored within each raw file. This is handled by a program embedded inside RawTherapee called "dcraw".
- Being able to interpret the image data. This can be further broken down into more items:
 - Measuring the black and white levels of the decoded image data, as every camera's idea of what is the blackest black and whitest white is different. This article covers this point.
 - Determining where the valid image lies on the "canvas" the "raw crop".
 - Correctly interpreting the colors. This is done by generating an input color profile (DCP or ICC), as described in the "How to Create DCP Color Profiles" article.

You will have to shoot the photos, but you do not heave to understand or carry out the measurement - we can do that for you.

What are black and white levels? A sensor is made of millions of tiny photo-sensitive elements called photosites or sensels. Each one measures the intensity of the light which falls upon it, and records that intensity as a number - the more light, the higher the number. The bulk of the raw file consists of these recorded measurements. Each photosite has a level below which it cannot sense any light - there might be some light falling upon it, but this light is too weak to register a signal. This is called the black level, and it is not always o. There is also a level beyond which the photosite will not register a change in light intensity even if the light does keep getting brighter - this is the white level. A photosite which cannot record any brighter light is said to be fully saturated, and in post-processing this state is called clipping.

Some camera models use the same white and black levels regardless of other settings, while for other models these levels depend on other factors, such as ISO sensitivity. We need photos from across the whole range of ISO values to determine this.

Some cameras have built-in noise reduction, often called LENR - Long Exposure Noise Reduction. It could affect the white and black levels. If you enable it, typically it only kicks-in when the exposure time is over 1s. The steps explained further on will explain when to turn it on and off.

The white levels for some camera models change depending on the aperture, but generally this only happens for wide-open apertures. To avoid this being a problem, set the aperture to f/5.6 or higher unless instructed otherwise.

What and How to Shoot

Shoot raw photos in manual (M) mode. If your camera has several raw modes, use the full one, uncropped, lossless compression if possible.

Each photo must be completely overexposed all across the frame. As such, it does not matter what you shoot since everything will be white anyway, but its easiest to achieve this while pointing the camera at the sky or at a white light bulb. It does not matter whether the sky is sunny or overcast, but don't point it at the sun as you might damage the sensor.

It does not matter what lens you use, but it will be easier to make the whole image overexposed if you do not use a wide angle lens.

Needed Photo Sets

We have broken down the image requirements into three sets, where each subsequent set would improve the quality of support, but would also require more effort from you and from us. In most cases, you only need to photograph the first two sets.

The sets:

- 1. If your camera has LENR, turn it OFF. Set the aperture to f/5.6 or higher. Take a series of photos, one photo for every ISO value your camera supports, making sure not to exceed an exposure time of 1 second, using an aperture of f/8. As an example, you could end up with about 8 photos for ISO 100, 200, 400, 800, 1600, 3200, 6400 and 12800. Most cameras include intermediate ISO values, e.g. ISO160 or ISO320, so if you wanted to improve white level accuracy you would need to photograph these intermediate values as well.
- 2. If your camera has LENR, turn it ON. Take a second series of photos, as described above but this time making sure that the exposure time in all cases is at least 2 seconds, not less. That's another 8 or more photos. In most cases these two sets should be enough.
- 3. Some cameras scale raw values for larger apertures, particularly Canon and Nikon models. The only way to know whether your camera does this for sure is to take a photo and measure it. Take one photo using your lens's widest aperture, e.g. f/1.7, at ISO100 with LENR turned OFF, and send it to us along with the rest of the shots. If we detect that there is raw scaling (or if you detect it yourself if you do your own measurements) then we will ask you shoot a series of photos using an exposure time of less than 1 second, from the widest aperture your lens supports down every 1/3 of a stop until such an aperture where raw scaling is no longer performed. This could mean many photos. Handling raw scaling caused by large apertures is not very important so don't feel daunted by it, you don't need to do it even if your camera does do raw scaling, but if you have the time and bandwidth then it would be better to check for it.

At the very least, you should end up with a series of about 8 photos from point 1. It is recommended that you take photos for both points 1 and 2, leading to 16 or more photos, plus the one raw scaling test photo from point 3. If it is found that your camera performs raw scaling, you could additionally take the needed series described in point 3, but since this could potentially mean many photos (over 50) it is not expected.

Compress all these photos, upload them to filebin.net (http://filebin.net/) and send us the full link either through our GitHub (https://github.com/Beep6581/RawTherapee/issues/new) page or in the Forum (http://rawtherape e.com/forum).

Completely clipped photos can have amazing compression, don't forget to compress them (7-Zip, ZIP, bzip2, whatever) before uploading! As an example, 10 completely clipped Sony 7M2 raw files with LENR disabled weigh 234MB but if you ZIP them you get a 1MB file.

More Information

For documentation detailing the required photos and instructions how to measure them, read the comments inside the camconst.json file:

https://github.com/Beep6581/RawTherapee/blob/dev/rtengine/camconst.json (https://github.com/Beep6581/RawTherapee/blob/dev/rtengine/camconst.json)

Storing and Reading This Information

RawTherapee looks for information regarding how to interpret the image data (the black and white levels and some other details) in three places:

- In the dcraw code which is embedded inside RawTherapee,
- In the raw file itself,
- In a text file on your system which is installed with RawTherapee called camconst.json

The needed information is gathered from all three places, and camconst.json is prioritized above the others.

File Paths

RawTherapee makes use of a "cache" folder to store temporary files which are safe to delete, and a "config" folder which stores your RawTherapee settings, custom processing profiles and other user-editable files. These folders reside in a special place, described below, and have a name that begins with the word "RawTherapee" optionally followed by a suffix. This suffix is set by the person who made the build of RawTherapee you're using. Some examples of what it can look like:

- RawTherapee
- RawTherapee4.2
- RawTherapee**5**
- RawTherapee5-dev
- RawTherapee_test
- And other possibilities always beginning with "RawTherapee"

The first part, "RawTherapee", is hard-coded. The second part, the suffix, is up to the person who made the build. It might be specific, like "5.0-gtk2-123-g87654321", it could be general, like "5", it could be anything else, like "_test", or it could be not set. We recommend that RawTherapee stable releases not use a suffix at all, while all development versions use "5-dev" - hopefully the person who made the build you're using took this into account.

Config

The RawTherapee config folder contains:

- the "options" file, which contains all of your settings from Preferences,
- the "batch" folder, which stores temporary processing profiles of the photos you sent to the Queue,
- the user-editable camconst.json file, where you can define details of how a specific raw format is to be treated (this overrides the values from the system camconst.json file),
- the dynamic profile rules,
- and the "profiles" folder where you can save your custom processing profiles to if you want them to appear in RawTherapee's drop-down list.

You could include this folder in your backups so that you can regain all of your settings and custom processing profiles if you install RawTherapee on a new system.

Default locations for the RawTherapee config folder (look for the "RawTherapee*" prefix as described above):

Windows XP

%USERPROFILE%\Local Settings\Application Data\

Windows 7, 8 and 10

%LOCALAPPDATA%

Linux

~/.config/

macOS

~/Library/Application Support/RawTherapee/config/

Under the Finder's 'Go' menu click 'Go to Folder' (shortcut Command+Shift+g), you can then type/paste any path you want to navigate to, even if it's hidden.

Cache

The RawTherapee cache folder contains sets of cached items, where each set consists of:

- a thumbnail,
- metadata,
- a sidecar file,
- and optionally an embedded profile.

By default, RawTherapee keeps up to 20 000 cached sets. Keep an eye on the "cache" folder as over time it may grow considerably in size! This is mostly due to the cached thumbnails which are stored in the "images" sub-folder. Deleting the "images" sub-folder is safe, you will not lose any image settings, RawTherapee will just have to regenerate the thumbnails.

A cached set used to include a 32kB histogram file, but as of RawTherapee 5.5 the need for storing the histogram was eliminated.

Default locations for the RawTherapee cache folder (look for the "RawTherapee*" prefix as described above):

Windows XP

%USERPROFILE%\Local Settings\Application Data\

Windows 7, 8 and 10

%LOCALAPPDATA%

Linux

~/.cache/

macOS

~/Library/Application Support/RawTherapee/cache/
Under the Finder's 'Go' menu click 'Go to Folder' (shortcut Command+Shift+g), you can t

Under the Finder's 'Go' menu click 'Go to Folder' (shortcut Command+Shift+g), you can then type/paste any path you want to navigate to, even if it's hidden.

Custom config and cache folders

Even though the path and name of the cache and config folders is hard-coded into RawTherapee, as of version 4.0.12.33 you can change them to any **absolute** path by using the RT_SETTINGS and RT_CACHE environment variables. How you do that depends on your operating system, so just search on the internet for "how to set environment variables in *your operating system*>".

Some examples:

Windows

```
RT_SETTINGS=%LOCALAPPDATA%\rawtherapee\4.2
RT CACHE=C:\Users\Bob\AppData\Local\rtcache
```

Linux and macOS

```
RT_SETTINGS=/home/bob/.config/rawtherapee/4.2
RT CACHE=/home/bob/junk/rtcache
```

Processing Profiles

If you create your own processing profiles, to have them appear in RawTherapee's "Processing Profiles" list you should save them to the "profiles" folder which you will find inside the "config" folder as described above.

Temporary Folder

The "Edit Current Image in External Editor" tool stores intermediate image files in a temporary folder:

Windows

The default location is the one stored in the TEMP environment variable, which is usually LOCALAPPDATA%/Temp

If you do not have the TEMP environment variable set, $C: \setminus$ is used.

Linux and macOS

The default location is the one stored in the \$TMPDIR environment variable, which is usually /tmp If you do not have the \$TMPDIR environment variable set, /tmp is used.

Preferences

You can access the Preferences window by clicking on the Preferences button which is either in the bottom-left corner of the RawTherapee window, or the top-right one, depending on your Editor tab mode layout.

About

The About button opens a window which contains a splash screen, technical details of the specific RawTherapee build you're running, credits, licence and release notes.

Please include these technical details when filing a bug report.

General Tab

Layout

Editor Layout

The layout of RawTherapee's user interface can be adapted to suite your taste and needs, specifically pertaining to whether you would like to have more than one raw file open simultaneously, and whether you use one monitor or more. The following modes are available:

- Single Editor Tab Mode
- Single Editor Tab Mode, Vertical Tabs
- Multiple Editor Tabs Mode
- Multiple Editor Tabs Mode (if available on second monitor)

Remember that if you simultaneously open several images each in its own Editor tab, each tab and instance will require a significant amount of RAM. Only use multiple Editor tabs if you have quite a lot of RAM - exactly how much depends on what resolution your images are, which tools you use, how many other programs you run in the background, and so forth, however a rule of thumb could be to not use multiple Editor tabs unless you have more than 8GB of RAM.

A restart is required for these changes to take effect.

■ Position of Curve Copy & Paste Buttons

Curves include adjacent buttons for copying, pasting, opening and saving the curve, and some include buttons for placing a node on the curve by picking a sample from the preview and for setting numeric in/out values. This option lets you decide where these buttons will be positioned relative to the curve widget.

A restart is required for these changes to take effect.

Histogram in Left Panel

Position the main histogram in the left panel above "History", or in the right tool panel above the tools.

Show Filmstrip Toolbar

The Filmstrip is a narrow panel which you can toggle to appear within the Editor tab. It contains thumbnails of the images in the currently opened folder, along with the filter and ratings toolbar. It can be useful to see the Filmstrip while working in the Editor tab, but you can hide it if you need the extra vertical screen space. Use this option to show or hide it. Note that you can also toggle its visibility from the Editor tab by using the "Toggle the visibility of the Filmstrip's toolbar" keyboard shortcut.

Compact Toolbars in File Browser

Enable this option if you have a high resolution screen to merge all the toolbars at the top of the File Browser tab into one.

If your screen's resolution is too low to fit the whole toolbar, some of the toolbar's contents (buttons, drop-downs, etc.) may become hidden. To see them, simply hover the cursor over the toolbar and use the mouse scroll-wheel to scroll the contents left and right.



The filter and ratings toolbar (4) can be shown in the Filmstrip.



Enabled, requires that the screen has a high enough resolution to fit all the elements horizontally.

■ Hide Vertical Scrollbar

You can hide the vertical scrollbar from the toolbox to save a little horizontal screen space. Use the mouse scroll-wheel to scroll when the scrollbar is hidden.

■ Tool Collapsed/Expanded State

If you have a few favorite tools which you would like to always see expanded, you can expand them now, hide the others, then return here to Preferences, disable "Automatically save tools' collapsed/expanded state before exiting", click "Save tools' collapsed/expanded state now", and click "OK" to close the Preferences window and commit your changes.

Alternatively, if you would instead like RawTherapee to remember which tools are collapsed and which are expanded at the end of every editing session, then enable "Automatically save tools' collapsed/expanded state before exiting".

Language

Select a language for the user interface. "Use system language" will try to auto-detect your language based on environment variables. You can override the auto-detected language by selecting one manually.

If you would like to help by updating one of the translations or creating a new one, see this post:

https://discuss.pixls.us/t/localization-how-to-translate-rawtherapee-and-rawpedia/2594 (https://discuss.pixls.us/t/localization-how-to-translate-rawtherapee-and-rawpedia/2594)

A restart is required for these changes to take effect.

Appearance

Theme

Choose a theme for the user interface. Although you will see theme-related changes as soon as you hit the "OK" button, you need to restart for the changes to take affect correctly.

The way human vision perceives colors depends on various factors, of particular importance to this paragraph are the properties of the area which surround the observed region. The way you perceive the colors of a photograph viewed on a screen depends in part on the colors of the area surrounding the photograph. You can read more about this in the CIECAMO2 article. In order to mitigate the errors the user makes while adjusting a photo, RawTherapee ships themes which use neutral background colors. While all of the themes are based on shades of grey, the theme which is most suited to avoid affecting human perception is "TooWaGrey - Average Surround", available from version 5.2 onward.



The Image Editor tab showing: (1) the preview background, (2) the crop mask, (3) lockable color pickers and (4) buttons which toggle the color of the preview background between black, white and theme-based.

Main font, and color picker font

Choose a custom main font, and a font for the Lockable Color Picker in the Image Editor tab, marked "3" in the screenshot.

Some users will find the default font size too small or too large due to their screen resolution and DPI setting. You can fix that by changing the font size.

Crop mask color

Adjust the color and transparency of the area outside of a cropped region, marked "2" in the screenshot. By clicking on the colored button, a new window appears where you can select a standard color or click on "Custom" to specify a new color. The vertical axis adjusts hue, while the horizontal axis adjusts transparency. Partial transparency is useful as it allows the cropped-off part of the photo to remain somewhat visible (2), so that you can move the crop around to find the best composition (hold the **Shift** key and move the crop with the mouse).

Navigator guide color

Adjust the color and transparency of the frame (marked "2" in the second screenshot) visible in the Navigator panel (marked "1") when the main preview is zoomed-in.

The Image Editor tab showing: (1) the Navigator panel, (2) the Navigator guide which marks the area currently visible in the main preview when zoomed-in.

Clipping Indication

The clipped shadow \triangle and \triangle highlight indicators in the Editor allow you to easily see which areas of the image are too dark or too bright. Highlighted areas are shaded according to the much they transgress the thresholds.

The thresholds for these indicators are defined in Preferences > General.

The clipped shadow indicator will highlight areas where all three channels fall at or below the specified shadow threshold. The clipped highlight indicator will highlight areas where at least one channel lies at or above the specified highlight threshold.

Clipping is calculated using data which depends on the state of the gamut button \not which you can toggle above the main preview in the Editor tab. When the gamut button is enabled the working profile is used, otherwise the gamma-corrected output profile is used.

Pan Rate Amplification

Imagine you are viewing a high resolution images while zoomed-in to 100%. In order to pan the image around the screen you would have to make multiple mouse movements (or have a very large mouse pad). RawTherapee saves you from this by using "pan rate amplification" - when set to 5, RawTherapee multiplies every pixel you pan by 5. If you'd normally move the cursor 500 pixels in one comfortable mouse movement, you will have panned 2500 pixels with this option set to 5.

The effect is most visible when you are zoomed in, and least visible when zoomed out.

When "Remember zoom % and pan offset" is enabled, when you open the next image RawTherapee will try to show the same area at the same zoom level as the current image. This only works in "Single Editor Tab Mode" and when "Demosaicing method used for the preview at <100% zoom" is set to "As in PP3".

External Editor

RawTherapee can send the processed image directly to an external program, e.g. an image editor, an image viewer or a script. This is done using the ** "Edit Current Image in External Editor" button in the Editor tab under the main preview - see the Saving article. It is here in Preferences where you can customize which program the processed image is to be sent to when you click the button.

Linux



- GIMP This option tries to run the following executables in the following order: gimp, gimp-remote. They should be accessible using the \$PATH environment variable.
- Custom command line Write the full path to the executable including the executable. You may need to enclose the whole line in double quotation marks if you need to pass arguments, see the example. Variables such as ~ and environment variables such as \$HOME are not supported.

Examples:

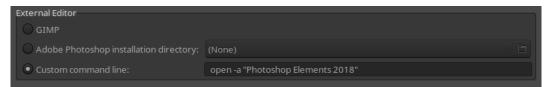
"/usr/bin/geegie --remote"

The above command opens the image in a single instance of Geeqie. Note that you need to enclose it in double quotation marks because you're passing the "--remote" option.

■ /home/bob/programs/luminance hdr/luminance-hdr

The above command opens the image in Luminance HDR. No arguments or options passed so no quotation marks needed.

macOS



■ GIMP - This option tries to run the following in the following order: open -a GIMP, open -a GIMP-dev.

- Photoshop This option is hard-coded to "open -a Photoshop".
- Custom command line Write open -a "External Program" where "External Program" is the name of the program to be used to open the image. Surround the name of the program in quotation marks if it contains one or more space characters.

Examples:

■ open -a "Adobe Photoshop CS6"

The above command opens the image in Adobe Photoshop CS6. Note that you need to enclose it in quotation marks because it contains space characters.

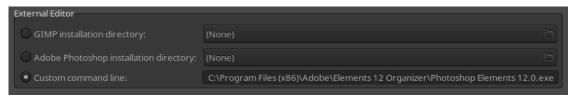
■ open -a "Affinity Photo Trial"

The command above opens the trial version of Affinity Photo. It too needed to be enclosed in quotation marks due to the spaces in the name.

■ open -a "/My stuff/Programs/Pixel Mixer"

The command above opens a program called "Pixel Mixer" in the "My stuff" folder. We have reports that it is not necessary to write the full path to the program even if it does not reside in the standard /Applications/ folder.

Windows



- GIMP Point RawTherapee to the folder which contains the bin folder which in turn contains the GIMP executable, gimp-2.*.exe. If you use an unofficial version of GIMP where the executable does not have that name, you may need to use the command line option instead. This option tries to run the following files in the following order: gimp-win-remote, gimp-2.12.exe, gimp-2.11.exe, gimp-2.10.exe, ..., all the way down to gimp-2.0.exe
- Photoshop Point RawTherapee to the folder which contains the Photoshop executable, Photoshop.exe.
- Custom command line Write the full path to the executable including the executable. Don't worry about spaces or about escaping backslashes. Environment variables such as %ProgramFiles% are not supported.

Examples:

- C:\Program Files\Gimp-2.9\gimp-2.9.exe
- C:\Program Files\Digital Light & Color\Picture Window Pro 6.0\pw60.exe

Image Processing Tab

Default Processing Profile

Specify which profile RawTherapee is to use when opening a raw or non-raw photo.

- The default processing profile **for raw photos** as of RawTherapee 5.4 is "Auto-Matched Curve ISO Low".
- The default processing profile **for non-raw photos** (such as JPEG, TIFF or PNG) is "Neutral". The "Neutral" profile just loads the photo as it is, without applying any changes.

To have processing profiles you have made yourself appear in the list, save them to the "profiles" subfolder within the "config" folder. You can find out where it is on the file paths page.

The special entry "Dynamic" activates the support for Dynamic Processing Profiles.

When you right-click on a thumbnail and select "Processing profile operations > Reset to default" RawTherapee will apply whichever processing profile is selected as default for that image type. If the default is set to "Dynamic", then RawTherapee will run through the dynamic profile rules to generate a profile dynamically.

Custom Processing Profile Builder

Executable (or script) file called when a new initial processing profile should be generated for an image. The path of the communication file (*.ini style, a.k.a. "Keyfile") is added as a command line parameter. It contains various parameters required for the executable or script to allow a rules-based processing profile generation.

This feature is very powerful; for example it allows you to set lens correction parameters or noise reduction based on image properties. It is called just once on the first edit of the picture, or called manually from the context menu when right-clicking on a thumbnail in the File Browser or Filmstrip

Note: You are responsible for using double quotes where necessary if you're using paths containing spaces.

Processing Profile Handling

Processing profile saving location

Choose whether you want RawTherapee to store the processing profiles next to the input file (the default behavior), to a central cache, or both.

It is a good idea to save the processing profiles next to the input files, as that lets you easily backup and handle your photos and their associated processing profiles.

Processing profile loading location

RawTherapee will look for processing profiles alongside the images, and in the central cache. If a profile exists in both places and they are not identical, this setting allows you to choose which one should have the deciding say.

Directories

Specify the location of your Dark-Frame, Flat-Field and HaldCLUT Film Simulation folders.

Crop Editing

This section lets you decide which guides are shown when the crop is __not__ being manipulated. "Original" means the guide type currently selected, so for example if you would like to see the "Rule of Thirds" guide while dragging a crop, and to have the guide automatically disappear once you are done dragging the crop,

then set this option to "Frame" or "None".

Once a crop is in place, RawTherapee can automatically zoom the cropped area to fit the screen if you enable this option.

Dynamic Profile Rules Tab

Here you can define your custom rules for creating Dynamic Processing Profiles.

File Browser Tab

Image Directory at Startup

At the top you can define the image directory to use at startup. It could be the RawTherapee installation directory, the last-visited directory, the home directory, or a custom directory.

File Browser / Thumbnail Options

These options determine which information is visible in the thumbnails and how it should be displayed.

Context Menu Options

Adjust the grouping of the right-click context menu in the File Browser (and Filmstrip).

Parsed Extensions

Choose which files are recognized as images and displayed in the File Browser. All supported extensions are set by default, except for PNG which is disabled by default.

If a desired extension is missing you can easily add it by clicking the "Add" + button.

Some users reported that their Parsed Extensions panel is empty. This could happen after updating from an unspecified older version of RawTherapee. If your parsed extensions panel is empty, we recommend you close RawTherapee, then find and delete the "options" file. The next time you run RawTherapee you will be using the latest defaults, and your list of parsed extensions will contain all supported formats.

Cache Options

To understand this section, first read the Cache article. The typical user should not need to change these defaults.

The maximum thumbnail height decides how large you can make the thumbnails. Each thumbnail is stored in RawTherapee's cache folder and requires disk space, so keep this in mind if you increase the default size.

The maximum number of cache entries decides how many of these cached files are kept before the oldest ones are deleted once the limit is reached.

You can manually clear elements of the cache using the "Clear" buttons.

You could "clear all cached files except for cached processing profiles" when updating RawTherapee to keep your disk clean and benefit from new cache-related improvements.

Color Management Tab

Use the "Directory containing color profiles" button to point RawTherapee to the folder which contains color profiles.

Standard locations where color profiles are stored:

Windows

C:\Windows\system32\spool\drivers\color

Linux

/usr/share/color/icc/

macOS

/library/ColorSync/Profiles/Displays/

Monitor

Set the "Default color profile" to the ICC file you generated when calibrating and profiling your monitor. You can have RawTherapee try to auto-detect the profile by using the "Use operating system's main monitor color profile" option, which relies on the <code>_ICC_PROFILE</code> X11 atom being set correctly. It supports only one monitor - if you have multiple monitors connected, it will always take the main monitor's profile.

Rendering intents and black point compensation are explained below.

The monitor profile must be of the "device" class in the RGB colorspace.

On macOS all displayed colors will be in sRGB space (https://en.wikipedia.org/wiki/SRGB), and then, if necessary, converted by the native macOS color pipeline to match the screen calibration, if any. This means that you cannot choose a monitor color profile on macOS. Colors will be displayed correctly, even over multiple screens, but if you have a wide-gamut screen RawTherapee's displayed colors will still be limited to sRGB. This will however not affect output, i.e. you can still produce images with colors outside the sRGB space. For more information, see: https://discuss.pixls.us/t/wide-gamut-preview-in-macos/2481 (https://discuss.pixls.us/t/wide-gamut-preview-in-macos/2481)

Printer (Soft-Proofing)

You can select here the color profile of your own printer or your print service in order to simulate the rendering of the printed image.

The printer profile must be of the "output" class in either the RGB or CMYK colorspaces.

See below for Black Point Compensation.

Rendering Intents

The "Rendering intent (https://en.wikipedia.org/wiki/Rendering_intent#Rendering_intent)" drop-down lets you choose how the ICC profiles are used for translation between gamuts or color spaces. When in the "Monitor" section, the "source" is the color space within which lies the image data at the end of the pipeline before being put into the monitor profile's color space, and the "destination" is the selected monitor profile's color space. When in the "Printer (Soft-Proofing)" section, the "source" is the image data at the end of the pipeline, and the "destination" is the selected printer profile's color space.

Relative Colorimetric

Colors from the source which lie outside the gamut of the destination color space will be shown using the nearest in-gamut color without affecting other in-gamut colors. The white point will be corrected. This is the default option and works with all profiles.

Perceptual

Colors from the source which lies outside the gamut of the destination color space will be compressed into the destination's gamut at the expense of also affecting in-gamut colors. How the compression is performed is up to the gamut mapping contained within the color profile - it usually involves desaturation, and sometimes even hue shifts. The perceptual intent only works with LUT profiles which contain the required gamut mapping tables - most ICC profiles do not, and in those cases "relative colorimetric" will be silently used instead (this is standard behavior across most software).

Absolute Colorimetric

Similar to relative colorimetric, but the white point will not be corrected. For this reason, it is used when you want to match paper whiteness to screen. You might want to use it when proofing, but not otherwise.

Black Point Compensation

When enabled, the Black Point level of the input image is moved to the Black Point level of the output image in a color transformation (e.g. from working profile to display profile). It means that the luminance channel alone is compressed or expanded to match the output capabilities. This feature will keep details in the shadows (avoid flat dark areas) at the expense of less color correctness.

Batch Edit Tab

Batch editing is making adjustments to more than one image at the same time. This is done through the Batch Edit tab in the File Browser.

The tool panel in the Batch Edit tab looks similar to the tool panel from the Image Editor tab, but it uses checkboxes to communicate which tool settings are consistent across the selected images and which are not. These checkboxes have three states:

[] Disabled	
[🗸] Enabled	
[-] Values differ acre	oss selected images

Batch editing is done by selecting multiple images in the File Browser (hold the 1 Shift) or ^ Ctrl key, then click the images you want to select), then you can edit those images using the tools in the Batch Edit panel on the right.

The controls (sliders, spinboxes, etc.) in the Batch Edit panel show the values of the processing parameters for the selected images. These can be the values of the default processing profile or the values from your last edit session of those photos.

If an image is currently being edited in the Editor, the editor's values will be reflected in real time in the Batch Edit panel, and vice versa, so take care what you're doing.

What happens to the tool values as you manipulate them depends on the "Behavior" setting in this Batch Edit tab.

The "Add" Mode

This mode may also be understood as "relative". Modifying sliders which are set to the "Add" mode will result in the value of the modification being added to the existing value. For example, if you select two images by holding the **Ctrl** modifier key, one image which has an

Exposure #Exposure_Compensation Exposure Compensation of -0.5 EV and the other which has +1.0 EV, moving the "Exposure Compensation" slider up to +0.3 will result in setting a value of -0.2 EV for the first image and +1.3 EV for the second one.

Using the "Reset" button will move the slider to its default (zero) position and will then bring back the initial value of that slider for each selected image.

The "Set" Mode

This mode may also be understood as "absolute". Modifying sliders which are set to the "Set" mode will result in the value of the modification being set, irrelevant of what the existing value was. If we use the same example as before, moving the slider up to +0.3 EV will result in setting a value of +0.3 EV for both images (one value for all images).

Using the 'Reset' button will move the slider to its default position (different for each slider), and will then reset this parameter for each image.

Performance Tab

The "Performance" tab is only for people who know what they're doing. It lets you poke under the hood and tweak some parameters depend on available RAM and CPU speed.

Preview Demosaicing Method

The "Demosaicing method used for the preview at <100% zoom" option sets which demosaicing method is used for the main preview in the Editor. By default, the same demosaicing method is used as specified in the Demosaicing section of the Raw tab, but if you are on a a very slow computer you can save a few hundred milliseconds by using the "Fast" demosaicing method. The trade-off is that the "Fast" method has the worst quality, though in most cases the difference is slight.

TIFF Read Settings

"Serialize read of TIFF files", enabled by default, can speed up thumbnail generation when opening for the first time a folder full of uncompressed TIFF files.

HaldCLUT Cache

The "Maximum number of cached CLUTs" setting lets you specify how many last-used HaldCLUT (Film Simulation) images are stored in RAM for faster access when switching back and forth between them in the Editor.

Inspect

Most raw files contain an embedded JPEG preview image. To show that image in the Inspect tab it needs to be extracted, which takes a fraction of a second. The "Maximum number of cached images" setting lets you specify how many of the last-viewed embedded images are kept in RAM, so that if you view the previous image in the Inspect tab, RawTherapee will not need to re-extract it, but just access it from RAM.

The "Image to show" option lets you decide whether to use the embedded JPEG image or to render one based on the real raw data using the "Neutral" processing profile. Using the embedded image is faster than rendering from the real raw data.

Threads

Splitting calculations and running them as concurrent threads (https://en.wikipedia.org/wiki/Thread_(computing)) allows them to complete faster, however doing so requires more RAM. By default, RawTherapee decides automatically how many threads to use. You can override this value here.

Most modern CPUs run two threads per physical core. Find out what CPU you have and how many cores it has, multiply that number by two, and you get the maximum number of threads it would make sense to run simultaneously. Let's call this number T_{max} . You would not benefit from running more threads than this - in fact you would likely suffer a small speed penalty.

Setting this parameter to "o" will let your CPU figure out what T_{max} is, and use that. If you experience crashes due to insufficient RAM, then you can calculate T_{max} yourself and use a number lower than that.

Sounds Tab

The "Sounds" tab lets you set an audible notification when a lengthy operation ends. It is currently only supported on Windows and Linux.

The "Queue processing done" sound is played after the last Queue image finishes processing. The "Editor processing done" sound is played after a lengthy in-editor operation that took longer than the specified number of seconds is complete.

Sounds can be muted either by disabling the "Enabled" checkbox or by setting fields with sound file references to blank values.

The "Queue" and "Editor processing done" text boxes can either point to wave (.wav) files, or can specify one of the following values:

Windows

- SystemAsterisk
- SystemDefault
- SystemExclamation
- SvstemExit
- SystemHand
- SystemQuestion
- SystemStart
- SystemWelcome

Linux

- bell
- camera-shutter
- complete
- dialog-warning
- dialog-information
- message
- service-login
- service-logout
- suspend-error
- trash-empty
- possibly the name of any file in /usr/share/sounds/freedesktop/stereo/ without the path or extension.

RawTherapee relies on libcanberra to produce sounds. In case of issues in Linux, you can trigger a sound to play from a terminal by invoking the following command (replace "bell" with any of the above):

canberra-gtk-play -i bell

Negative

Negatives are images with reversed lightness and colors, such as those produced by film cameras. RawTherapee does not yet have a flawless single-click solution for dealing with them, so this page serves to inform you of the possible workaround solutions:

- 1. Invert a diagonal tone curve either in the Exposure tool, or all of the curves in the RGB Curves tool. In the Color Management tool select "No profile" as the input profile. The downside is that there are tonal shifts when doing this.
- 2. Use a negative Hald CLUT via the Film Simulation tool. The "RawTherapee Film Simulation Collection" contains one, get it from the Film Simulation page. The downside is that some controls might operate in reverse, such as the Exposure slider, and you may experience clipping in the shadows and/or highlights as these tools are not designed to work with negatives.
- 3. In addition to using the neutral negative Hald CLUT as described above, if you have a successful workflow of not only inverting negatives but also toning them to your liking in RawTherapee or in other software, you could make your own negative Hald CLUT which reproduces the whole look including negative inversion. To do that, apply the same steps to the "identity Hald CLUT image" shipped with the RawTherapee Film Simulation Collection as you would to a photo negative, save it under a new name, then open a photo negative in RawTherapee and apply that new Hald CLUT image. This lets you instantly achieve not only negative inversion but also your own toning with the click of a button, leaving you only needing to expand the histogram by adjusting the Exposure slider or using curves.
- 4. Currently the best method is to use the DCP (DNG Camera Profile) for your camera model but edited in DNG Profile Editor so that the diagonal tone curve is inverted, then manually loading this DCP in RawTherapee for all negative shots. Method outlined below. A sample inverted DCP if you don't want to make your own can be downloaded from here: Linear inverted.dcp

Creating a DCP for Negatives

- 1. Get DNG Profile Editor (http://www.adobe.com/support/downlo ads/detail.jsp?ftpID=5494). It runs fine in Linux through wine (htt ps://www.winehq.org/).
- 2. Convert one of your camera's or scanner's raw photos (it can be the photo of the negative) to DNG by following the guide "How to convert raw formats to DNG".
- 3. Open the DNG image in DNG Profile Editor.
- 4. In the Color Tables tab, see whether the Base Profile called "Adobe Standard (<your camera model>)" is available. If it is, then select it. If it is not, then select "Choose external profile" and find the file titled "<your camera model> Adobe Standard.dcp". The guide How to get LCP and DCP profiles explains how to get them and where to find them.
- 5. In the Tone Curve tab invert the diagonal so that it goes from topleft to bottom-right (move the top point to the bottom).
- 6. Still in the Tone Curve tab, there are three "Base Tone Curves" for you to choose from. "Base Profile" and "Camera Raw Default" are usually identical and have more contrast, while "Linear" makes the image look flat. We recommend you save one DCP which uses "Base Profile" and another which uses "Linear", and see for yourself which one better suits your needs in RawTherapee, Both I

In: 255 Out 0 Custom

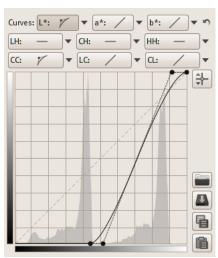
Screenshot of DNG Profile Editor's Tone Curve tab showing an inverted diagonal tone curve for use on negatives.

yourself which one better suits your needs in RawTherapee. Both DCPs will require further image tweaking in RawTherapee, but the "Linear" one will require more tweaking than the "Base Curve" one, though the latter might over-saturate colors - keep an eye on that.

7. To create the DCP, click "File > Export < your camera model > profile...". As recommended in the previous step, save two versions.

To use this new DCP for negatives, once you have your negative raw opened in RawTherapee go to the Color tab > Color Management section > Input Profile, and select Custom, then find this new DCP file. Enable "Use DCP's tone curve".

When tweaking images in RawTherapee using these DCPs, remember that using tone curves which operate on RGB channels (Tone Curve 1 and 2 in the Exposure tool) will change not only the lightness but also the color saturation the stronger your curves. You can control color saturation by using the "Weighted Standard" and "Saturation and Value Blending" modes, or you can avoid the problem by working in L*a*b* space. The advantage of working in L*a*b* space as the L* curve does is that colors are not changed while you modify the lightness, so you can fix the lightness first using as strong an L* curve as you need, and then fix the color saturation using for example the CC curve.



L* curve used to treat a negative raw with an applied negative DCP. The left spike is caused by the background around the scanned photo, which this curve clips to black. The right spike is the actual image on the photo.

GIMP Plugin

You can open raw images in GIMP using RawTherapee as a plugin.

Requirements

- RawTherapee 5.3
- GIMP 2.9.6
- The RawTherapee executable must be found from the \$PATH environment variable. This will be the case if you installed RawTherapee system-wide.

Usage

Just open a raw file from GIMP. A RawTherapee editor window should open automatically, which you can use to tweak your raw file. When you close the window, the image gets imported into GIMP.

Processing Profiles

Sidecar Files - Processing Profiles

Introduction

Processing profiles (with a PP3 extension for version 3 or PP2 for the older version 2) are text files which contain all of the tool settings which RawTherapee applies to the associated photo. If you are familiar with other raw processors, you may know their equivalent as "presets". They are stored alongside their associated photos, which is why they are also called sidecar files (https://en.wikipedia.org/wiki/Sidecar_file).

When you open a folder with photos in RawTherapee's File Browser for the first time, none of the images will have PP3 sidecar files. The thumbnails shown for images which have no processing profile assigned (images which have never been opened or edited) are created from the JPEG image embedded in each raw file. A processing profile is assigned to the image the moment one of these actions are taken:

- You open the photo for editing.
- You apply a processing profile manually, by using the right-click context menu in the File Browser or Filmstrip.
- You apply a dynamic processing profile.

When you open an image for editing, or when a processing profile is assigned, RawTherapee will convert the real raw data into a viewable image. In order to do this, there are many settings which need to be set to *something*, and what these specific values are depends on:

- Your default processing profile.
- Your dynamic processing profile rules, if any.
- Or on the processing profile you selected from the right-click context menu if you right-clicked on a thumbnail.

Sources

Processing profiles come from three quite different sources, though they work in exactly the same way:

■ "Bundled profiles".

RawTherapee comes with a bundle of profiles. Their purpose is to give you a good starting point, to demonstrate how the tools can be used together. They are the ones you see in the Processing Profile Selector drop-down list's "Bundled profiles" section, in the Image Editor.

"My profiles".

When you make a processing profile which you want to re-use, for example one which works well with your camera and your style, you can save it so that it also appears in the Processing Profile Selector drop-down list, in the "My profiles" section. To have it appear there, save it to the "profiles" folder within the "config" folder - see the File Paths article to find it.

Automatically generated profiles.

Whenever you edit an image, the tool settings you want applied to that image are stored in a processing profile that is particular to that image (ranking information, the history panel contents and snapshots are not stored in these files yet, see issue #473 (https://code.google.com/p/rawtherapee/issues/detail?id=473)).

Saving

As simply viewing the image requires processing, RawTherapee stores the settings it used to show you the image in a sidecar processing profile. That processing profile also stores all the tool tweaks you made in the Editor tab.

The processing profile is written to disk:

- When you apply a processing profile manually or using a dynamic profile.
- When you close the current image (the Editor tab) if using Multiple Editor Tabs Mode (METM).
- When you close the current image by opening a different image if using Single Editor Tab Mode (SETM).
- When you close the current image by closing RawTherapee.
- When you manually save the processing profile using the Processing Profile Selector panel in the Editor tab.
- When you use the "force saving current settings to the processing profile" keyboard shortcut from the Editor tab.

If a photo has an associated processing profile, a green check mark will appear over its thumbnail.

If you have a photo opened in an Editor tab and you make changes to it from the File Browser, the changes are reflected immediately in the Editor tab.

Storage

Where the processing profile is stored can be configured in Preferences > Processing Profile Handling.

By default, the processing profile for an image is stored alongside the input image (if you open kitty.raw, a new file kitty.raw.pp3 will be created next to it), but they can also be stored in a central cache. You can choose whether RawTherapee should use the cache, write the processing profile alongside the image, or both, from "*Preferences > Image Processing*". We suggest you store these files alongside your input image files so that if you decide to move the images you can move the processing profiles easily along with them.

When saving an image you have the option of ticking the "Save processing parameters with image" checkbox. If it is ticked, and if you are working on kitty.raw and saving to a JPG file, then the processing profile used to develop that image will be stored in a file called kitty.jpg.out.pp3. The ".out" part if there so that conflicts do not occur if you are working on a non-raw file.

Defaults

The default processing profile used when opening **non-raw** images is called "Neutral". This profile has all tool settings at their neutral values, so they have no effect. Since non-raw images usually have already been processed and are ready for viewing, having RawTherapee not introduce any tweaks by default is the desirable behavior.

The default profile for **raw** photos is called "Auto-Matched Curve" (from RawTherapee 5.4 onward). This profile makes your raw image look like the out-of-camera JPEG, which is usually a desirable starting point.

Furthermore, most tools in the Editor tab have a reset button.

- Clicking the reset button resets the tool to its hard-coded neutral value, usually zero.
- Ctrl+clicking the reset button resets the tool to whatever value it had when you opened the image, i.e. the way it was if you rewind the history stack to the top.

Partial Processing Profiles and Fill Modes

Processing profiles can be copied, pasted, loaded, saved and applied in full or partially. Clicking the buttons in the Processing Profile Selector handles them in full, while ctrl+clicking those buttons handles them partially - a windows pops up letting you choose which parameters to include. This feature lets you, for example, copy only the noise reduction parameters from one image to another, or only the white balance, or only both.

The processing profile fill mode lets you decide what happens when you apply a processing profile to an image when that profile does not contain information about all of the tool settings.

■ "Fill" mode 🖹

When in "Fill mode", the missing values will be taken from RawTherapee's hard-coded default values (typically neutral). For instance if you apply a partial profile which contains only sharpening settings, all of the remaining tools (such as Exposure, Tone Mapping, Noise Reduction, Resize, etc) will pop into their default positions.

■ "Preserve" mode 들

When in "Preserve mode", only those values in the profile will be applied, and the missing ones remain unchanged.

Creating Your Own Processing Profiles

Using certain tools in certain ways may make your processing profile only usable with that specific image. For example if you set a white balance, cropping and rotation, you won't get good results if you apply that profile to an image taken under different lighting with the camera rotated any other way. See the article Creating processing profiles for general use for advice on how to make processing profiles which can be used on many images.

Compatibility

Processing profiles evolve from one version of RawTherapee to the next. We strive to ensure backward compatibility (e.g. a profile created in 5.3 and opened in 5.4 should look the same), but this is not always possible.

Processing profiles can gain new parameters or lose ones which became obsolete. Tool behavior can also evolve, wherein default values change or in extreme cases the meaning of a value is interpreted differently; an example of this is the noise reduction tool, where a luminance noise reduction value of 10 in RawTherapee-3.0 would lead to a different result in RawTherapee-4.0.10 as the whole noise reduction engine has been greatly improved.

Consolidating processing profiles into a cache allows one to store isolated copies of the processing profiles per specific version of RawTherapee. In such a case, the cache can be used to re-process photos in order to get the same output as originally intended (but e.g. with a new size or output color space) using the same version of RawTherapee in which the image was originally edited. Whether this is desirable is debatable. Consider that you want to squeeze as much out of your raw files as possible. If two years later you want to go back to an old raw file, perhaps getting the same result as you did two years ago is not the best idea, because RawTherapee's capabilities would have greatly improved in that time, you may have acquired a better monitor, and your taste and skill would also have evolved. Nevertheless, by backing up the whole cache folder, when installing a new version of RawTherapee you retain the option of going back to an older version of RawTherapee in order to get the exact same result.

The File Paths article describes where you can find the "cache" and "config" folders on your system.

When releasing a major new version of RawTherapee, it may happen that we use a new suffix for the "cache" and "config" folders. This means that the new version of RawTherapee will not see your old configuration or processing profiles. Though this sounds undesirable, there are good reasons we may (rarely) choose to do that.

- Backward compatibility. There may be changes in behavior between old and new versions of a specific tool. For instance, the effects of the Auto Levels tool have changed (for the better) between versions 4.0.11 and 4.0.12, so if your old processing profiles had it enabled, the results in 4.0.12 will be a little different and may require tuning your old profiles.
- Some users have not checked "Preferences" in a long time, and their program is tuned for what worked best long ago, not for what works best now. Our defaults are good defaults, we keep them up to date to make RawTherapee look and function well out-of-the-box, so sometimes having RawTherapee start with fresh defaults is a good thing, and it will motivate users to look into "Preferences" again.
- Some users have never looked inside "Preferences" in the first place, and are unaware of some of the features that can be unlocked there. As above, fresh defaults will activate these things.
- Some old cache and config files can cause RawTherapee to crash. While we patch the specific cases made known to us, it is safe to assume there will always be cases unknown to us which will still cause instability. Starting with clean cache and config folders mitigates this problem.

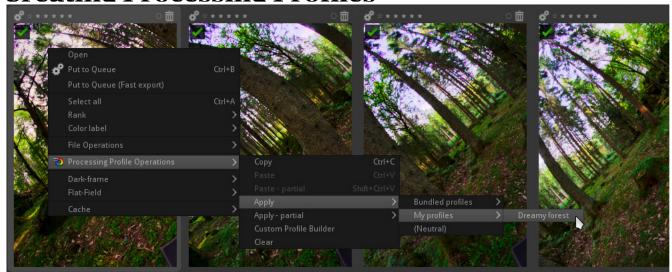
Creating processing profiles for general use

In RawTherapee, we call the sidecar files 'processing profiles'. We supply a bunch of processing profiles with RT, so that you can start off with an existing look and modify it to your liking, saving you some time. One example of such a profile is "Pop 1" - it will make your photo vibrant and lively, lifting the shadows and bringing out detail.

You can see the whole list of processing profiles in the *Image Editor* tab, if you expand the *Processing Profiles* list. You can also see them if you right-click on a thumbnail in the *File Browser* tab and move your mouse over to "*Processing Profile Operations > Apply Profile*".

Read the short "Processing Profile Selector" article to make sure you understand how to make full use of the selector, partial profiles and the fill mode toggle button.

Creating Processing Profiles



Apply a custom-made processing profile using the context menu in the File Browser (or Filmstrip).

You can create your own processing profiles and have them shown in the Processing Profile Selector drop-down list.

- Open a photo you want to create a good starting point profile for.
- You could start off with the 'Neutral' profile, or make changes to any of the other profiles that come bundled with RawTherapee. Just apply the desired profile to your photo.



Apply a custom-made processing profile using the dropdown menu in the Processing Profile Selector in the Image Editor.

- Make the changes you like, remembering that the more specific your tweaks, the fewer photos they will work well with because every photo is different so what works well for one may not work well for another if they differ significantly. For example if your camera has a very low-noise sensor and your lens is not very sharp you could probably enable sharpening, or conversely if your camera has a noisy sensor you may want to apply a certain level of noise reduction by default. Maybe you'd like to set your name in the IPTC "Author" field and have RawTherapee copy your metadata changes to the saved files. You will generally want to leave the white balance set to "Camera" since your photos will be taken under various lighting.
- When you are done tweaking, click the *Save Current Profile* icon ▲ in the *Processing Profiles* panel. Enter any name; you don't need to specify the extension RawTherapee will add it for you. To have it appear in the drop-down list you need to save it to the "profiles" sub-folder in the "config" folder refer to the File Paths page to find out where this folder is on your system.
- Restart RawTherapee, and now your new processing profile will appear in the drop-down list under "My profiles".

Partial Processing Profiles

Sometimes, you will want to save only a subset of the parameters available, e.g. to avoid storing geometric parameters like rotate, crop and resize. In this case, hold the *Control* key while clicking on the *Save* button. When you select the output file name and click *Save*, a window will let you choose which parameters to select. You can then share these profiles with your friends or in our forum.

Remember that in order for a profile to be universally applicable to all photos of the same scene and situation (baby portrait photos in this example), you need to think of all the variations in all of the baby portrait photos you might want to apply it to. Remember that exposure will vary between shots, even if you shot the baby in a



Save Partial Processing Profile window

studio, as the little one is likely to be crawling around, and even more so if you upload your profile on the internet for other baby photographers with different cameras and different lighting gear to use, so instead of setting a specific exposure, such as +0.60, you should rather turn on *Auto Levels*. This applies to all other settings - remember to set just the bare minimum number of options to achieve the effect you want. Leave the rest untouched, as it is very likely that if you had set those other options, they will not apply well to other photos. If your processing profile is meant to make baby face photos look soft and cuddly by a clever mixture of highlight recovery, auto exposure, Lab and RGB tone curves, then don't enable noise reduction (as photos might be shot at different ISO values), don't set custom white balance (as light may have changed between shots), don't rotate the photo, and so forth. All these superfluous parameters are likely to change between photos and not influence your soft baby look in any way, so turning them on will just litter your profile. Double-check these things before sharing your profiles.

Default Raw/Non-raw Processing Profile

To use your own processing profile as the default for raw or non-raw photos, set it under "*Preferences > Image Processing > Default Processing Profile*".

If you want to use a dynamic profile, then set "Default processing profile for (non/)raw" to "(Dynamic)".

Dynamic processing profiles

Sometimes a single, "static" default processing profile is not enough to cover all use cases. For example, the amount of noise reduction to apply varies according to the camera and ISO setting used. Another example is the kind and amount of lens corrections needed, which is obviously dependent on the lens used.

In order to handle such cases, RawTherapee provides a feature that allows to create a default processing profile "dynamically", based on the metadata of the image being processed (such as camera and lens name, shutter speed, ISO value, and so on).



The Dynamic Profile Rules panel in RawTherapee-5.4.

This is done by defining a set of "dynamic profile rules". Each rule has a (partial) processing profile attached to it, plus some conditions on the image metadata that define whether the rule is applicable. When a picture is edited for the first time, the list of rules is scanned, and all the profiles that match are combined (in the order given, so later rules can override earlier ones) to build the initial processing profile.

In order to activate the functionality, the default processing profile must be set to "(Dynamic)". Rules are defined in the Dynamic Profile Rules section of the Preferences window.

Dynamic profile rules work on the following image metadata:

Camera

the camera name (including brand) as shown in the image info overlay of the Image Editor. If active, by default this entry will cause the rule to apply only to pictures taken with the exact camera specified here (the name is case-insensitive). However, if the entry starts with the re: prefix, then the rest of the string will be interpreted as a regular expression (https://en.wikipedia.org/wiki/Regular_expression) to use for the matching. For example, a rule with the Camera value set to re: SONY ILCE-[56].00 will be applied to all Sony Alpha a5xxx and a6xxx cameras.

Lens

The full lens name. As above, a regular expression can be used by starting with the re: prefix.

ISO

The range of ISO values.

Aperture

The range of apertures of the lens, measured in f-stops (https://en.wikipedia.org/wiki/F-number).

Focal length

The range of focal lengths used, in mm.

Shutter

The range of shutter speeds, in seconds. For example, enter 0.03 for a speed of 1/30".

Exposure compensation

The range of exposure compensation values, in stops.

Processing Profile Scripts

RTProfileSelector

RTProfileSelector is a RawTherapee plugin that automatically selects custom processing profiles (.pp3 files) based on user-defined rules. The rules are sets of Exif fields and values which are matched against the actual values extracted from the raw files the first time they are opened in RawTherapee.

A few things you can automate in RawTherapee through RTProfileSelector:

- Assign your own custom processing profiles to approximatelly match your camera settings (such as "monochrome"/"black-and-white", "vivid color", "film modes", etc.)
- Set noise reduction parameters in RawTherapee according to camera model and ISO value
- Assign lens correction profiles (LCP) based on the lens and camera model used

RTProfileSelector is written in C++11 and compiles on both Windows and Linux. Source code and Windows executable can be downloaded from its GitHub repository (https://github.com/marcapelini/RTProfileSel ector). RTProfileSelector uses ExifTool (http://www.sno.phy.queensu.ca/~phil/exiftool/) for extracting metadata from images.

For installation procedures and online documentation, please go to the project's wiki section (https://github.com/marcapelini/RTProfileSelector/wiki).

RTbatch

RTbatch is a Python (https://en.wikipedia.org/wiki/Python_(programming_language)) script which allows you to easily process an input image with all of RawTherapee's built-in processing profiles and export either full-size pictures or lower resolution previews. RTbatch can also export markdown (https://en.wikipedia.org/wiki/Markdown) code for displaying galleries of your images on websites that parse markdown (currently this is only implemented for Octopress (http://octopress.org/) using the *imageap* or *imag* tags).

The script also comes with a command-line utility, *RTbatch_cli*, with which you can call it and pass all arguments in a single line from your terminal.

Source: https://github.com/TheChymera/RTbatch (https://github.com/TheChymera/RTbatch)

Bundled Processing Profiles

Default

Default

Faded

The **Faded Pack profiles** allow the application of nostalgic, thematic, sentimental looks to images with a general appearance of faded colors and additional color toning that is varied across the luminosity range of the image. Most profiles come in three variants:

- 1. The base version,
- 2. a TM (tone-mapped) version, and
- 3. a slightly brighter tone-mapped version.

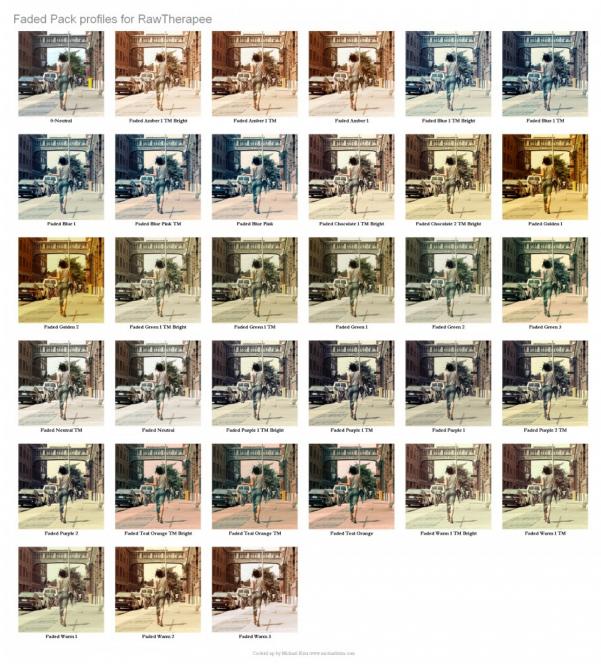
A side benefit of these profiles is that they handle burnt highlights very gracefully, making their appearance more natural and seamless with the rest of the image.

The essence of the toning effect is created mostly by settings of the Channel Mixer and RGB Curves (both in Color Tab). Channel mixer adds the fading effect and RGB curves apply the further color toning.

To tune profile to personal taste, one could adjust the RGB curves to alter the coloring effect. If color toning is acceptable, one could continue image adjustments with the tone curves in the exposure panel.

Faded Pack profiles for RawTherapee Faded Warm 2

Faded profile pack, city example. Click for larger view.



Faded profile pack, woman example. Click for larger view.

Tool Description

General Comments About Some Toolbox Widgets

This article is important - some things may sound obvious if you are an experienced user, but some things are guaranteed to be unknown to you and knowing them will let you use RawTherapee faster and easier. Take five minutes to read it!

The panel on the right side of the preview contains the controls for all the tools available in RawTherapee. They can do a lot, perhaps even more than you ever may want to! If you're new to RawTherapee or new to raw processing in general, don't feel overwhelmed, as there is no need to touch all those sliders to get decent results. In this section you'll find a brief description of what all those tools are for, tab after tab.

Panels

A panel is a foldable element which contains buttons, tool, histograms, etc. The Editor tab has three main panels: the one on the left which contains the history, the one on the right which contains the tools, and if you're in "Single Editor Tab Mode" the one on the top which contains the Filmstrip. They can be hidden away from view with the little arrow buttons or using keyboard shortcuts, making more room for the image preview.

You can use the mouse scrollwheel to safely scroll up/down in the panels without fear of accidentally changing a slider, because RawTherapee requires that you hold the Shift key while using the mouse scrollwheel if your intention is to manipulate a slider or cycle through the options in a drop-down menu (called a "combobox").

Tools

RawTherapee has many tools, and each tool has one or more controls, or widgets (https://en.wikipedia.org/wik i/Widget_(GUI)), which you can turn or move or push or slide to make the program do something. They are divided up among the Exposure, Detail, Color, Wavelet, Transform, Raw and Metadata tabs, for example the Exposure tab has tools which deal with exposure, such as the Exposure tool and the Shadows/Highlights tool, etc., the Color tab has the White Balance tool, and so on.

Because there are so many tools, you will notice that vertical screen space is precious, and you will at times hide some tools to see others. RawTherapee makes this easy. Right-clicking over a tool's title expands it and collapses all other tools of the same tab. You'll learn to love this shortcut when you consider the time you would have spent on folding tools manually...

To the left of most tools' titles is a button which lets you enable or disable the corresponding tool. The concept of being enabled or disabled doesn't always make sense - for example what would it mean for the White Balance tool to be disabled? There must always be some white balance being used - so these tools instead of an enabled/disabled button just use a triangle symbol to let you expand or collapse them. Lastly, you will notice that in the File Browser tab if you select several photos then the power button can have a third state which looks only half-enabled - this state, called "inconsistent", means that the tool in question is not enabled in all of the selected photos.

- Tool disabled.
- ტ Tool enabled.
- Tool inconistent.
- Tool collapsed.

Tool expanded.

Sliders

Each slider has three values in memory:

- The current value, when you move the slider to any position,
- The 'default' value, the one that the programmer has set as default. It can be called back by clicking on the 'Reset' button.
- The 'initial' value, which is the value of the profile used when the image was loaded in the editor. It can be called back by **Control-clicking** on the 'Reset' button.

As written above, if you want to move a slider using the mouse scrollwheel, or cycle though a combobox, make sure you hold down the Shift key while doing so, otherwise instead of moving the slider you will scroll up/down in the panel.

Curve Editors

RawTherapee has three types of curve editors – for Threshold Curves, Tone Curves and Flat Curves. They are discussed below along with some general pointers.

Threshold Curves

Threshold Curves are the simplest. They are used to tell a RawTherapee tool the tones (or hues or saturations values) that you want processed (or processed differently).

As an example, consider the Threshold curve editor on the Detail -> Sharpening tool. The setting shown is telling the Sharpening tool to phase in sharpening quickly in the black areas (the steep line up on the left), maintain full sharpening through mid-tones (the plateau area) and then phase out sharpening slowly in the highlights (the long slope down). Dragging one of the control points will move the slope leading up or down from the point. To just move the point, but not the slope as a whole, hold down the shift key while you drag.



On the top right is a reset button that will reset back to default.

Warning: resetting the curve is considered a curve modification, so if you've just modified the curve and mistakenly pressed the *Reset* button, there's no way to bring back your curve. (Ctrl-z will go one step before in the *History* list, not in the curve's edition). This comment applies to Tone and Flat Curves as well.

You'll also find Threshold Curves used in the Contrast by Detail Levels and the Vibrance tools.

You may have noticed that Threshold Curves actually consist of a few straight lines rather than a curve. If this bothers you, you might want to take a break before moving on to Flat Curves.

General Comments on Tone and Flat Curves

Tone and Flat Curves are more powerful than Threshold Curves and have some controls that aren't needed on Threshold Curves.

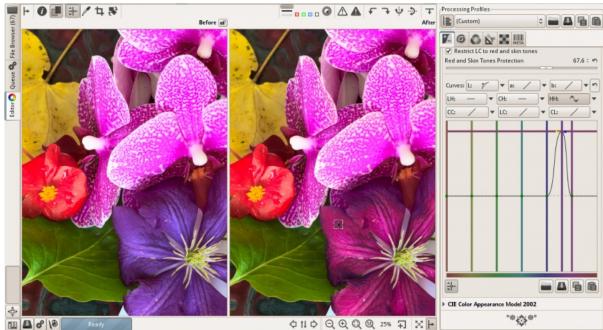
Each Tone or Flat curve editor has a button to select its type. It's a so-called 'Toggle' button, i.e. it will stay pressed or stay released after each click on it. Toggling on/off the curve's button will respectively display/hide its associated editor. This is very handy and saves a lot of space when handling groups of curves (e.g. see the Lab curve editor).

These curves have a Reset button on the right which will reset the displayed (or pressed button) curve only.

They also have a pipette tool and a node in/out value input tool, which deserve sections of their own.

Pipette

Most



The pipette tool pinpointed the exact hue of the purple flower to easily make it more red using the HH curve.

curves in RawTherapee have a pipette button . This feature is a great way to target specific hues, tones and areas of a particular saturation, and using it will make tweaking your images easier and faster. Let's say you want to change the purple hue of a flower, to make it more red. Without the pipette, you would have to guess the exact shade of the hue, else the affected range of hues would be too wide and you might end up changing something you did not want to change. With the pipette, you simply click on the flower, and a point appears in the relevant curve. You can adjust this point to your liking, knowing it represents the exact hue you want affected.

The pipette tool must be activated for each curve, simply by clicking on it. Now when you hover the cursor over the main preview you will notice that a vertical line (or even four of them) appears in the curve panel. This line represents the value of interest of the pixels you're hovering over. To place an adjustment point in the curve for the value you're hovering over, Ctrl+leftclick in the preview, and a dot appears in the curve. You can adjust that point without leaving the preview area, just keep holding the left mouse button after you placed the curve point, and moving the mouse up and down moves the point up and down. Remember that holding Ctrl while editing a curve point decreases your mouse speed so that you can very finely adjust points, but usually you will not need this, so after Ctrl+leftclicking to place the point, let go of Ctrl, but keep holding the left mouse button.

To deactivate the pipette, either right-click anywhere in the preview area, or left-click on the same pipette's button again.

You do not need to deactivate the previous curve's pipette to use it on a new curve, just activate it as usual, and the old one automatically deactivates.

Curve Node In/Out Value

Introduced in version 4.2.192, each curve has a node in/out value input tool. You can use this tool for example to match node values on a photo of a color target with known patch values.

The tool works with nodes, and the most likely way you will create these nodes is by using the pipette. For this example, we will start with a curve without any nodes and create some using the pipette. Click the node value editor button next to the curve, and also click the pipette's button. You will now see "I" (in) and "O" (out) values displayed under the curve. They correspond to the point under the mouse cursor if you hover it over the curve. Hover the cursor over the image preview. Since you activated the pipette, you can Ctrl+click on a spot in the preview area to place a node which corresponds to that spot's value in the curve (whatever that value may be - for example for the L* curve, the value is the internal value of the pixel you've placed the spot over at that point in RawTherapee's internal pipeline, and the output value is that value affected by your curve). Do that, Ctrl+click on a spot in the preview. A node appears in the curve. To edit the node's in/out values, right-click on the node. It turns red with a red ring around it. Now you can edit the in/out values and see the node move in real-time. Once you are done editing, either right-click anywhere inside the curve area other than on the node to go out of node editing mode, or just click the node value editor button again to deactivate it.

Tone Curves

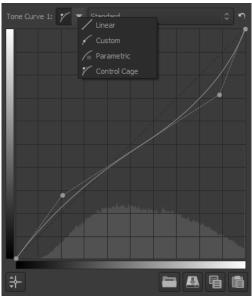
Tone Curves are somewhat misnamed since, while some are used to adjust tones, others are used to adjust saturation, chromaticity or other properties. Their purpose is to map an input value (on the horizontal or X axis) to an output value (measured on the vertical or Y axis). If Math isn't your first language, don't worry – after a little playing with them, everyone seems to quickly develop an intuitive understanding.

As an example, consider the figure to the right – part of the Exposure tool. The left part of the graph represents the darker tones, the right part represents the brighter tones of the photo. You can see that the bottom left of the curve has been moved up. This will cause dark areas to be boosted. Similarly the top right of the curve has been pulled down, cutting back the bright areas.

The figure shows the curve type drop down. Tone curves allow four different ways of manipulating the curve:

- **Linear:** The default type a straight line that results in no change to input values. The mathematically inclined may observe that it is a graph of y=x. The rest of us just set the control to linear to "turn off" the curve.
- **Custom:** The type most commonly seen in other software. Click to drop a control point anywhere on the curve and then drag the control point to change the curve's shape. The top-right point represents the brightest areas in the photo. Drag that point vertically down to make the highlights less bright; move it horizontally to the left to make bright areas brighter, perhaps at the cost of some overexposure. The bottom-left point represents the darkest areas in the photo. Move that point horizontally to the right to make the photo darker, perhaps at the cost of some underexposure. Move it vertically up to make the darks lighter.

- Control Cage: At first sight this curve type looks very much like the Custom curve, but there are some differences. With the Custom curve, the curve touches all the control points. This is not the case with the control cage curve the control points attract the curve towards them but the curve doesn't actually go through them. Another difference is that the control cage allows for a straight section of the curve, while you can't do this with the custom curve. The cage curve needs at least three points for that (so five in total). Holding down the Shift key while dragging a point will help you to easily create a straight line by snapping the point to the line made by the previous and next point (displayed in red by the 'snap to' tool). Many users prefer Control Cage type curves to the alternatives.

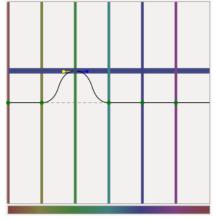


Tone Curve showing type drop down

The Flat Curve

A number of tools in RawTherapee use the *flat curve*:

- Lab Adjustments
 - LH
 - CH
 - HH
- Defringe
 - Hue
- HSV Equalizer
 - H
 - S
 - V



The input color green (vertical line) was changed to the output color blue (horizontal line).

It's very simple to use once you understand it, so let's use the HSV Equalizer in the Color tab as an example. Click on the drop-down icon next to the H(ue) button and choose "Minima/Maxima"

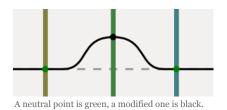
control points" ... You'll see six dots on the horizontal line in the middle and six vertical lines that cross these dots. Notice that those lines are colored; from left to right: red, yellow, green, aqua, blue and magenta. Now click on the very left dot (the cursor changes into a little hand) and move it slightly upward and downward. Result: red colors quickly change to green, blue and magenta as the cursor is moved up, and to pink, blue and green when moved down.

Notice that a new horizontal line appears when you start dragging a color point, and see how its color changes. The vertical axis represents input colors, and the horizontal axis output colors.

When you click and drag a vertical line (the line, not the point!), the very first movement will determine the kind of move: vertical or horizontal (so take care with this first movement if you want to have a predictable result). If you want to move the point in both directions at the same time, then click and drag the point itself. To move the point only in one direction (only horizontally or only vertically) you can use the 'snap to' function by holding down the Shift key while moving the point.

It's easy to see if a point is on its neutral value (i.e. on the middle line) because the color of the point will be green. As soon as you move a point off its neutral value, it changes color to black.

The HSV Equalizer wraps around on the horizontal axis, so the very right vertical line equals the very left line. You can see this by dragging the red line on the left side a bit to the left. Now the left point of the graph is at the same position as the very right point.



Holding the **Shift** key while dragging a point prevents it from wrapping around the horizontal axis, which can be useful in preventing accidental curve steps in hard-to-see places at the edges.

You can delete points by dragging them out of the editor field. You can add points by clicking somewhere on the curve. When you place the mouse on one of the points, you see a yellow and blue indicator. Place the mouse on the yellow one and the cursor changes into a left arrow. Now you can drag this point to the left, to change the slope of the curve. Ditto for the blue indicator.

To get an idea how this editor works, delete all but two colors (e.g. red and yellow) and move the graph around, change its slope and see what happens to your photo.

Reset the *Hue* curve to "*Linear*" (no changes) by clicking on the reset icon next to the *Value* button. To compare the effects of the *Hue* curve with linear: switch between "*Linear*" and "*Minima/Maxima control points*" in the drop-down menu next to this button, or use the history list on the left side of your screen.

You can save a curve for later use by clicking on the disk button. Note that only the actual (shown) H, S or V curve is saved, not all three at once, so don't give your curve a name like my_hsv because it doesn't describe whether the curve inside it is H, S or V, but instead name your saved curve files something like my_hue, my_sat and my_val. The extension will be added automatically, ".rtc".

The Preview Area

The preview is designed to show you the most realistic result possible, however one must keep in mind that the larger an image is, the longer it takes to process. For speed reasons, the preview of the effects of most tools is calculated not on the full-sized image (which would take exactly as long as saving the image, making using sliders and curves impossible), but on the preview image which is of the size of your preview area. Many tools, such as the Exposure tool, can be applied to an image of any size, and their effects will be identical regardless of the size of the image they are applied to. However some tools are sizedependent, for example all of the tools in the Detail tab, which means if you apply one of these tools to a full-sized image and to a smaller version of that image, and then downscale the full-sized image to the smaller image's size, the two images won't match. For reasons of speed, RawTherapee must use the small preview image so that your tool tweaking experience may be fast and fluid, but this means that the effects of size-dependent tools would not be accurate when applied to a small zoomed-out preview. We made the decision to either disable the preview effects of these tools entirely at zoom levels less than 100%, or to keep the preview effects active but to warn you that what you see at zoom levels less than 100% may be inaccurate depending on the tool settings (for example Tone Mapping and Wavelet may be accurate at zoom levels less than 100% or they may be inaccurate, depending on their settings). You will know which tools these are because they are marked with a "1:1" icon 14 next to their names. RawPedia explains how accurate the preview is for all affected tools on each tool's page.

Exposure Tab

Exposure

Auto Levels

The *Auto Levels* tool analyzes the histogram and then adjusts the controls in the Exposure section to achieve a well-exposed image.

It uses only the "Exposure compensation", "Highlight compression", "Highlight reconstruction", "Black", "Lightness" and "Contrast" sliders.

Think of the adjustments Auto Levels comes up with as a good starting point. It has been tuned to work best with "typical shots", so the result should often be aesthetically pleasing, but as the program doesn't know your taste or expectations this will not always be the case. For example you might be going for a high-key look (i.e. not typical), in which case you should adjust the values yourself.

You can reset all of the sliders in the Exposure section by clicking on the Reset button. The Tone Curves will not be touched.

Clip %

Auto Levels uses the "Clip %" value to adjust the exposure. This number defines the percentage of pixels allowed to clip to the white and black points of the raw histogram. The minimum value is 0.00, the maximum is 0.99. Higher values increase the contrast, lower decrease it.

As of RawTherapee 5.5, while the Auto Levels tool in the Editor still uses the Clip % value, the thumbnails use a fixed value of 0.2% - this was done to reduce the number of files needed to be stored in the cache.

Highlight Reconstruction

Use *Highlight Reconstruction* (HR) to try to restore overexposed highlights in raw files. It attempts to restore clipped (blown-out) regions in the raw image relying on the fact that the three channels in a raw file do not clip at the same time and so a missing (clipped) region in one channel could be guessed from the present data of one of the other color channels. Using the *Color Propagation* method, it can also guess clipped data using nearby data from unclipped channels, if present. Remember, this tool is used for the *reconstruction* of clipped highlights, while if you just want to *compress* highlights which were not clipped in the first place but became clipped due to the use of, for example, *Exposure Compensation*, then use the Highlight Compression slider. The *Auto Levels* button will automatically enable *Highlight Reconstruction* if necessary.

Four different methods of highlight reconstruction are available:

1. Luminance Recovery

Recovered details will be neutral gray.

"Highlight Compression" values under 100 are recommended in most cases.

2. Color Propagation

This is the most powerful recovery method. In addition to restoring luminosity, *Color Propagation* tries to restore color information by 'bleeding' the surrounding known color into the missing clipped area. This method works best on small overexposed areas, and can work wonders on overexposed skin. Its weakness is that it may sometimes 'bleed' the incorrect colors, depending on the image elements surrounding the blown highlights, or the colors can bleed into undesirable patterns. It is also computationally intensive and is therefore slower than the other methods.

Works well even with very high "Highlight Compression" values (under 500).

3. CIELab

Reduces the luminance channel and tries to restore colors afterwards. "*Highlight Compression*" values under 100 are recommended in most cases.

4. Blend

Attempts to guess clipped color channels by filling in their values from the closest match from unclipped highlight regions nearby.

"Highlight Compression" values under 100 are recommended in most cases.

Exposure Compensation

The values of the Exposure Compensation slider (EC) are ISO values. This means a value of +1 equals one stop of overexposure (+1 EV, exposure value; also known as +1 LV, light value)). If you make two photos, one without correction (EV = 0) and one underexposed by one stop (EV = -1), you can make both photos match by setting exposure compensation for the overexposed photo to -1, or for the underexposed photo to +1.

Take a look at the histogram while moving this slider. Moving it to the right shifts the whole histogram to the right. This means this slider changes the black point (on the very left of the histogram) and the white point (on the very right).

If you try decreasing the exposure of a photo which contains clipping, you will notice that the clipped areas turn flat-gray. Enabling "*Highlight Reconstruction*" will prevent this flat-gray look to some degree by recovering highlight information from the remaining non-clipped channels.

If you've used RawTherapee for a while you may notice that you almost always need a positive Exposure Compensation - as if all your photos are underexposed. Don't worry - this is normal. Most cameras deliberately underexpose to preserve highlights (though some other raw processors disguise it).

For the user interested in the internals follows a technical description of how EV=0 relates to the raw data: while EV=0 equals 1.0 in gain (that is, no change), there is a white balance dependent base gain applied before exposure. This base gain is calculated such that the color channel with the smallest range just can reach the maximum value. Although all raw channels have the same range in the raw file, white balance will re-balance them (lower temperatures mean more gain on red, higher more gain on blue) so they do not clip at the same level. The base gain raises the smallest channel just enough to make sure that at EV=0 no highlights will be based on clipped raw data. As white balance changes raw channel gains the base gain will thus change as a side effect when white balance is changed. For large white balance changes you may therefore see a slight change of brightness in the image. Note that the base gain relates to the maximum values channels can represent, that is if there is no highlight clipping in the raw file no highlight will reach the maximum level at EV=0.

Highlight Compression

The Highlight Compression slider (HC) can be used to recover the highlights in a photo by compression, useful for 'dimming' (or burning) slightly overexposed areas. While the Highlight Reconstruction tool lets you try to reconstruct missing data in the raw file relying on the fact that the three channels in a raw file do not clip at the same time and so a missing (clipped) region in one channel could be guessed from the present data of one of the other color channels, this *Highlight Compression* tool only works on data which is already there - which has not been irreversibly lost at the time of shooting. If the original image has no clipping, but due to the action of, for example, exposure compensation you caused clipping, then you can use HC to compress these clipped regions back into view. As such, it works not only on raw files but on normal images too.

To see if your photo contains overexposed areas, click on the clipped highlight indication icon ∧ on the top right of the image window. Overexposed areas will show up in black.

By dragging the *Highlight Compression* slider to the right, the intensity of the highlights will decrease. For highlight compression to work best you must also enable Highlight Reconstruction. Each of the HR methods has its strong and weak points, as explained above. Color Propagation is the most likely method to produce good-looking results when the HC slider is significantly over 100. For the other methods you will usually want to keep the HC slider around or below 100 - watch the histogram and the preview!

An overexposed photo. No highlight compression,

the clouds are clipped.

To find out the optimal HC value, you can make use of the histogram. In the screenshots above, what you see are overexposed clouds above the Teide volcano in Tenerife. When hovering the mouse cursor over the overexposed area, the pixel values indicator (in the Navigator panel, under the little preview) shows that lightness (L) is at 100, and the histogram shows that all channels are clipped (see the red, green and blue squares in the top-right corner of the histogram, they mean that there are so many pixels of maximum value that they are offscale). Increase the HC slider until the red, green and blue channels in the histogram no longer squash up against the right end of the histogram - you want them to touch it, but not to cram up against it. You can enable the *Clipped highlight indication* icon ∧ before moving the HC slider up. Once the indicator's black areas disappear from the white parts you want compressed, which is also when the lightness of those pixels drops from L=100 to L=99, you stop. Don't increase the HC slider any more, because now the hopelessly-lost white areas would start turning gray. You don't want them to turn gray. That would make the photo look dull. In this example the indicator's black areas disappeared when I set HC to 125.



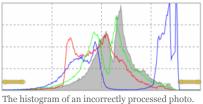
The Clipped highlight indicator has been turned on to make the clipped areas easier to see.



A correctly compressed photo - the histogram touches both ends.

As a rule, the histogram of a correctly developed image should touch both ends - the black and the white end. Not doing so means the image was incorrectly developed. This is true for the vast majority of photos, the only exceptions being photos of scenes which lack dynamic range, such as misty scenes. If you increase the HC slider too much, then whites turn gray as your histogram no longer touches the maximum end. Examples of photos that have been over-compressed can be easily found on the internet. They look horrible, don't do that! Recover what you can, but what's clipped beyond repair should stay white.

RawTherapee offers more ways to deal with blown highlights. The side-effect of all these methods is that they also take away some of the brilliance of the photos, as they get more 'flat' or 'dull' as a result. Highlight compression is very useful when used with moderation, but remember that you cannot recover what is not there to begin with, so once you notice that the completely clipped white areas become gray you should reduce the compression amount until these areas become white once more. To create the best possible output, feed RawTherapee the best possible input - so expose well in the first place!



The histogram of an incorrectly processed photo. It should touch both ends instead of having gaps.

Highlight Compression Threshold

The Highlight Compression Threshold slider sets the point where the HC slider starts implementing compression. A value of o means the threshold is zero: data compression occurs over the whole range of tonalities. 100 sets the threshold at one stop below the white point, so all the compressed highlights are squeezed into the top stop. In practical terms, more highlights are compressed when this slider is set to o.

Black

Use this to set the black point. See the left side of the histogram move when you touch the slider. Values greater than o will make the image darker, negative values will lighten up the shadow parts of the photo.

Shadow Compression

The *Shadow Compression* slider 'dampens' the effect of the *Black* slider. The maximum value of 100 gives a less dark image. This slider only has effect when the *Black* slider is set to a value other than 0. Practical use of this *Shadow Compression* slider is to fine-tune the shadow intensity of the image.

Lightness

This slider applies a hard-coded tone curve to lift or lower the tonalities of the photo, resulting in a more or less light image. The same tone curve is applied separately to each R, G and B channel. The black point and the white point keep their positions.

Contrast

This slider increases or reduces the contrast of the photo. It applies a contrast curve centered at the average luminance level. Tonalities above the average are lifted (lowered), while tonalities below the average are lowered (lifted). The same contrast curve is applied separately to each R, G and B channel.

Saturation

This slider makes the photo more or less saturated. In more technical terms, it adjusts the saturation of the image by applying a multiplier to the saturation level of pixels in the HSV color space.

Tone Curves

Here you can construct your own tone curves. They work on all three R, G and B channels at the same time (so you can't work on the R channel only).

There are two tone curves available, which can be designed with various curve types, and applied in several different modes all explained below. Clicking on the curve button hides the curve from the interface - it does not disable the curve.

The histogram displayed as the curve's background shows you the levels of the data as it flows into the curve at that point in the processing pipeline. You will notice that it differs from the main histogram which shows you the levels of the final image, at the very end of the pipeline.

While you are free to use only one tone curve to make your adjustments, you can gain much finer tonal control if you use two curves at once. The typical use of both curves is to lower values using the first curve, and to raise values using the second one. It is similar to creating an S curve in one of them, but you should be able to make finer adjustments by using both without entering too fast in the "danger zone" where your colors becomes unrealistic.

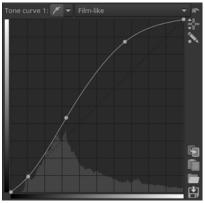
You can save a curve to disk. Click on the *Save current curve* button next to the graph and give it a name. Use the *Load a curve from file* button to apply this curve later to another file. Use the *Reset curve to linear* button to delete all the points you created and to reset the curve to neutral/linear. You can also *Copy* and *Paste* curves to/from RawTherapee's own clipboard, which comes in very useful if you want to quickly apply an identical curve to a different tool. Saved curves have the .rtc filename extension.

The curve used in some of RawTherapee's bundled DCP profiles is the same as the default curve used by Adobe Camera Raw. You can get the curve from this .rtc file. In the Exposure tool set either Tone Curve 1 or Tone Curve 1 to "Custom", set the curve mode to "Film-Like", and apply this file: File:DCP tone curve.rtc

You can use as many control points in a curve as you like.

You can define different curves for all curve types if you like, but only the one that is selected in the dropdown menu will be applied to the photo.

The curve and histogram is always displayed with sRGB gamma, regardless of working or output profile. This means that the shadow range is expanded and highlights compressed to better match human perception.



DCP (ACR) tone curve, DCP tone curve.rt

Linear Curve

This represents the unaltered (or linear) image, so without any tone curve applied. It disables the curve.

Custom Curve

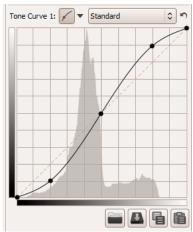
This is a classic type of curve, seen in many other programs as well. The left part of the graph represents the darker tones, the right part represents the brighter tones of the photo. Click on the curve to mark a point and drag it with the mouse to change tonalities. Dragging the point down makes the image darker, while pushing it up makes it brighter. The dotted diagonal line marks the linear or unaltered state of the photo. Press and hold the **Control** key to slow down the movement. Hold the **Shift** key to snap the point to key elements: maximum value, minimum value, middle value (i.e. snapped to the diagonal or

horizontal dotted line), same value of the preceding point, same value of the next point, and for the *Control Cage* type, the line going through the previous and next points. Delete a point on the curve by dragging it out of the editor area.

The top-right point represents the brightest areas in the photo. Drag that point vertically down to make the highlights less bright; move it horizontally to the left to make bright areas brighter, perhaps at the cost of some overexposure.

The bottom-left point represents the darkest areas in the photo. Move that point horizontally to the right to make the photo darker, perhaps at the cost of some underexposure. Move it vertically up to make the darks lighter.

Flip the diagonal line (from bottom-left and top-right to top-left and bottom-right) to produce a negative image.



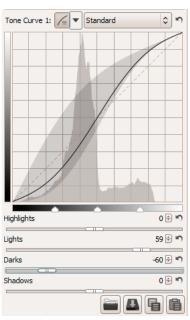
S-shaped custom curve.

A typical usage of the custom curve is to construct a so-called *S-curve*. Mark three points at the 'coordinates' (1,1), (2,2) and (3,3) respectively. Drag the point at (1,1) somewhat lower and the point at (3,3) a bit higher. Your image will get more 'punch' this way. If your S-curve is symmetrical, i.e. if you move the point you first placed at 1,1 by the same amount as the one you placed at 3,3 but in the opposite direction, then the effect will be identical to manipulating the *Contrast* slider.

Parametric Curve

This curve presents four sliders and three control points. The sliders are used to control highlights, lights, darks and shadows respectively (shadows mean deep darks here). Move the mouse over the four sliders and a dark area under the curve tells you which slider alters what part of the curve. Move the Highlights slider to the left to make highlights less bright, move it to the right to make them brighter. The *Lights* slider moves the lights but not the highlights in the same way as above, as does the *Darks* slider: moving it to the right lightens the dark tones, moving it to the left darkens them. The *Shadows* slider works as the *Darks* slider, but only on the darkest parts of the photo. You can again construct the above mentioned stylized S-curve, although the parametric curve gives you less 'extreme' control over the form of the curve. This mode, however, has its own benefits, as curves can be shaped in a well controlled manner. Note that using these sliders can have a profound influence on the overall contrast of the image.

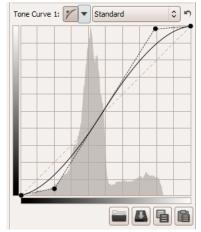
Use, if needed, the three control points under the curve. They determine what point of the curve will be affected when moving the sliders. Moving the middle control point to the right makes the image darker (the form of the curve changes again, as does the dark area around the curve), moving it to the left makes the image brighter. Moving the left control point to the right darkens the dark areas somewhat, moving it to the left lightens them, again somewhat. Moving the right control point to the right makes the highlights brighter, moving it to the left darkens the highlights.



Use the *Reset to default* buttons next to the sliders to reset individual sliders, use the same button at the top of the tone curve section to reset all four sliders and the control points to linear (zero).

Parametric curve.

Control Cage



Control Cage curve.

At first sight this curve type looks very much like the *Custom* curve, but there are some differences though. With the *Custom* curve, the curve touches all the control points. This is not the case with the *Control Cage* curve. To see this, click somewhere on the line and move the black point to the left or to the right. Now the curve passes nearby the black point, but doesn't touch it. Another difference is that the *Control Cage* allows for a straight section of the curve, while you can't do this with the *Custom* curve. The *Control Cage* curve needs at least three points for that (so five in total). Holding down the **Shift** key while dragging a point will help you to easily create a straight line by snapping the point to the line made by the previous and next point (displayed in red by the *Snap to* tool). Now make a new point between the two most left ones and move it. As you can see, only the part on the left side moves, not the rest of the curve.

Curve Mode

Next to each curve type, you'll find a *Curve Mode* combobox selector. This will let you choose the algorithm that will be used for the corresponding curve. The curve mode will have a strong effect on the appearance of colors, especially if you use a contrast-enhancing curve (S-curve). This can be used for creative effect, but can for some purposes or styles cause undesired color changes depending which mode you choose. Choose a mode that suits your specific taste and needs for the photo at hand. By combining two different curves in tone curve 1 and 2 you can further fine-tune the look.

The choice of working profile has an influence on the effect of the curves in all modes except for perceptual - in that mode, changing the working profile will not alter the effect of the curve.

Standard

This is the most basic mode (and the only one available in older versions of RawTherapee and is found in some shape or form in most image-related software): the values of each RGB channel are modified by the curve in a basic "correspondence" method, that is the same curve is applied to all channels.

The drawback of this mode is that e.g. considering an S-curve shape to get more contrast, an orange color with a high value of red and green and a low value of blue will tend to shift toward yellow, because the red and green component will be raised, while the blue one will be lowered.

In general an S-curve will increase separation of the channels and thus increase saturation, which is a similar behavior to how color film reacts to contrast. This together with the simplicity of implementation has made the curve type popular in raw converters in general and is often the only alternative available in less flexible software.

Weighted Standard

You can use this method to limit the color shift of the standard curve, even if it won't suppress it entirely. Keeping the previous example, this method will raise the first component (red), and will also linearly alter the green and blue component by raising them too. We end up with 3 values (R, g and b) while we have only processed the red component.

This process is then done for the green and blue component, so at the end of the process, we end up with 9 values (R,g,b / r,G,b / r,g,B). Values of the same component are then mixed together, which will produce the resulting color with a smaller color shift.

Film-Like

The film-like curve provides a result highly similar to the standard type (that is strong saturation increase with increased contrast), but the RGB-HSV hue is kept constant - that is, there are less color-shift problems. This curve type was designed by Adobe as a part of DNG and is thus the one used by Adobe Camera Raw and Lightroom.

Saturation and Value Blending

This mode is typically better suited for high-key shots, but can be used for creative effect in other photos as well. The average value of the three component is computed, and then the curve is applied to this value, giving a positive or negative gain. The color is converted to its Hue, Saturation and Value representation, then if the gain is positive, the pixel is linearly targeting Value = 1 and Saturation = 0, the Hue is preserved. If the gain is negative, the pixel is linearly targeting Value = 0, Saturation and Hue are preserved.

The result is highly similar to a luminance curve in Lab space (that is change contrast without affecting hue or saturation). For contrast-increasing curves the look will typically be slightly desaturated. This is not really because the curve desaturates the colors but because that in human vision contrast and saturation is tightly coupled, so the same image with higher contrast requires higher saturation to appear to have the same.

Luminance

Each component of the pixel is boosted by the same factor so color and saturation is kept stable, that is the result is very true to the original color. However contrast-increasing curves can still lead to a slightly desaturated look for the same reason as described for the Saturation and Value Blending curve mode. If you want to manually counter-act the desaturation, using the L*a*b* Chromaticity slider is a more neutral way of compensating for it than using the RGB-based saturation slider.

Despite showing the R, G and B histogram (merged) in the background of the curve, the curve operates on luminance values, where Relative Luminance (https://en.wikipedia.org/wiki/Relative_luminance) Y = R*0.2126729 + G*0.7151521 + B*0.0721750 First the relative luminance value of a pixel is obtained, then the curve is applied to that value, the multiplication factor between before and after luminance is calculated, and then this factor is applied to each R, G and B component. This is in contrast to the other methods where the curve is applied to each R, G and B component separately.

Perceptual

This mode will keep the original color appearance concerning hue and saturation, that is if you for example apply an S-curve the image will indeed get increased contrast, but the hues will stay the same and the image doesn't look more or less saturated than the original. It's specifically useful to establish a pleasing baseline contrast without distorting the colors provided by a camera profile which doesn't apply a curve itself (if you use a third-party profile that does apply a curve it's typically already perceptually mapped with similar techniques as described here).

The choice of working profile has an influence on the effect of the curves in all modes except for perceptual - in this mode, changing the working profile will not alter the effect of the curve.

The algorithm works the following way: it analyses the curve to get a contrast value, which is used as base to scale chroma (saturation) such that more contrast leads to more saturation and the other way around. As contrast and saturation is tightly coupled in human vision this scaling is necessary to make saturation *appear* constant. There are further fine-tunings such as increase saturation more in the shadows, and less for colors that are already highly saturated, also this corresponds to human vision phenomena so the net effect it that the colors appear constant. In the extreme highlights close to the white point the algorithm blends over to white (like the standard curves) which is less true to color but more practical for real output as the brightest color of the output media (screen or paper) is white.

However do keep in mind that the perceptual model is not perfect and cannot be perfect. This is only a curve, image content is not analyzed and no localized changes are made. This means for example that for an S-curve a large flat blue sky (low local contrast) may appear slightly more saturated than the original. If you want to make A/B comparisons don't compare side by side as the eye will then be confused by the two contrast levels viewed simultaneously and then saturation will not appear the same, but instead swap and let the eye adapt for a few seconds.

If you want to further fine-tune the saturation manually it's generally best to use the L*a*b* chroma tools.

Due to the many components in the algorithm it's considerably slower than the other curve modes so refresh rate may suffer.

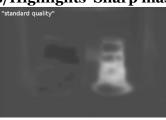
Shadows/Highlights

Use this tool to raise the shadows or darken the highlights of an image.

Sharp Mask

Shadows/Highlights 'Sharp mask' effect







The source image.

"Sharp mask" turned off.

"Sharp mask" turned on.

In order to separate the dark areas from the light ones, a lightness mask (invisible to the user) is created. There are two algorithms for doing this; one blurs the image, while the other retains sharp edges between light and dark zones. Neither one is "better", both have their own merits. The soft mask approach can lead to halos, but it is quick. The sharp mask is slow, but it doesn't cause halos, though it can cause edge artifacts at close inspection.

Shadows

Allows you to brighten the darkest parts of the image

Highlights

Allows you to darken the brightest parts of the image.

Tonal Width

Shadows/Highlights Tonal Width allows you to control how bright an area must be for it to be affected by the highlights slider, and how dark an area must be for it to be affected by the shadows slider. Though the underlying math is a bit more complicated, think of a histogram - the highlights tonal width specifies the range of tones from the white end of the histogram which the highlights slider will affect, and the shadows tonal width specifies the range of tones from the black end of the histogram which the shadows slider will affect. The higher the tonal width value, the more tones are affected.

Radius

The value of the Radius slider influences the effective area of the Shadows and Highlights sliders.

Tone Mapping

The effects of this tool are visible at any zoom level. However, due to the nature of the algorithm, only the 1:1 (or more) preview zoom will match the saved image perfectly. If you are zoomed out at less than 1:1, you should be aware that the preview can match the saved image very well or not so well, depending on the "Edge Stopping" and "Scale" sliders. Read the "Getting the preview to match the saved image" section below. Use a detail window (click on the "icon under the main preview panel) to inspect a part of the image, or zoom the main preview to 100% (also called 1:1) [1].

The Tone Mapping tool can be used to lift the dark areas of your photo in a way that prevents halos from appearing, and it can be used to bring out or suppress detail, to make the photo more crisp or more 'dreamy'. Tone mapping adjusts the global contrast of an image differently from the local contrast. Specifically, it's very useful for decreasing large scale contrasts while preserving (or boosting) small scale contrasts. The method used is taken from *Edge-Preserving Decompositions for Multi-Scale Tone and Detail Manipulation* with some modifications.



Tone mapping was used on this high dynamic range (HDR) 32-bit TIFF image to make it look aesthetically pleasing and viewable on low dynamic range media, such as a computer screen or on paper.

Note: tone mapping requires a lot of memory (RAM) and is CPUintensive.

Getting the preview to match the saved image The effects of this tool highly depend on the size of the input image in the engine's toolchain. To keep the preview fast, RawTherapee feeds The effects of this tool highly depend on the size of the input image in the engine's toolchain. To keep the preview fast, RawTherapee feeds

the engine's toolchain. To keep the preview fast, RawTherapee feeds each tool with the image you see in the preview at the same resolution as the preview, not the original, huge one (however, when you save an image, it uses the original, huge one, which is why saving takes longer). Processing a 900x600 image is much faster than processing a 7360x4912 one, for instance. The side-effects of this are that the zoomed-out preview may not match the saved image, depending on the "*Edge Stopping*" and "*Scale*" sliders.

The default values for this tool are Edge Stopping=0.5 and Scale=0.10. These values generally lead to good results and a preview which quite closely matches the saved image. If you wish to adjust these values, you will need to process the full image every time you change them to make sure that their effects on the saved output are what you had in mind. An easy way to simplify this task is by setting an image viewer as the "External editor", then you just hit the keyboard shortcut Ctrl+e



Edge Stopping=1.40 and Scale=0.10 while editing, but Scale is set to 1.00 just before saving. Using this trick, the preview matches the saved image very well.

to have RawTherapee fully process the image and automatically view it in your image viewer.

Remember that if you zoom the preview to 100%, it will perfectly match the saved image regardless of what slider values you use.

Interface Description

Strength

This controls the strength of the overall effect. As of version 4.2.156, as you increase the strength, shadows are raised and highlights are slightly lowered in a way which aims to preserve the average luminance of the image thereby preventing the washed-out look.

Gamma

Gamma moves the action of tone-mapping to shadows or highlights.

Edge Stopping

This parameter affects sensitivity to edges. The greater it is the more likely an illumination change is to be considered an "edge". If set to zero tone mapping will have an effect similar to unsharp masking (https://en.wikipedia.org/wiki/Unsharp_masking).

Scale

This control gives meaning to the difference between "local" and "global" contrast; the greater it is the larger a detail needs to be in order to be boosted. See the note above on the influence this setting has on the preview vs the saved image.

Reweighting Iterates

In some cases tone mapping may result in a cartoonish appearance, and in some rare cases soft but wide halos may appear. Increasing the number of reweighting iterates will help fight some of these problems. When more than zero reweighting iterates are used, the best results will be had if the edge stopping parameter is set to one (technical detail: this results in a 1-norm approximation of the smoothness using iteratively reweighted least squares). Artifacts in high contrast edges may begin to appear when this value is set to about 1.8 or higher.

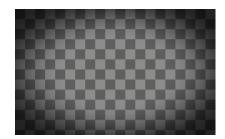
HDR Tone Mapping

HDR Tone Mapping

Vignetting Filter

The vignette filter is intended for adding artistic vignetting to your image. This vignetting filter is placed relative to the crop, if cropping is used.

For correcting vignetting caused by the lens light fall-off (as opposed to this filter which is not for correction but for artistic effect), use the Vignetting Correction filter in the Transform tab, in the Lens/Geometry tool. Even better, use the Flat Field tool.



Strength

The amount of darkening the filter will apply, in stops. The full strength is reached in the corners of the image. If you apply a negative amount the corners will be brightened instead of darkened.

Feather

The feather slider controls the width of the feathering. If at o only the corners will be feathered and the rest of the image will not be affected by the filter. At 50 the feather reaches halfway to center and the rest is unaffected, and at 100 the feather reaches all the way into the center.



feather=o

feather=100

Roundness

The roundness slider controls the geometry of the filter. At 0 the shape is rectangular (with rounded corners), at 50 it is a fitted ellipse, and at 100 it's circular. Note that if your image is square the fitted ellipse will of course be a circle, so the shape will then not change in the range 50 to 100.



roundness=0

roundness=100

Graduated Filter

The graduated filter tool simulates a real neutral density graduated filter. These can be used used in for example landscape photography to limit the brightness of the sky. It applies an exposure adjustment in the form of a gradient.

Strength

The strength of the filter, in stops.

Angle

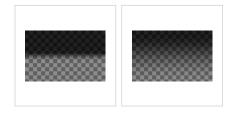
The angle slider allows to rotate the direction in which gradient is applied.



angle=45

Feather

The feather slider controls the width of the filter's feather (gradient). At o there is no feathering, just a sharp edge which can be useful as a temporary setting when choosing position and angle of the filter. The slider value (o - 100) specifies the width of the feather as percent of the image diagonal.



feather=10

feather=50

Center X and Center Y

The Center X and Center Y adjustments moves the center of the filter's feather (ie the rotation anchor point) horizontally and vertically respectively.

Lab Adjustments

Lab (https://en.wikipedia.org/wiki/Lab color space) (also called CIELAB or L*a*b) is a three dimensional color space designed to approximate human vision, as opposed to the RGB color space which models the output of physical devices rather than human visual perception. It keeps the tone (also called lightness or value) separate from the color, so that you can adjust one without changing the other.

- The L component closely matches human perception of lightness.
- The a component defines how green/magenta the color is.
- The b component defines how yellow/blue the color is.

Lightness

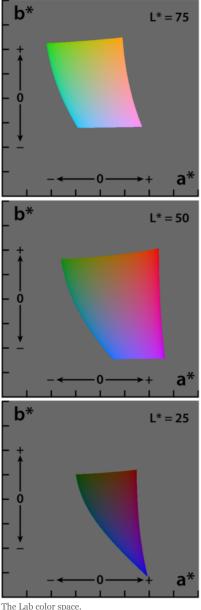
When using the Lightness slider in the Lab section, a tone curve is applied to the L-channel of the Lab color space. As with the brightness slider in the Exposure section above, the black point and the white point do not move.

Contrast

The contrast slider in Lab increases or decreases the contrast of the photo, again applied to the L-channel. In developer's terms: this slider applies a contrast curve centered at the average lightness level. Tonalities above the average are lifted (lowered), while tonalities below the average are lowered (lifted).

Chromaticity

The Lab Chromaticity slider increases or decreases the chromaticity of the image, by applying a contrast curve to the a- and b-channels of Lab space. Setting this slider to -100 removes all color, making the image black and white. The best way to convert an image to blackand-white is by using the dedicated and powerful Black-and-White tool in the Color tab.



The Lab color space.

B&W Toning

The "B&W toning" checkbox is deprecated from version 4.0.12 and is replaced by the Black-and-White tool located in the Color tab. For backwards compatibility, when opening processing profiles where "B&W toning" was used, the *Chromaticity* slider will get automatically set to -100, providing the same effect.

Avoid Color Shift

Fits the colors of the image into the gamut of the working color space and applies Munsell correction (http s://en.wikipedia.org/wiki/Munsell_color_system) to retain color purity.

Restrict LC to Red and Skin Tones

When enabled, it restricts the effects of the *Lightness According to Chromaticity (LC)* curve, so that you can make skin fairer (by increasing the lightness of the skin) without affecting the model's clothing or background.

Red and Skin Tones Protection

When enabled, the effects of the Chromaticity slider and the CC curve will not be applied to skin colors, so that you can increase the chromaticity of your photo without causing skin to appear oversaturated.

Curves

Lab Adjustments provides a wealth of curves to alter the look of the image. Below are illustrated explanations of each curve.

L Curve



The L curve was used to improved the lightness without causing any changes to the color.

The L curve allows to control output lightness based on the input lightness, L=f(L). The histogram on the L curve reflects lightness prior the Lab adjustments. This curve allows you to control the lightness without affecting color.

An S-shaped curve applied to the L channel increases image contrast. At the same time this leads to a perceptually desaturated look. Chromaticity adjustments can be used to compensate for this effect.

"a" and "b" Curves

The "a" and "b" curves allow to control output "a" and "b" channels based on the input "a" and "b" channels respectively, a=f(a) and b=f(b).

As indicated by the color bars, the "a" curve allows one to shift colors between green and magenta, and the "b" curve to shift between blue and yellow. This can be used to apply color toning effects.

Black-and-White Color Toning

Color-toning a black-and-white image can be done using one of two methods: the recommended and most intuitive method is by using the Color Toning tool along with the Black-and-White tool. The other, less powerful method is using the a* and b* curves of the L*a*b* Adjustments tool once the image is desaturated. The reason we still describe how to do it without using the Black-and-White and Color Toning tools is that these tools are relatively new additions to RawTherapee, and maybe you're stuck using an older version which lacks these tools, or you're just curious what your options are.

Read about color-toning a black-and-white image the recommended way on the Color Toning tool's page; this section describes how to do it by using the a* and b* curves.

First you need to make the image black-and-white. Do this using any of the available methods: either using the Black-and-White tool, or by decreasing saturation in the Exposure tool, or by decreasing chromaticity in the L*a*b* Adjustments tool. Each tool will lead to a different effect since they work in different ways and in different color-spaces. It's just a matter of taste. Once the image is reduced to

grayscale you can give the image a tone by using the a* and b* curves. To copy just color toning from one image to another, copy the current processing profile to clipboard \blacksquare , then partial-paste it either by right-clicking on a photo in the *File Browser* and selecting "*Processing Profile Operations > Paste - partial*", or from the *Image Editor* tab by Ctrl+clicking on "*Paste profile from clipboard*" \blacksquare to paste only the $L^*a^*b^*$ *Adjustments* section of the profile. Note that other adjustments in the $L^*a^*b^*$ *Adjustments* sections will be pasted as well. Alternatively, the a* and b* curves can be copied and pasted individually. This is another reason for using the recommended method, because it's easier, more precise, to copy and paste the Color Toning and Black-and-White tools.

LH Curve



The LH curve was used to change the lightness of the chosen colors.

The LH curve (lightness according to hue) allows to modify the lightness based on hue. To lighten the colors of the particular hue, move the desired point on the LH curve up, and to darken - down.

CH Curve



The CH curve was used to desaturate all colors except shades of red.

The CH curve (chromaticity according to hue) allows to control output chromaticity based on the input hue, C=f(H). Using it you can very easily boost or mute only a selected range of colors.

HH Curve



The HH curve was used to change the hue of shades or red into shades of purple.

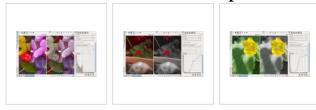
The HH curve (hue according to hue) allows to alter the hue for a specified hue. For example, one could shift reds to be more orange by moving the red point up until the thick horizontal line that appears as the point is being dragged becomes the color you desire. RawTherapee has two HH curves, one in the Lab tools in the Exposure tab, and one in the HSV tool in the Colors tab. The HH curve in the Lab tools has a more restricted range compared to the HH curve in the HSV equalizer, to allow finer adjustments. The range is between the previous and next color, e.g. green could be changed

within the range of yellow and blue (as you can see in the curve on the screenshot above). This is useful, for example, for fine-tuning skin tone appearance, removing a greenish pale look by shifting reds and yellows a little towards magenta.

CC Curve

The CC curve (chromaticity according to chromaticity) allows to control output chromaticity based on the input chromaticity, C=f(C). The histogram on the CC curve reflects chromaticity before the adjustment. This allows you to separately adjust the chromaticity of pixels of low and high saturation, so you can boost saturation where needed without causing already saturated zones to clip.

CC curve effect examples



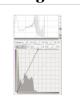
Boost low chromaticity, mute high chromaticity. Mute low chromaticity.

Mute low chromaticity.

You can use the Show/Hide chromaticity histogram button • besides the histogram to help you see the effects of your CC curve tweaks on the histogram, and to help you find the maximum value before you start clipping colors.

Chromaticity histogram





Smooth chromaticity histogram with neutral CC curve.

Spiked chromaticity histogram with too strong CC curve.

The screenshots show what the chromaticity histogram looks like for the untouched image, and then what happens when you increase chromaticity too much (you can do this using the Chromaticity slider, or, as in the screenshot, by sliding the top-right point of the CC curve to the left. Holding the Shift key while you slide the point will help you keep the point at the top).

To find the maximum chromaticity boost you can apply without causing nasty spikes, which will appear as sudden flat regions of color in the image, similar to posterization, all you need to do is click Show/Hide chromaticity histogram if you haven't done so already, and then slowly boost chromaticity until you notice the histogram begins to spike. The curve does not have to be linear of course.

LC Curve

The LC curve (lightness according to chromaticity) allows to control the output lightness based on the input chromaticity, L=f(C). You can use it on portraits to lighten skin

LC curve effect examples





Lighten lowchromaticity zones.

Lighten skin tones.

The LC curve's action is modulated by the *Restrict LC to red and skin tones* checkbox. Thus the LC curve provides a complex image control, altering lightness based on image chromaticity and also targeting a specified range of hues. With this option enabled, the lightness of only red and skin tones is affected, for example allowing you to make skin fairer and conceal wrinkles and blemishes while preserving the color of the model's clothes and background. When it is disabled, the LC curve acts on other colors as well.

The coloring of the bar on the horizontal LC curve axis changes to reflect which colors the curve applies to, as chosen by the state of the Restrict LC to red and skin tones checkbox.

CL Curve

The CL curve (chromaticity according to lightness) allows to control the output chromaticity based on the input lightness, C=f(L). It allows you to separately control the chromaticity of regions of the image based on their lightness, so you can for example decrease the chromaticity of shadows if they are noisy or for artistic purposes, or increase the chromaticity of dark and mid-tones without affecting the bright sky.

CL curve effect examples





Increase chromaticity of light areas without saturating

shadows.

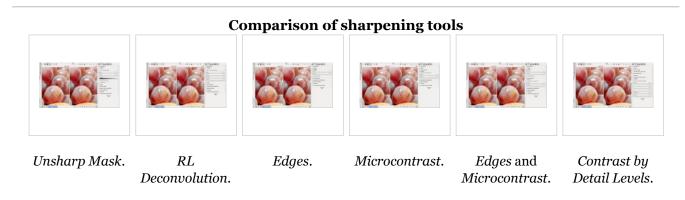
Chromaticity of dark and midtones increased without saturating the

sky.

Detail Tab

Sharpening

The effects of this tool are only visible at a preview scale of 1:1 or more. Use a detail window (click on the icon under the main preview panel) to inspect a part of the image, or zoom the main preview to 100% (also called 1:1) .



This article describes the tool called "Sharpening", however RawTherapee contains other tools which can be used to perform various types of sharpening - see Edges and Microcontrast and the Wavelets tools.

The Sharpening tool is applied to the full image, **before** the Resize tool. If you would like to apply sharpening **after** resizing, use the Post-Resize Sharpening tool which you will find inside the Resize tool.

Unsharp Mask

Unsharp masking (https://en.wikipedia.org/wiki/Unsharp_mask) (USM) is a technique used to increase the apparent acutance (https://en.wikipedia.org/wiki/Acutance) (edge contrast) of an image, making it appear clearer, even though it technically does not really sharpen the image. It makes use of several phenomena of the human visual system in order to accomplish this effect, such as the Cornsweet illusion (https://en.wikipedia.org/wiki/Cornsweet_illusion) and Mach bands (https://en.wikipedia.org/wiki/Mach_bands). Though unsharp masking in other software is easily prone to causing halos (https://en.wikipedia.org/wiki/Haloing), RawTherapee has a unique threshold slider which allows you to achieve a superb sharpening effect with a minimal risk of halos.

Radius

The Radius determines the size of the details being amplified and, consequently, relates to the width of the sharpening halo. In general the quality of sharpening is best if the sharpening radius is smaller. For low ISO images that are in focus and without motion blur a value of 0.5-0.7 is satisfactory.

Amount

The Amount parameter controls the strength of the sharpening.

Threshold

The *Threshold* tool helps to suppress noise amplification and haloing and to confine sharpening to a desired tonal range. The Threshold tool allows one to create a curve via which the sharpening is applied. The vertical axis corresponds to opacity: 0% at the bottom (transparent, sharpening not visible), 100% at the top (opaque, sharpening visible). The horizontal axis corresponds to luminosity: select the tonal range that will get sharpened - the darkest tones are on the left, progressing to white tones on the right. As mentioned in the tooltip, to move each of the points in the threshold tool individually, hold the Shift key before clicking on a point with your mouse. Holding the Ctrl key while moving a point with the mouse allows for very fine movements.

When moving the right pair of sliders to the left side, sharpening is reduced in the highlights. When moving the left pair of sliders to the right side, sharpening is reduced in the shadows and minimizes amplification of dark noise.

The default threshold values will protect from over-sharpening and haloing in most cases and limit the sharpening effect to mid-tones. In the example screenshot, the blackest tones have no USM applied, then USM is applied to a broad range of tones from dark to light, and the strength of USM gradually drops off from maximal at the mid-tones to none at the whitest tones, so as to prevent noise amplification and haloing.

Sharpen Only Edges

If you activate "Sharpen only edges" then uniform areas will not be sharpened. This is useful when sharpening noisy photos.

Two new sliders appear as well:

Radius

The *Radius* is used for noise detection. If the noise is low, a lower radius can be used, and vice-versa. A higher radius slows down the image processing.

Edge Tolerance

Edge Tolerance determines how much a pixel has to differ from its neighbor to be considered as an edge and not as noise. It is very similar to the USM *Threshold* parameter and has a high impact on the visual quality. For low ISO (low noise) images use 1000 or less, for high ISO images use 2500-3000 or even more.

Halo Control

Halo Control is used to avoid halo effects around light objects when sharpening too aggressively. When activated, a new slider appears:

Amount

At 100 it works at maximum, reducing the visual impact of the USM filter.

RL Deconvolution

RL deconvolution (https://en.wikipedia.org/wiki/Richardson%E2%80%93Lucy_deconvolution) is named after the makers of this algorithm, Richardson and Lucy. It uses the point spread function (https://en.wikipedia.org/wiki/Point_spread_function) (PSF) to deconvolve (reverse) the effects of Gaussian-like blur. In reality, the blur produced by the lens and by motion may differ from Gaussian blur significantly, therefore some artifacts, such as halos, may appear when the radius diverges too far from the type of blur in the actual image, and when then effect is too strong.

Radius

The radius defines the standard deviation (https://en.wikipedia.org/wiki/Standard_deviation) (sigma) of the Gaussian blur (https://en.wikipedia.org/wiki/Gaussian_blur) in the image. Find the right value through trial and error.

Amount

Controls the blend factor between the unsharpened image and the sharpened one.

Damping

Damping reduces the effect of the deconvolution at each iteration. It has the effect of preventing sharpening of the finest details. Use it if the sharpened image has too much "bite" at the finest level.

Iterations

RL Deconvolution is an iterative algorithm; it requires being repeated in order to achieve the intended results. Each repetition of the process is called an "iteration", and the result of one iteration is used as the starting point of the next iteration. While each iteration removes blur, it also increases processing time and the likelihood that halo artifacts will appear, so you need to find the perfect balance through trial and error - the default value should be fine for most cases.

Local Contrast

Local Contrast

Edges and Microcontrast

The effects of this tool are only visible at a preview scale of 1:1 or more. Use a detail window (click on the icon under the main preview panel) to inspect a part of the image, or zoom the main preview to 100% (also called 1:1) .

Comparison of sharpening tools













Unsharp Mask.

RL Deconvolution.

Edges.

Microcontrast.

Edges and Microcontrast.

Contrast by Detail Levels.

Unlike *Unsharp Mask*, *Edges* is a true sharpening algorithm. It does not introduce halos, it can be used to some degree on noisy images, and it works in the Lab space. *Edges* sharpens just the edges, and it can be complimented by *Microcontrast* to enhance texture.

Edges

Iterations

How many passes the algorithm makes. Higher numbers produce a more posterized effect. **Quantity**

Quantity

How many adjacent pixels will be searched for an edge. Larger values lead to sharper edges.

Luminance only

Sharpens the L* channel only; a* and b* channels are untouched.

More information here:

https://web.archive.org/web/20110625093654/http://www.rawness.es/sharpening/?lang=en(https://web.archive.org/web/20110625093654/http://www.rawness.es/sharpening/?lang=en)

Microcontrast

Quantity Uniformity

A 3x3 matrix is better suitable for noisier images.

Impulse Noise Reduction

The effects of this tool are only visible at a preview scale of 1:1 or more. Use a detail window (click on the icon under the main preview panel) to inspect a part of the image, or zoom the main preview to 100% (also called 1:1) .

Suppresses salt-and-pepper noise - sudden white and black pixels, which remind one of salt and pepper sprinkled over a photo. This is done after demosaicing.

Whereas salt-and-pepper noise is typically just white or black, hot pixels can be of a pure, saturated color, while dead pixels are black. Hot and dead pixels occur for a very different reason than salt-and-pepper noise and should be handled using the Hot/Dead Pixel Filter, which works before demosaicing.

The slider adjusts the threshold which must be exceeded for the suppression to be applied.

Noise Reduction

The effects of this tool are only visible at a preview scale of 1:1 or more. Use a detail window (click on the icon under the main preview panel) to inspect a part of the image, or zoom the main preview to 100% (also called 1:1) .

Introduction

Photography is based on recording light which falls on a medium during an exposure. The medium is typically film or a digital sensor. The light signal *recorded* on the medium is not an ideal representation of the signal which *fell upon* that medium - these differences constitute noise. Film and digital photographs alike are susceptible to noise (called "grain" in film), however there are various types of noise from various sources specific to each medium.



Before and after noise reduction.

In order to effectively deal with mitigating noise it is useful to understand what types of noise exist and where they come from. The topic of noise is well explained in this paper by one of RawTherapee's developers and physics professor at the Enrico Fermi Institute, Emil J. Martinec: Noise, Dynamic Range and Bit Depth in Digital SLRs (http://rawtherapee.com/mirror/noise/)

Light consists of packets of energy called photons. A digital sensor (https://en.wikipedia.org/wiki/Image_sensor) comprises of millions of light-sensitive elements called *photosites* (also known as *sensels* - sensor elements). Each photosite is capable of recording a signal from a certain range of photons - too few and the photosite will not register anything; too many, and the photosite will "clip" to pure white (completely overexposed). Think of it as a bucket collecting water - despite there being moisture in the air, if it doesn't rain then the bucket is empty, but if it rains too much then it overflows.

Note that the idea of "pixel" does not exist at this point yet - information from several photosites will later be combined into one pixel during a process called demosaicing. Also note that some sources do write "pixels" when they mean "photosites".

The physical sensitivity of the sensor is constant, however the photographer can amplify the recorded signal by modifying a setting you know as ISO (see the Film Speed (https://en.wikipedia.org/wiki/Film_speed) article on Wikipedia). Since the signals recorded by the sensor are not perfect, using a higher ISO amplifies not only the desired signal, but also the noise. Sensors are susceptible to noise at every ISO level, however the higher the ISO the more apparent the noise.

There are different tools for dealing with different types of noise:

■ The Noise Reduction tool is best at dealing with photon shot noise (Gaussian (https://en.wikipedia.org/wiki/Gaussian_noise) and Poisson (https://en.wikipedia.org/wiki/Shot_noise) and film grain, and some

sensor read noise.

- Sensor read noise and thermal noise are best handled by the Dark-Frame tool.
- Salt and pepper noise (sudden white or black pixels) is best handled by the Impulse Noise Reduction
- Hot and dead pixels are best dealt with using the Hot/Dead Pixel Filter.
- Pattern noise (periodic, anisotropic) is best handled by the Line Noise Filter. You can also fix pattern noise (de-screen) after RawTherapee in GIMP, using the Fourier transform in G'MIC.

Regardless the source, noise will manifest itself as blotches of deviating color - "chrominance noise", and deviating brightness - "luminance noise".

- 1. Chrominance noise is endemic to digital images, it is generally unattractive and something you will always want to remove.
- 2. Luminance noise, on the other hand, looks like film grain and can be attractive, so it's not uncommon to want to remove chrominance noise but keep luminance noise.

Examining Noise



A noisy test photo taken at ISO 6400.



AMaZE demosaicing leads to small maze-like patterns.



LMMSE demosaicing avoids maze-like patterns while preserving detail.



This is what chrominance noise looks like. Luminance detail was obliterated to make the chrominance noise more clear. Notice the color blotches in what should be a smooth wall.



Enabling chrominance noise reduction eliminates the colored blotches.



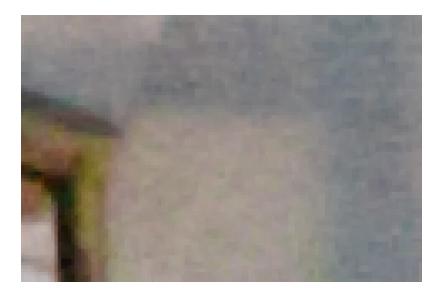
This is what luminance noise looks like. Chrominance noise was removed to make the luminance noise more clear.



Both chrominance and luminance noise were removed.



Tiny pixel-sized artifacts are left-over from noise reduction.



These artifacts can be removed using the median filter.

Not everyone's requirement for good noise reduction is the same. Some like a completely clean, smooth result, while others prefer to have some grain left over to give the photo a more film-like quality. RawTherapee's powerful *Noise Reduction* tool caters to all your needs - it lets you eliminate noise while retaining detail. It uses wavelets (http://en.wikipedia.org/wiki/Wavelet), a Fourier transform (https://en.wikipedia.org/wiki/Fourier_transform) and a median filter (https://en.wikipedia.org/wiki/Median_filter) to work its magic. Read on to learn how to use it efficiently.

Usage

This section details the order of operations for removing noise.

- 1. Start by ensuring you're using the optimal demosaicing algorithm. AMaZE is recommended for general RawTherapee use, however, when working with very noisy, high-ISO images, it is recommended to use the LMMSE or IGV demosaicing methods instead. AMaZE can lead to tiny maze-like artifacts appearing in very noisy images, whereas LMMSE and IGV are designed to prevent that from happening.
- 2. Check the sharpening tools to make sure that you're not sharpening any fine detail, because your noisy photo has no fine detail! If you're using Contrast by Detail Levels or Wavelets, make sure that the first one or two fine-detail contrast sliders are at o to prevent these tools from amplifying noise.
- 3. Zoom into the photo to 100% or more and find an area that has both sharp, in-focus parts as well as large, plain or out-of-focus ones, so that you have a good overview of the effects of the tool.
- 4. Enable the Hot/Dead Pixel Filter if you notice salt-and-pepper noise (black and/or white pixels).
- 5. Enable the Noise Reduction tool. Chrominance noise is automatically removed and usually does not require any tweaking. At this point the remaining noise looks more like film grain. If you are happy with keeping it then you are done, else keep reading.
- 6. To remove luminance noise, set the *Detail recovery* slider to 0, and increase the *Luminance* slider until the noise has been smoothed-away.
- 7. Increase the *Detail recovery* slider until you regain a satisfactory level of detail.
- 8. You may notice some small artifacts remain from the noise reduction process. Use the Median filter to remove them.
- 9. While it is generally not recommended to combine sharpening with noise reduction, RawTherapee-5.5 has a "contrast threshold" adjuster in the Sharpening tool, thanks to which you can sharpen details while preserving the smoothness of uniform, flat areas.

Noise Reduction Step by Step



The noisy image.



Changing the demosaicing method to LMMSE eliminates the small maze-like patterns and makes the salt-and-pepper noise more clear.



Enabling the Hot/Dead Pixel Filter eliminates the salt-and-pepper noise.



 $\label{lem:eq:chromaticity} \begin{tabular}{ll} Enabling automatic chromaticity noise reduction renders a \\ pleasantly-grainy image. \end{tabular}$



Luminance noise was smoothed-away using the Luminance slider.



Detail was restored using the *Detail recovery* slider.



The median filter was used to eliminate left-over artifacts.



Sharpness was restored using an unsharp-mask with a contrast threshold to prevent sharpening areas which should be smooth.

Interface

General

Closely examine large areas of strong saturation with fine detail - such as the pattern on a colored shirt or the petal of a flower - as you switch between the RGB and $L^*a^*b^*$ spaces.

The following images demonstrate the effects of various types of noise reduction, exaggerated for clarity. While the source image does not contain any very-low-frequency *noise*, it was chosen because it does display the effects (and side-effects) very well.

Comparison of Noise Reduction - Color Spaces



The noisy image. Detail recovery will be intentionally left at o, and chromaticity strength will be intentionally set very high, to emphasize effect.



Luminance noise reduction in L*a*b* space, conservative.



Luminance noise reduction in L*a*b* space, aggressive.



Luminance noise reduction in RGB space, conservative.



Luminance noise reduction in RGB space, aggressive.



Chrominance noise reduction in $L^*a^*b^*$ space, conservative.



Chrominance noise reduction in L*a*b* space, aggressive. Notice how colors bleed one into another in areas where one hue meets the other - navy blue bleeds into cyan, green bleeds into dull-red, dull-red bleeds into crimson-red, etc.



Chrominance noise reduction in RGB space, conservative.



Chrominance noise reduction in RGB space, aggressive. Low-frequency detail is lost.

Color Space

Noise reduction can be performed in the L*a*b* and RGB color spaces.

When working in the L*a*b* space, the L* channel is used for luminance and the a* and b* channels are used for chromaticity.

When working in the RGB space, the Y from the CIE XYZ color space is used for luminance and (X-Y) and (Y-Z) are used for chromaticity.

Mode

There are two general noise reduction modes which control whether only high frequency or also low frequency noise is removed. *Low frequency* noise is noise whose blotches cover a large area; conversely, *high frequency* noise has smaller blotches which cover fewer pixels.

- 1. Conservative removes all except very low-frequency noise, so color detail is better preserved at the expense of not removing very large blotches. Use in most cases.
- 2. Aggressive removes also very low-frequency noise at the expense of being more aggressive with higher frequency noise. Use only on extremely noisy photos.

Gamma

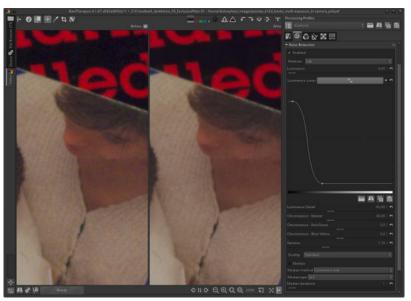
Gamma varies noise reduction strength across the range of tones. Smaller gamma values let noise reduction affect all tones emphasizing the action on shadows, while higher gamma values limit the effect to brighter tones only.

Luminance

"Luminance control" lets you choose whether you want to manipulate the luminance noise reduction via sliders or a curve.

Adjusting the "Luminance" slider is equivalent to manipulating the amplitude of the luminance curve - both affect how strong the noise reduction effect is. The curve has the additional advantage of letting you control noise reduction strength as a function of the pixels' luminance - e.g. it allows you to have strong luminance noise reduction in the shadows and none in the highlights.

The "Detail recovery" slider allows you to recover structure while not reintroducing noise, unless you set this value too high.



The Luminance Curve in RawTherapee-4.1 was used to completely remove noise in dark areas while fully retaining structure of lighter areas.

Chrominance

Method

Chrominance noise reduction can be performed using one of three methods:

- Manual
- Automatic global
- Preview

Preview noise

The "Preview noise" indicator gives the estimated chromatic noise values after "Chrominance" processing:

- Mean: estimates the average noise value across all channels.
- High: estimate the highest noise value across all channels.

Manual

The three sliders and the curve act on the full image. You control the settings manually.

Master

Controls the strength of chrominance noise reduction. Functions as an offset independently to the redgreen and blue-yellow values. For example if master=50, red-green=-50 and blue-yellow=-50, the end result is 0; no effect.

Red-Green

Reduce/boost noise reduction in the red-green channel (a* in L*a*b*).

Blue-Yellow

Reduce/boost noise reduction in the blue-yellow channel (b* in L*a*b*).

Chrominance curve

The chrominance curve lets you control chrominance noise as a function of the pixels' chrominance - e.g. it allows you to have strong chrominance noise reduction in areas of low saturation and weak noise reduction in areas of high saturation. This curve modulates the "Master", "Red-Green" and "Blue-Yellow" sliders' action by multiplying their values by the curve ordinate. For example, if the master slider is set to 30 and the curve is at mid-height, the equivalent result will be about 45. It can be useful to boost noise reduction in grey or drab areas, as we distinguish noise more easily in areas of low saturation than we do in areas of high saturation. When using the "automatic global" noise reduction method, the automatically-calculated parameters are an average for the whole image, and they might be insufficient to remove noise in these low-saturation areas - the chrominance curve can help.

Automatic Global

The algorithm splits the image into multiple cells. For each cell the following are calculated:

An average noise level for the red-green channel and the blue-yellow channel.

• A maximum noise level for the same channels.

Preview

This method is only operational when zoomed to 100% or more. It analyzes the areas currently visible in the preview (if you are zoomed to 100% or more) and calculates:

- An average noise level for the red-green channel and the blue-yellow channel.
- A maximum noise level for the same channels.

The three sliders - Master, Red-Green and Blue-Yellow - as well as the "Preview noise" values, are updated accordingly.

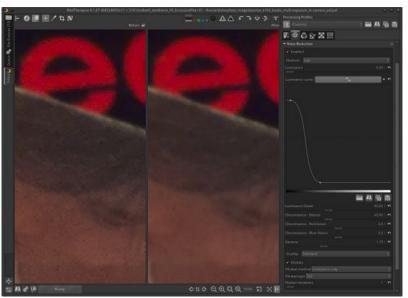
If you would like to keep the currently calculated values then you should switch back to "manual" method, else the values will be re-calculated when you pan or when you copy the profile to other images.

Median

Use this filter to remove tiny, sharp-looking artifacts left-over from noise reduction.

Median Type

The median filter (https://en.wikipedia.or g/wiki/Median_filter) replaces each pixel with the median value of its neighboring pixels. The contiguous group of pixels being sampled is called the "windows". This window slides pixel by pixel over the entire image. You can choose the size of this window using the "Median type" drop-down. The larger the size, the longer it takes.



The Median filter in RawTherapee-4.1 was used to eliminate artifacts which fit in a 3x3 window left over from the noise reduction process.

Available window sizes:

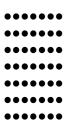
- 3x3 soft: treats 5 pixels in a 3x3 pixel window.
 - 0 0
 - •••
 - 0 0
- 3x3: treats 9 pixels in a 3x3 pixel window.
 - •••
- 5x5 soft: treats 13 pixels in a 5x5 pixel window.
 - 00•00



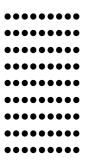
■ 5x5: treats 25 pixels in a 5x5 pixel window.

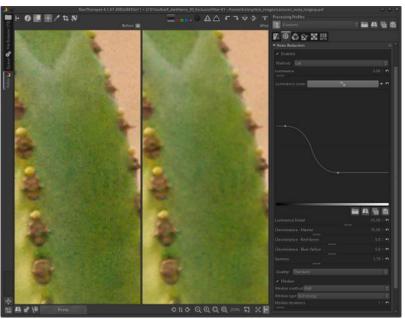


 7x7: treats 49 pixels in a 7x7 pixel window.



9x9: treats 81 pixels in a 9x9 pixel window.





Comparison of noise reduction without and with median filtering, zoomed to 200%. While the filtered image on the right may appear less sharp, notice that the unfiltered image on the left does not really contain any more detail - the noise gives the illusion of sharpness. At the same time, the filtered image when saved to JPEG/92/standard was almost 40% smaller than the unfiltered one (350kB vs 215kB). RawTherapee-4.1.

Sometimes it is possible to achieve higher quality running several iterations with a small window size than one iteration with a large window size.

Median Method

You have five methods at your disposal:

- Luminance only: works in the L*a*b* space, but only affects the L* channel.
- Chroma only: works in the L*a*b* space, but only affects the a* and b* channels.
- Weighted L* (little) + a*b* (normal): affects all channels in the L*a*b* space, but acts more weakly on the L* channel.
- L*a*b*: affects all channels equally.
- RGB: works in the RGB space, and the window size choice is limited to 3x3 soft, 3x3 and 5x5.

When using the "Luminance only" and "L*a*b*" methods, median filtering will be performed just after the wavelet step in the noise reduction pipeline. When using the RGB color space, it will be performed at the very end of the noise reduction pipeline.

You may wonder what other uses median filtering has apart from the elimination of pixels which strongly differ from their surrounding neighbors for aesthetic reasons. One of these benefits is a reduction in file size when saving to compressed formats such as JPEG and PNG. Median filtering removes variations which you will lose anyway if you downscale the image. You are also likely not to see these variations if you print the image. Removing them using median filtering can reduce the file size by even 40% (tested using JPEG compression strength 92 with "balanced quality" chroma subsampling (https://en.wikipedia.org/wiki/Chroma subsampling), so give it a try if output file size is a factor.

Finally, the "chroma only" median filter method can be used as a complement to automatic chrominance noise reduction calculation - by reducing sharp outliers it can soften the calculated values, thereby avoiding fading out color detail too much.

Defringe

The effects of this tool are only visible at a preview scale of 1:1 or more. Use a detail window (click on the icon under the main preview panel) to inspect a part of the image, or zoom the main preview to 100% (also called 1:1) .

Purple fringes are a form of axial (or longitudinal) chromatic aberration, and appear along dark edges adjacent to bright areas due to incorrect focus, lens imperfections, or simply (but more technically) due to the nature of lenses not focusing all colors on the same plane. As lenses are optimized to focus visible light of longer wavelengths on the same plane, those shorter wavelengths farther away from the ones the lens was optimized for (i.e. purple, violet - wavelengths on the shorter side of the visible spectrum) can visibly tint dark regions when the bright areas are of sufficient intensity. This tool should be able to effectively remove most of them.

Defringing is applied either in the Lab or the CIECAMo2 space (if CIECAMo2 is enabled). Consequently, enabling CIECAMo2 may lead to slightly different defringing results, especially when using the *Hue* curve.

Interface Description Radius

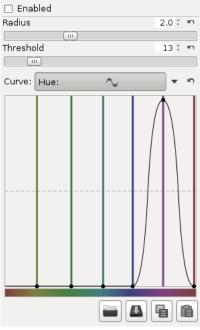
Strong chromatic edge fringes are suppressed by averaging over a neighborhood of the specified radius.

Threshold

Sets a threshold for the application of defringing.



Defringe removed purple fringing.



Defringe Hue curve to target only purple fringing.

Hue

You can use the *Hue* flat curve to specify which colors *Defringe* should target. The horizontal axis represents the range of colors, and the vertical axis the strength of fringe removal. This allows you to limit the action to a specific range of colors without affecting colors of other hues.

If you place a purple dot at the top and keep the rest of the colors at the bottom, purple fringes will be removed with a maximal strength while other colors will not be affected.

Contrast by Detail Levels

The effects of this tool are only visible at a preview scale of 1:1 or more. Use a detail window (click on the icon under the main preview panel) to inspect a part of the image, or zoom the main preview to 100% (also called 1:1) .



Unsharp Mask. RL
Deconvolution.

Edges. Microcontrast.

Edges and Microcontrast.

Contrast by Detail Levels.

Contrast by Detail Levels uses wavelet decomposition to decompose the image into six levels, each adjusted by a slider. Slider o (Finest) has a pixel radius of 1, sliders 1 to 5 have a pixel radius of approximately 2, 4, 8, 16 and 32 pixels. Giving a slider a value less than 1.0 decreases local contrast at that level, while giving it a higher value increases it. Thus you can use it to increase perceived sharpness of an image, to increase local contrast, or to mitigate certain levels of detail.

You should remember that resizing an image has a direct impact on perceived sharpness, as does viewing distance. In practical terms this means that you should use this tool while zoomed more or less to a level representative of your intended final image size and viewing distance, so if you intend to print the high resolution image on a 90x60cm canvas and admire it from 30cm away then it makes sense to zoom in to 100% and tweak the "o (Finest)" slider. However, in real life such large prints usually hang on the wall and are appreciated from the couch a few meters away - from there the finest detail level setting will have no effect whatsoever - your eyes cannot make out the detail from that distance. The same goes for images you intend to resize (downscale) for use on the internet or email to friends or clients - not only do you lower the resolution by downscaling, but they will also be viewed on low resolution devices, probably not even fullscreen, e.g. on a laptop, tablet or phone. In this case too playing with the "o (Finest)" detail level will make no difference to the end result. Most of the time only sliders "3" and "4" will have a practically useful effect.

For example, in order to remove skin blemishes while retaining skin texture on this 10 megapixel photo, where your intention is to view it at its full size and from up close (e.g. art gallery), zoom to 100% and start by setting the sliders as follows, then fine-tune to taste:

```
0 (Finest) : 1.4

1 : 1.4

2 : 0.4

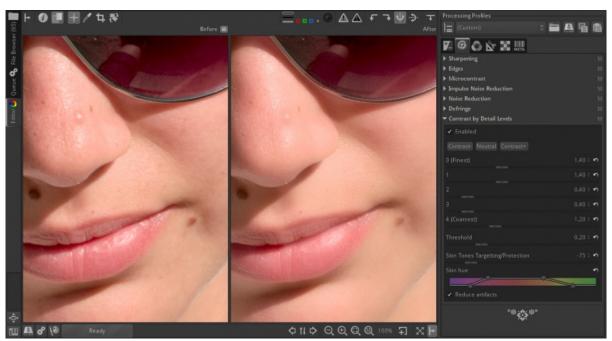
3 : 0.4

4 (Coarsest): 1.2

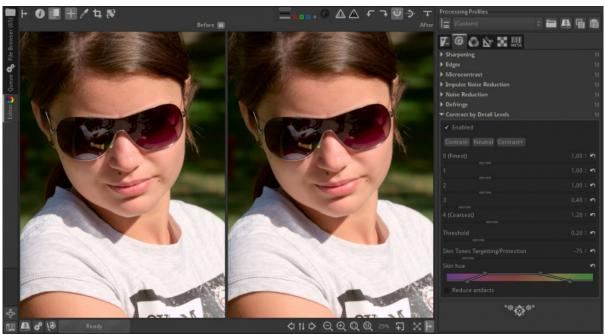
Skin Tones Targetting/Protection: -75
```

If you want to downscale the photo to use it on the web, you should zoom out to about 25% which is more or less the size at which the image will be viewed. Using the settings above we can still see that the skin is smoothed, but if you reset the first three sliders to "1.0", you will see doing so made no difference! The reason is that at this smaller resolution, changes to those levels are lost in the downscaling process. The

result



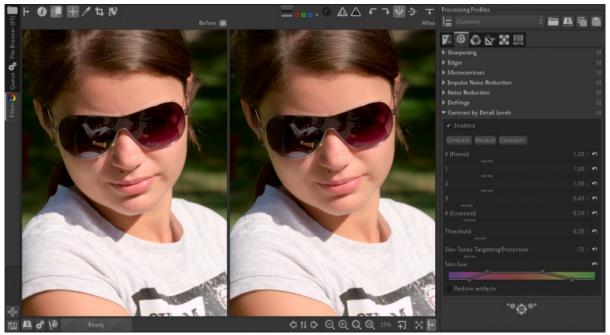
At 100% zoom, Contrast by Detail Levels removed skin blemishes and enhanced skin texture.



At 25% zoom, the effects of the first three sliders are lost by the downscaling process.

would be identical if you saved the full-sized image and then downscaled it to 25% afterwards in some other program, so don't think this is RawTherapee's shortcoming - this is simply how sharpness (resolution and acutance (https://en.wikipedia.org/wiki/Acutance)) work. Knowing this, if you only intend to share the downscaled version of your photo, you can save yourself time by simply ignoring the first three sliders.

Furthermore, you may find when you zoom out that the sliders had effects which were not immediately apparent at 100% zoom. For example you may find that using the "4 (Coarsest)" slider can pleasantly



At 25% zoom we saw what was not clear at 100% zoom, that the harsh shadows can further be softened using the "4 (Coarsest)" slider.

mitigate large, harsh shadows, so set it to 0.5 and fine-tune to taste.

Process Locate Before/After Black-and-White

This combo-box lets you decide when in the pipeline the CbDL tool will run. This tool had been added to RawTherapee long ago, and more recently the Black-and-White tool was added, placed before CbDL in the pipeline. An unforeseen result was that if you enabled the B&W tool then you could not use CbDL's Skin Targetting/Protection because a black-and-white image has no skin color information. This combo-box was added to remedy the situation. Running the CbDL tool before the B&W tool lets you target skin tones before conversion to black-and-white. We recommend you stay with the default option, "Before Black-and-White".

Contrast +/- and Neutral

Use the "*Contrast*-" button to move all five sliders by preset amounts to the left (noise reduction). Use the "*Contrast*+" button to move all five sliders by preset amounts to the right (sharpening). Use the "*Neutral*" button to reset all sliders to o.

Feel free to move individual sliders as well, and inspect the results in the detail window; you may want to zoom in to 200% or more to see better what this filter does.

For high ISO shots (1600+), try for example clicking on the "*Contrast*-" button twice and using Sharpening > Unsharp Mask with an amount of 80.

Threshold

The "*Threshold*" parameter is used to prevent the sharpening of noise: if a pixel's luminance differs only a bit from its neighbors (the difference is less than the threshold), then it is not sharpened. You can set the threshold also to o but then everything will be sharpened (even the noise).

Color Tab

White Balance

Introduction

Digital images generally consist of a mixture of the three primary colors: red, green and blue. For various reasons which you can read about in-depth elsewhere, the red, green and blue values which serve as the starting point in any raw photo development program need to be corrected in various ways before they resemble the photographed scene. One of these corrections is performed by adjusting the white balance ensuring that neutral-colored (white) objects in the photographed scene still appear neutral on the photograph. Adjusting the white balance affects all colors, though it is easiest to discern whether the white balance is correct if an object you know to be of a neutral (white, gray) color looks non-neutral.

White balancing works by multiplying each of the primary colors by a different amount, until a satisfactory result is reached. In order to make this operation more human-friendly, instead of operating on the three multipliers directly, the user is presented with an abstraction in the form of a temperature slider which adjusts colors along a blue-yellow axis, and a tint slider which adjusts them along the magenta-green axis.

A neutral color is one whose red, green and blue values are equal. For example, R=G=B=65% and R=G=B=90% are both neutral, the former being darker than the latter. You can tell whether the white balance of a spot which should be neutral is correct by checking whether that spot's RGB values match, or whether the a* and b* values in the L*a*b* color space match, or whether the RGB indicator bars under the main histogram are directly over each other. You can do this even if you have a very miscalibrated monitor. Your perception of color changes depending on the color of the surroundings and of the illumination in your room, so don't always trust your eyes - verify using the method described above.

Having an incorrect white balance results in the image having a color tint, typically warmer (orange) or colder (blue). Some people use this for creative effect, however there are various tools and operations which rely on the assumption that the white balance of the image is correct (for example highlight recovery in the Exposure tool, skin targeting in the Contrast by Detail Levels tool, sky targeting in the Wavelets tool, the CIECAMo2 tool), so you should not misuse the white balance tool to create a color cast for artistic effect but rather use it to ensure that neutral areas remain neutral, and then use Color Toning or any of the other tools to render a creative color tint.

The white balance tool can be turned on/off. When off, the multipliers are set to R=1 G=1 B=1 when working with raw files. This can be useful for diagnostic purposes or when working with UniWB images.

Interface Description

Method

White balance can be set in different ways: Camera, Auto, Custom, or a host of presets for different light sources.

Camera

Takes the white balance used by the camera. If you shoot only in raw (so no raw+JPG), put the white balance settings of your camera on Auto. This should generally give good results.

AWB Auto

Automatically corrects the white balance, by assuming that the average color of the scene is neutral gray. Works well for a wide range of scenes, and can be a good starting point for manual adjustments.

Custom

Set your own color temperature and green tint by moving the two sliders and/or using the Spot WB tool.

- Light source presets
 - ** Daylight (Sunny)
 - **a** Cloudy
 - **1** Shade
 - Underwater
 - Tungsten
 - **#** Fluorescent
 - **ﷺ** Lamp
 - ∰ LED
 - **5** Flash

Pick

When you click on the Pick button \mathcal{I} (shortcut: **w**), the cursor changes into a pipette when it's over the preview. Click on a neutral area to set the correct white balance for the whole image based on the clicked area.

Pick a spot which should have a neutral tone - gray or white. This spot should not be clipped in any of the three channels, as clipping means that information from the clipped channel is missing. As far as white balancing is concerned, "white" does not mean R=100%~G=100%~B=100%~as that would be clipped, but instead means a shade of gray - even a very light one, but still one without any clipping. The picked spot should also not be black, as black means that insufficient data was captured for that area, and so a correct white balance calculation cannot be performed.

You can use the picker multiple times on different places in the photo until you find an ideal spot. Use the *Size* drop-down box to change the size of the pipette.

This tool can be used as well inside a detail window. Right-click to cancel the tool and to get the regular cursor back.



This tank wagon was the whitest object in the scene. We can tell that the white balance is wrong because the RGB levels are not equal, and the RGB indicator bars directly under the histogram are spread apart.



Picking the white balance off the side of this tanker, which we know to have been the whitest object in the scene, adjusts the colors across the whole image such that the RGB levels in that spot are now equal.

Temperature and Tint

The temperature slider adjusts colors along the blue-yellow axis. Moving it to the left makes the image cooler (bluish); moving it to the right makes it warmer (yellowish).

The tint slider adjusts colors along the magenta-green axis. Moving it to the left makes the image more magenta; moving it to the right - more green.

Blue/Red Equalizer

The red/blue equalizer allows to deviate from the normal behavior of "white balance", via increase or decrease of the ratio between red and blue. This can be useful when shooting conditions are far from the standard illuminant, e.g. underwater, or are far from conditions where calibrations were performed, for which the color matrices in the input profile are unsuitable.

AWB Temperature Bias

The auto white balance temperature bias slider allows you to specify how much the automatically-calculated temperature should deviate. Use this if you would like the automatically-calculated white balance to be cooler or warmer.

White Balance Connection to Exposure

The white balance is described in temperature and tint, but will for raw images be translated into weights of the red, green and blue channels. The weights will be adjusted so that the channel with the smallest weight reaches clipping in the working space (usually ProPhoto RGB) when the raw channel is clipped. In other words, with exposure set to 0.0 and no highlight recovery enabled the full visible range is fully defined by the raw backing. As white balancing changes the weights you may see a slight exposure change if you make drastic changes.

Vibrance

Vibrance is an intelligent saturation adjustment tuned to correlate with the color sensitivity of human vision. The vibrance effect is applied with higher accuracy on a correctly white balanced image when the RGB *Saturation* and Lab *Chromaticity* sliders are set to 0. You can separately control the vibrance of *pastel* tones (tones of low saturation) and *saturated* tones (as the name implies, tones of high saturation).

Interface Description

Pastel/Saturated Tones

These *Pastel Tones* and *Saturated Tones* sliders let you individually control saturation of saturated and pastel tones.

Pastel/Saturated Tones Threshold

The threshold adjuster lets you differentiate between pastel and saturated tones - the threshold between what is considered a pastel tone and what is considered to be saturated.

Protect Skin Tones

When enabled, colors closely resembling natural skin tones are not affected by the vibrance adjustments.

Avoid Color Shift

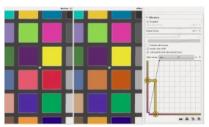
When enabled, exercises extra control to avoid hue shifting.

Link Pastel and Saturated Tones

When enabled, the vibrance level is adjusted with a single slider equally controlling saturation of both pastel and saturated colors.

Skin Tones - Hue According to Hue

This H=f(H) curve lets you change the hue of skin tones. It behaves just like the tone curve you know from any image editing program, but instead of working on a whole color channel it just works on a range of common skin tones. The x axis represents the input hue, and the y axis the output hue. Decide on the specific skin tone hue you want to change (the one your subject has in the photo), find it on the x axis, and then change the curve to map it to a nicer tone on the y axis. In practical terms, you can use this tool to map a British-pink suntan skin color to a nice Maldives-brown one.



Using the *HH* curve, a pinkish skin hue was changed to a tanned-brown one.

Channel Mixer

The *Channel Mixer* is used for special effects, for color and black-and-white alike. The *Channel Mixer* is divided into three sections: Red, Green and Blue. Those sections represent the three available color output channels in a RGB image. All values shown here are percentages. The mixer makes a new R channel from the specified percentages of the existing R, G, and B channels; similarly for the other channels.

The *Channel Mixer* can be used to adjust the white balance of the photo, if the proper *White Balance* tool does not suffice. For the sliders that have a starting value of 100, enter values other than 100 and you are in effect changing the white balance (i.e. only changing the R slider on the R channel, only the G slider on the G channel, and only the B slider on the B channel). Get the image into the right ballpark, then use the *Temperature/Tint* sliders of the usual *White Balance* tool to fine tune. In principle this allows an arbitrary range for *White Balance*.

Black-and-White

The



Before/After view of a color image processed with the Black-and-White tool.

Black-and-White tool gives you great control over the process of converting a color image into a black-and-white one.

For black-and-white color toning, see the Black and White Color Toning section described in the Lab Adjustments chapter, Exposure tab.

Black-and-White addon

The different methods

General remarks

The Black-and-White tool is organized in three methods, each producing a different black and white result.

- 1. Desaturation;
- 2. Luminance Equalizer
- 3. Channel Mixer

Please note that Rawtherapee can produce black-and-white images without the use of this tool:

- 1. by setting the Saturation slider in the Exposure tool of the Exposure tab to -100;
- 2. by setting the Chromaticity slider in the Lab Adjustments tab to -100;
- 3. by enabling Film Simulation in black and white (films Ilford, Kodak, Fuji...)

Nevertheless only the methods in the current tool gives you the maximum of possibilities for a black-and-white conversion.

For a perfect gray tone, except in the case of Color Toning, as "Ajustements Lab" are treated at the end of the pipeline, the values of "a" and "b" in the Lab Adjustments of the "Exposure" tab are set to zero.

Note also the interaction with the Color Toning tool, see Color Toning section below.

Desaturation

This method works in such a way that for each pixel (R=G=B) is given an equivalent luminance value of L=0.299*r + 0.587*g + 0.114*b. This ensures a totally neutral gray image.

Note: the other 2 methods of desaturation in Rawtherapee (quoted above) give other results due to different algorithms. In "Exposure" it is channel "S" from HSV that is set to o. in "Lab Adjustments" it is chromacity C=sqrt(a*a+b*b) that is set to o.

Luminance Equalizer

This method use a flat curve, that allows to modify the luminance based on hue.

The algorithm uses a conversion rgb==>LCH [modifying L based on H]==>rgb with gamut control.

Unlike some other commercial software that act only on a limited number of colors with sliders, you can interact on the whole color palette with Rawtherapee.

Finally, the R, G, B values are set at the same level to ensure a perfect gray tone.

There is a gamut control, but it doesn't prevent you to obtain very good special effects by pushing the curve to extreme values.

Here, the pipette is very useful. For example, choose with the pipette a tone area you want to darken. This adds a control point on the curve. Move this control point downwards to darken (or upwards to lighten) this tone.

Channel Mixer

At first sight this method seems to be very complex!

There is however a simple explanation: this method uses a channel mixer in order to carefully manage the balance between the different color components of the image, to reconcile the distribution of the lights, mid-tones and shadows. It amounts to take a percentage of each channel (R,G,B) and put them together!

The sensible reader with a mathematical mind will notice that the sum of the 3 channels should be 100% to avoid clipped highlights. The same reader will look only for positive values (because they are logical) and no negative values.

But don't let this stop your creativity, open your mind to:

- a) values over 100%;
- b) negative values.

With these two possibilities that you have to experiment with, you can create special effects and color filters such as infrared, but also some common settings such as landscapes, portrait, contrast, etc.

Components of the channel mixer

Presets

It allows to choose between:

- predefined settings (Normal contrast, High contrast, Luminance, Portrait, Landscape, High and Low Sensitivity, Panchromatic, Hyper Panchromatic, Orthochromatic, Infrared);
- settings at user disposal based on four criteria:
- 1. **Absolute RGB**: it offers to the user to mix the three channels R, G and B without any control for the limits. You can enter positive or negative values with a sum that is lower, equal to or higher than 100%.
- 2. **Relative RGB**: it offers to the user to mix the three channels R, G and B, but with control for the limits. You can enter positive or negative values, but the sum will always been forced to 100%. E.g. if you set R=10%, G=10%, B=30%, this is translated to R=20%, G=20%, B=60%. This mode is the default one for all the predefined settings as "Landscape": R=66% G=24% B=10%. Relative RGB is the most intuitive and simple setting, especially if one doesn't set negative values.
- 3. **Absolute ROYGCBPM**: (for Red Orange Yellow Green Cyan Blue Purple Magenta). This "special" mixer offers 2 interesting options:
 - Tweak the complementary colors: in this case, if one acts on a "OYCPM" slider, a correction is automatically made on the basic colors (R,G,B);
 - algorithm OYCPM: if set to "Linear", it has a strict proportional response to the wanted strength and acts directly and proportionally on the 2 basic colors, e.g. Orange acts on Red and Green; If set to "Special effects", it introduces a funny effect in the conversion of Orange to red and Green (non-linear response, and possibly also on Blue depending on the values of the sliders).

This is the least intuitive setting but with maximal creativity possibilities

4. Relative ROYGCBPM: as above, but with a limit control to 100% for the 3 basic channels R,G,B

Color Filter

The color filter simulates shootings with a colored filter placed in front of the lens. These filters reduce the transmission of a specific color and have therefore an impact on the luminance. E.g. a red filter darkens a blue sky and lightens the reds. This filter works as a multiplicator for the settings made with "presets".

Auto

This button activates an algorithm that calculates on the entire image and strictly balance the 3 basic channels R, G, B to give them the same relative weights.

Warnings

- You will notice the incidence of the final tuning of the channel mixer after you have made all tunings in Presets (including "Adjust complementary color" and ""Algorithm OYCPM"). The line under Presets displays 4 numbers, e.g.: R=37,2% G=-82,3% B=126,6% Total=155%. In this case, the global image lightness will exceeds the original one of 55%, and each pixel will have its own values multiplicated before mixing by the 3 previous data.
- For the positive values in relative mode, the result is predictable... It's the usual mode for this channel mixer. You can find on the web values for black-and-white film simulating, e.g. "Ilford Delta 100: 21,42,37", etc.
- On the contrary, in the absolute mode, negative values, the use of the sliders "OYCPM" and the algorithm "Special effects" can lead to unexpected results: black screen, artefacts, ...

Gamma Correction

One can change the rendering of the tones for each channel (R, G, B) using the gamma sliders. This command roughly simulates the rendering of the paper under the enlarger (hard, normal, soft). Pushing the sliders to the left (negative values) darkens the image and gives more contrast, pushing to the right (positive values) softens the image. Notice: there isn't a bug in the Blue channel: when using Prophoto as working space, the blue channel is less active. Use the sRGB working space to see the effect.

'Before' curve and 'After' curve

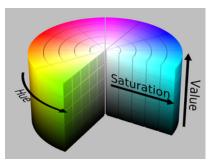
These curves are used the same way and have the same final result as those described in the Tone Curves section of the Exposure tool. They allow customizing the Black-and-White tool, making it independent of the tunings made elsewhere. Notice: the "After curve" has only one mode, as the image is then in black-and-white!

Color Toning

- You can use Color Toning with the Black-and-White tool for special effects. You can also use Color Toning with black-and-white film simulations, but provided the black-and-white tool is enabled.
- The architecture (the various tools order in the processing pipeline), the algorithms "Color toning" and "Black-and-White" have been adapted to give you the maximum of the joined effects.
- You can act simultaneously on all the possibilities in "Color toning", nevertheless, Color Balance Shadows/Midtones/Highlights gives the most possibilities.
- Try switching between Color Balance Shadows/Midtones/Highlights and e.g. L*a*b* blending, try the gamma sliders and the curves in the "Black-and-White" tool.
- Of course you have to walk through a number of trials and errors iterations if you are looking for special effects.

HSV Equalizer

The *HSV Equalizer* (Hue, Saturation and Value) works in the HSV color space (https://en.wikipedia.org/wiki/HSL_and_HSV). Use it to selectively change colors, e.g. make the sky more dramatic by making it somewhat darker or more saturated without touching other colors. It can be very useful for correcting skin tones as well. 'Incorrect' use may produce effects resembling the analogue cross-processing technique. When used for black-and-white, be sure that the raw is desaturated by setting the Chromaticity slider in the Lab Adjustments tool to -100, or by enabling the Black-and-White tool. The *HSV Equalizer* does not work when the raw file is desaturated by setting Saturation in the Exposure tool to -100.



The HSV color space.

There is a description of how the flat curve type works in the General Comments About Some Toolbox Widgets chapter.

Film Simulation

The Film



The Film Simulation tool can change the look of a photo to match a film stock with a single click.

Simulation tool allows you to match the colors of a photo to a reference image with a single click. This tool requires the use of reference images in the HaldCLUT pattern, in either PNG or TIFF format. Each HaldCLUT image corresponds to one "look". Though the look can be based on anything, most of the reference images we ship are based on classic film stock, ergo the tool's name.

RawTherapee needs to be able to access these HaldCLUT images in order that you can use the tool. You can download our collection, called the RawTherapee Film Simulation Collection, or you can design your own - more on that later on. The first time you run this tool you will see a message informing you that you need to point RawTherapee to a folder which contains the reference images. Once you have them ready in a folder, go to "Preferences > Image Processing > Directories" and set "HaldCLUT directory" to the appropriate folder.

Startup time

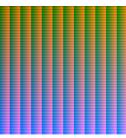
RawTherapee will scan the designated HaldCLUT folder every time you start the program, that's why it is important that you create a folder which you will use only for storing these HaldCLUT images and nothing else, so that RawTherapee does not waste time scanning unrelated files. As a safety measure, you will be warned if scanning the folder on startup takes more than 10 seconds. Should that happen, just click the

button in the popup to stop the scan, then go to "Preferences > Image Processing > Directories" to see which folder is being used, and either point RawTherapee to a folder which contains only HaldCLUT images and nothing else, or to an empty folder if you don't want to use the Film Simulation tool.

To give you an idea of how the startup time is affected, the difference between having 0 images in the HaldCLUT folder and having 500 images in it (that's more than in our collection) results in a 100ms startup time difference - that's nothing. If, however, you were to accidentally tell RawTherapee that the HaldCLUT folder is $C: Program\ Files\ (x86)$, then the startup time could even take several minutes, as that folder contains hundreds of thousands of files. As you can see, there is no reason to worry when using HaldCLUTs as long as you use a dedicated folder as suggested, keeping only HaldCLUT images in it.

How It Works

The Film Simulation tool uses specially prepared images in what is called a HaldCLUT pattern. "CLUT" means "Color Look-Up Table", while what "Hald" means is anyone's guess. A HaldCLUT image contains a set of gradients of various hues arranged in a matrix. The neutral, unaltered state is called an "identity HaldCLUT image". Applying an identity HaldCLUT to a photo will result in no change at all. To create a "look", the identity image is opened in an image editing program - any image editing program - and the colors are modified in some global way, for example using levels, curves, adjusting hue and saturation, etc. Only global adjustments such as the ones listed should be applied - local adjustments are incompatible with how this works; for example you cannot make a HaldCLUT image give a tone-mapped look, it cannot denoise, it cannot sharpen, but it can make skin tones look more tan, and foliage more vibrant.



The level 12 Hald_CLUT_Identity.tif image.

Though the HaldCLUT image looks like a set of gradients, what it really is, is a graphic representation of a matrix of numbers. These numbers are the "out" or "result" values. The "in" or "source" values are known to all programs which support HaldCLUT images. These programs known that a pixel at position x=143 y=0, for example, should be pure red, RGB=(100%, 0%, 0%). If you edit the identity image and change that pixel to be less red, say more orange, then save this modified identity image as a new HaldCLUT file, then whenever you use that file RawTherapee will make the same change to all pure red pixels. This means that if you want to make a certain shade of red more orange you don't need to edit boring numbers to accomplish it - you just open the HaldCLUT identity image in an image editor such as RawTherapee or GIMP, adjust that shade of red (or, more realistically, perform some global adjustment which affects all colors or all shades of red, such as though manipulating RGB curves), save the HaldCLUT image, and then use it in RawTherapee on any photo to repeat that same adjustment.

It should not make a difference what color profile is assigned to a HaldCLUT image. What matters are pixel values, because, remember, a HaldCLUT image is just a matrix of "out" numbers. "Assigning" a color profile does not alter stored pixel values, but "applying" or "converting" does alter stored pixel values, so don't do that. Having said that, the assigned color profile could have an influence, because an image editing program might change the working space depending on the assigned color profile.

Color transformation precision depends on the number of levels (columns) in the HaldCLUT image, and on the bit depth used to store the HaldCLUT image. If working on a 32-bit image (or an any-bit image in 32-bit space) using a 12-level HaldCLUT stored in an 8-bit file, for example, the HaldCLUT stores color values with less precision than your image (or your working space) can store. This should not be a problem though, as unknown colors are interpolated from known values, so posterization should not

occur. The bit depth of the image format you use to store the HaldCLUT depends on your needs, and is not strictly related to the number of levels the HaldCLUT contains. Photographers will be fine using 8-level HaldCLUTs stored in 8-bit images even when processing higher bit-depth photographs.

For more information, refer to Eskil Steenberg's page on HaldCLUT: http://www.quelsolaar.com/technology/clut.html (http://www.quelsolaar.com/technology/clut.html)

To generate an identity 12-level 16-bit HaldCLUT image using ImageMagick, run this command in a console:

convert hald:12 -depth 16 -colorspace sRGB hald12_16bit.tif

Caveat

If you need to generate your own identity HaldCLUT, do not use the program for generating HaldCLUT images from www.quelsolaar.com as it has a bug which causes issues with highlights. Use ImageMagick or GraphicsMagick instead. You can of course use the identity file we provide here - it is bug-free.

Make Your Own

This section explains how to put the HaldCLUT identity file to use so that it reproduces a specific rendering of color and lightness.

- 1. Open a photo in RawTherapee or some other image editing program and tweak it to your liking. Remember that the Film Simulation tool can only reproduce global tonal changes, so make no local changes no local contrast, no tone-mapping, etc.; make no changes which move pixels no distortion correction; use no sharpening or noise reduction; make only global tonal adjustments color and saturation changes, curves, levels, L*a*b* adjustments. Save the sidecar file or write down the changes you made so that you can reproduce the changes in the next step.
- 2. Open the identity HaldCLUT image in the same program and apply the same sidecar file or re-do the same tweaks as you did in the step above.
- 3. Save this image as an 8-bit sRGB TIFF or PNG in the HaldCLUT folder you pointed RawTherapee to. It's now ready for use. Restart RawTherapee so that your new HaldCLUT appears in the list.

Even if your HaldCLUT image contains colors in only 8-bit precision, missing values will be interpolated so that posterization will not occur in your photo. As such, since there is no visual loss of quality, we recommend using the level 12 identity file, or even an 8-level one, and storing your self-made HaldCLUT images in an 8-bit per channel PNG or TIFF file.

The color profile assigned to the saved HaldCLUT image does not matter. What matters are the pixel values of that image. You may be familiar with the terms "assign color profile" vs "convert to color profile" from GIMP or Photoshop - which color profile is assigned does not matter because it does not alter pixel values, but converting matters because then the pixel values change. When editing the HaldCLUT image in RawTherapee, the choice of output color profile matters because it alters pixel values of the saved image. As the identity image provided by us or generated according to our recipe uses the sRGB primary chromaticities, so you should use RTv2_sRGB or RTv4_sRGB when saving it in order to preserve the colors.

If you apply this HaldCLUT to a photo in RawTherapee and the photo unexpectedly and unintentionally becomes considerably darker or lighter than it should have, then it's likely that the program which you ran it through did something with the gamma. To remedy, you have to undo what that program did. Try generating your own 12-level HaldCLUT but instead of using the "sRGB" colorspace use just "RGB".

Advanced - Identity DNG

This is experimental, it may or may not work.

Some programs might not let you open a TIFF image. If the program supports DNG files, and demosaiced ones at that (what Adobe DNG Converter refers to as "Linear (Demosaiced)"), then you can use this trick. Using ImageMagick, ExifTool and the commands below, making use of the fact that DNG is just a form of TIFF, you can generate an identity HaldCLUT in DNG format:

```
convert hald:12 -depth 16 -colorspace RGB -gravity NorthWest -splice 4x4 -gravity SouthEast -splice 4x4 foo.tif
exiftool -DMGVersion=1.4.0.0 -PhotometricInterpretation='Linear Raw' foo.tif
env -v foo.tif Hald_CLUT_Identity_12.dng
```

Raw editing programs will discard a certain number of pixel rows and columns from the image edges for technical reasons to do with demosaicing. How many rows and columns get discarded depends entirely on the program. You need to figure this out. A 12-level identity HaldCLUT will have precisely 1728x1728 pixels. When you process that CLUT in a program whose color effects you're trying to emulate, the saved image must have precisely 1728x1728 pixels. Since you're fooling the program into thinking it's working on a raw file, and since it will probably discard some pixels around the edges, you need to figure out exactly how many rows and columns of padding are needed and add them around the image.

RawTherapee cuts off 4 pixels all around when reading demosaiced DNG files, so the command above first adds a 4 pixel row and column to the bottom and right edges, then another 4 pixel row and column to the top and left edges. When you open this image in the target program, zoom into each side's edge and figure out whether you need to add more (or remove some), then modify the command accordingly.

Once you have the borders figured out, merely open this DNG in the target program and follow the steps above in the "Make your own" section.

Browsing in RawTherapee

There are many HaldCLUTs in our collection, and you will probably want to try them all out. You don't need to click on each one with the mouse! There is an easy way of checking them all out, one by one, which requires no mouse work. Assuming you had already set RawTherapee up to point to the HaldCLUT folder and you have a photo open and the Film Simulation tool is ready for use, just select any HaldCLUT image using the combobox, and use the "up" and "down" keys on your keyboard to apply the previous and next one. Using this trick, you can even jump across sub-folders - for example from Black-and-White to Color if using our collection.

If you would like to compare the effects of several specific HaldCLUTs, and they are not consecutive in the combobox, an easy way to switch between them is by taking a snapshot after applying each, and then just clicking on the snapshots.

RawTherapee Film Simulation Collection

This archive contains a collection of film simulation HaldCLUTs which you can apply to your photos to instantly match their colors to the film stocks the HaldCLUTs are based on. Unless otherwise noted in the filename, they are all in the sRGB color space, 8-bit per channel, in the PNG image format. Most of them are designed to mimic the results of various film stocks, pushed and pulled in various ways or faded over time.

The suffixes +, ++, +++, -, --, --- refer to the strength the film was pushed or pulled (https://en.wikipedia.org/w iki/Push_processing) during development (non-linear), and "generic" refers to the film type usually sold for rebranding.

Download (http://rawtherapee.com/shared/HaldCLUT.zip) (402MB!)

Changelog

2015-09-20

Added the "CreativePack-1" color collection.

Converted all TIFFs to PNG (except for the identity image).

2015-03-25

The identity CLUT had a bug causing cyan colors in the highlights, it has been replaced with a fixed one.

Numbered the files so they are sorted in the correct order when pushed or pulled (--, -, normal, +, ++).

2014-08-25

The first public release.

Re-organized into *Color* and *Black-and-White*, sub-folders sorted by brand.

2014-08-15

Expanded README.txt and added disclaimer.

2014-07-05

The first internal release.

All images re-compressed with maximum lossless compression.

Learn more about HaldCLUTs here:

http://www.quelsolaar.com/technology/clut.html

http://blog.patdavid.net/2013/08/film-emulation-presets-in-gmic-gimp.html

http://blog.patdavid.net/2013/09/film-emulation-presets-in-gmic-gimp.html

Credits:

Pat David - https://discuss.pixls.us/u/patdavid (https://discuss.pixls.us/u/patdavid)
Pavlov Dmitry

Michael Ezra - https://discuss.pixls.us/u/michaelezra (https://discuss.pixls.us/u/michaelezra)

Disclaimer:

The trademarked names which may appear in the filenames of the HaldCLUT images are there for informational purposes only. They serve only to inform the user which film stock the given HaldCLUT image is designed to approximate. As there is no way to convey this information other than by using the trademarked name, we believe this constitutes fair use. Neither the publisher nor the authors are affiliated with or endorsed by the companies that own the trademarks.

RGB Curves

RGB curves allow to apply a curve adjustment individually to each of the RGB channels. This provides a fine control over image color balance when curves applied to each of the RGB channels are different. Using RGB curves one could make warmer highlights or colder shadows, simulate film cross-processing effect, etc.

Luminosity Mode

The purpose of *Luminosity Mode* in the *RGB Curves* tool is to alter image luminosity by changing the contribution of the RGB channels to it, while keeping image color the same.

The effect is somewhat similar to *V* changes in the HSV Equalizer, but is smoother and broader across hues, not as selective. When working on black-and-white images, similar adjustments could be made via the Channel Mixer, but *RGB Curves* allow a finer control.

Color Toning

About Color Toning

The first question that arises is: What is the definition of, or what do you hear by, "Color Toning" or "Split Toning"? Indeed, when consulting the Web, we generally find out something like: "Color Toning consist in coloring a black & white image in a different way according to the brightness, e.g. coloring highlights in yellow and shadows in blue.

Extending the concept, we can put under the same definition:

- 1. the toning of a color image that allows to add a dominant color to the image. It will be possible to tweak this dominant color in the image highlights and/or shadows.
- 2. to extend the toning to the whole luminance spectrum, and not only, in a restrictive way, to the highlights and shadows.

In RawTherapee, two types of algorithms try to meet the principles defined above:

- i. "blending from target colors": in this case, a chromatic value is weighted according to a formula like: "output hue" = "input hue" + ("target hue" "input hue") * balance where balance is a coefficient between 0 and 1. We can easily find on the internet references to this kind of lagorithm.
- ii. "RGB channels adding and reduction": in this case, according to the luminance (highlights, midtones, shadows), each channel is amplified at the same time the two others are reduced. e.g. an action on the red channel for a given luminance range, will increase the "R" channel by X%, and at the same time, "G" and "B" channels will be reduced by X%. Note it is not a "Channels Mixer". I did not find any references to this kind of algorithm, but studying the Photoshop "Color Balance" module behaviour, I think I figured out an algorithm that gives similar results.

From experience, the first algorithm type will give good color toning results for color images but is not easily predictable, and it is not as good for black and white images, even if, of course, it gives satisfactory results. This algorithm is embedded in two different ways - no one is better than the other - that give various results:

- 1. RGB mode: each R, G, and B channel has the algorithm explained in (i) applied.
- 2. Lab mode: each color component "a" (red/green channel) and "b" (blue/yellow channel) have the algorithm explained in (i) applied. This mode allows, according to the suggested choices (menus), a normal predictability or an important creativity.

The second algorithm type can have three usages based on the "Strength" slider:

- 1. using low values, the user can simulate a "color balance" and accurately tweak the tone color
- 2. using high values, the user will be able to, in "color" mode, get similar results as the "blend" algorithm, but with less creativity
- 3. using high values, the user will be able to, in black and white mode, get strongly specials effects

The different methods

"blend" Methods

The "blend" methods are divided into "L*a*b* blending" that uses both lab chromatic components "a" and "b", and "RGB-sliders / RGB-curves" that use the same RGB algorithm but differ in the user interface. The Lab method isolate the color component from the luminance, whereas the two RGB methods indirectly act on the luminance. This difference partly explain the behaviour gap between these methods.

Even if the interface is different (sliders or curves), the two RGB methods use only one type of opacity (the color blending management), whereas the lab mode offers four of them. The first one "Standard chroma" is similar to the one used in "RGB-curves". The other three allow special effects.

The used curves are special flat curves.

- 1. the color curve displays luminance in abscissa and target hues in ordinate. the two vertical lines delimit the main resulting areas. by moving:
 - the vertical lines positions;
 - the curve shape;
 - the target hues choices;

you will obtain different results

2. The Opacity curve (L*a*b* blending > Standard chroma or RGB curves) displays luminance in abscissa and opacity in ordinate (also called Balance) that translate the way the original hue (image) and the target hue are assemblied, in this case,the opacity value varies from 0 to 1. The highest the curve will be, the more the blending near the target hue. When setting the opacity curve to 0, the image stay unchanged.

The saturation setting (the effects maximum intensity) can be adjusted:

- 1. Manually, in this case the box "Automatic" is not checked. You can move the two sliders "Threshold" and "Strength".
- 2. Automatically, in this case the box "Automatic" is checked. An algorithm takes into account the color space (sRGB, Adobe, Prophoto...) and the image pixels saturation to determine the best values for "Threshold" and "Strength".

This settings are obviously without any effect for images converted into black and white.

"L*a*b* blending" particularities

Special chroma

Here, the flat curve is replaced by a diagonal curve. The two chroma components "a" and "b" (Lab) are modified with the same amplitude. If you move the curve under the diagonal, you introduce negative opacity values, which will bring special effects often unpredictable.

Special a* et b*

Here, the flat curve is replaced by two diagonal curves. The two chroma components "a" and "b" (Lab) are individually modified by two different curves. The first one only acts on the "a" component (Lab), i.e. the red-green dimension. the second one only acts on the "b" component (Lab), i.e. the blue-yellow dimension. If you move the curve(s) under the diagonal, you introduce negative opacity values, which will bring special effects often unpredictable.

Special chroma '2 colors'

Here, like in "Special chroma", the flat curve is replaced by a diagonal curve. The two chroma components "a" and "b" (Lab) are modified with the same amplitude. If you move the curve under the diagonal, you introduce negative opacity values, which will bring special effects often unpredictable.

The difference with "Special chroma" lies in using the color curve. In "Special chroma", the whole curve hue=f(Luminance) is used, in the "Special chroma '2 colors'" case, only the two hues focused by the vertical lines are used.

RGB-sliders particularities

The searched ergonomic aims to be closed to the Lightroom one (like by the way "Saturation 2 colors")

You have two sliders with two levels at your disposal, the first one for highlights, the second one for shadows. For each of both sliders you can set the wished hue and strength: When set to 0, the two strength sliders doesn't change anything to the image.

The "Balance" slider allows to set the equilibrium between high and low lights. By moving it to the left (negative values) the action on highlights is increased, whereas to the right (positive values) the on action shadows is increased.

"Adding" Methods

Color Balance Shadows / Midtones / Highlights

This method is very close to the Photoshop module "Color Balance", both in its operating mode and its rendering.

You can act differently on the highlights, the mid-tones and the shadows.

Each slider acts on a color and its complementary color: Red and Cyan, Green and Magenta, Blue and Yellow

The "Strength" slider allows to set the system sensitivity:

- 1. with low values less than 50 you can use this tool to tweak the image color balance, thus modifying the whole blending to give a generalised chromatic correction,
- 2. with medium values, you can use this tool as a color toning,
- 3. with high values, you can use this tool as a black and white toning, interacting with the Black and White tool (internal algorithm parameters are different for a color or a black and white action)

Select "Preserve Luminance" to prevent any change of the lightness values in the image when modifying the color. This option allows to preserve tone balance in the image.

Saturation 2 colors

This method is close to ACR and Lightroom, both in its operating mode and its rendering.

It is mostly intended to color toning, even if it may be used in interaction with the Black and White tool.

Two sliders with two levels are at your disposal, the first one for highlights and the second one for shadows. For each of the both sliders, you can tweak the desired hue and strength: if set to zero, a strength sliders prevent any change to the image.

The "Balance" slider allows to balance the action between high and low lightness. Moving it to the left (negative values) increase the action on highlights, to the right (positive values) increase the action on shadows.

The "Strength" slider allows to set the system whole sensitivity.

Select "Preserve Luminance" to prevent any change of the lightness values in the image when modifying the color. This option allows to preserve tone balance in the image.

Black and White

It is thanks to going to and fro from the Black and White tool - particularly Luminance Equalizer- to the "Color Toning" tool - particularly Color Balance - that you will get the most pronounced (black and white) special effects.

Film Simulation

- In the case of color film simulation, all the "Color Toning" tools are directly available.
- In the case of black and white film simulation, it is mandatory to enable the "Black and White" tool. The desaturation method is almost neutral and allows a direct use of the black and white simulations in all the "Color Toning" tools", but without being able to use the special effects of the "Black and White" tool.

Color Management

Input Profile

No Profile

No input color profile will be applied. The color matrix will use "1" along the diagonal and "0" everywhere else.

- Raw files will show the camera's native RGB color. They will only be demosaiced and white-balanced.
- Non-raw files will be displayed without any embedded input profile applied, including no gamma correction, which means they will look bright.

This feature is generally only useful for didactic and scientific purposes. For example if the camera has recorded colors far outside of the conventional gamuts, using no input profile ensures that no color clipping occurs.

Camera Standard

Looks for and uses a color matrix from the DNG file, from camconst.json, hard-coded in RawTherapee, or from dcraw (https://en.wikipedia.org/wiki/Dcraw), whichever one it finds first, in that order.

A *color matrix* is a matrix of 3x3 constant values which is multiplied with the camera's native RGB colors to convert them to colors which are as natural as possible. A color matrix works best (i.e. provides more accurate colors) when the white balance is close to what the matrix was calibrated for. The camera standard matrix is calibrated for D65 (https://en.wikipedia.org/wiki/Illuminant_D65), i.e. 6500K. Do not worry if the white balance is quite far off from that though, color will be reasonably accurate anyway.

For applications where the most accurate and fine-tuned color is not of highest importance, such as landscape photography, the color matrix will provide good colors. An advantage of color matrix processing compared to lookup table-based DCP and ICC conversions is that it's purely linear, i.e. a dark and a bright color of the same hue and saturation is translated the same way. This makes it robust and may be the best choice if you will be exporting images for processing in an HDR application or other application when a predictable linear color response is important.

Auto-Matched Camera Profile

Uses RawTherapee's camera-specific DCP input profile that can provide more accurate colors than the standard matrix (and fall back to legacy ICC profiles if no DCP is available). Available for some cameras, these profiles are stored in the /dcpprofiles directory (or legacy /iccprofiles/input) and are automatically retrieved based on matching the exact make and model of the camera as it appears in the info section in the Editor to the filename, e.g., "Canon EOS 5D Mark III.dcp".

In other words, if "Auto-matched camera profile" is selected, RawTherapee will try to do the following, in this order:

- 1. locate a DCP profile in /dcpprofiles
- 2. if DCP is not found, locate an ICC profile in /iccprofiles
- 3. if DCP and ICC are not found, revert to the camera standard color matrix.

If you want to contribute a camera profile, DCP is the preferred format.

Some of RawTherapee's profiles are single-illuminant (Daylight/D50), while others are double-illuminant (Daylight/D50 and Tungsten/StdA). Some include a tone curve, others do not. They strive for accurate colors (i.e. not a specific "look"). Most accurate colors will be achieved for white balances close to the calibration illuminants.

Camera profiles work in the normal range, from black up to clipping. If you enable highlight reconstruction, new data is added above the clipping level and if you bring it into visible space (by negative exposure for example), that range will not be naturally covered by the profile. However, RawTherapee will linearly extend the profile to cover this range too, colors there will get the same correction as the brightest colors of the same hue and saturation in the normal range.

Custom

Specify a custom DNG Profile (DCP) or ICC camera input profile stored on your computer.

DCP is a format specially designed for camera profiles and RawTherapee should support the most recent DNG standard (where DCP is defined), so you can for example use all those provided via Adobe's DNG converter.

ICC profiles on the other hand are more tricky. ICC profiles can be used for a multitude of purposes (printers, displays etc) and since they're not designed specifically for camera profiling, different vendors have chosen different approaches for their ICC profiles. In practice this means that the input image must be pre-processed in some specific way for the profile to work. The profile itself lacks information of how to do this pre-processing, which means that if you are using a third-party profile RawTherapee may not do the expected pre-processing; results will vary.

Third-party DCP support

A DNG Camera Profile, DCP, is the preferred camera profile format for RawTherapee. All elements of the 1.4 DNG specification is supported, with the exception of the black render tag (see below). A DCP can be a pure matrix profile, it can have a LUT (typically 2.5D) to improve the colorimetric accuracy, and then it can have an embedded curve and a separate "look table" on top. It may also add an exposure offset. All those elements can be toggled via checkboxes. However, although it is possible few third-party profiles have been designed to produced the intended color with anything else than all their elements enabled. For example, the tone curve itself changes color appearance so if you disable an embedded tone curve to get a linear profile you can't count on that the color is as intended.

The typical third party profile would come from Adobe Camera Raw / Lightroom, and RawTherapee supports them. Many of Adobe's profile lack tone curve, but in Adobe's world that does not mean that no tone curve should be applied but that Adobe's default curve should be applied. RawTherapee will therefore identify Adobe profiles (from the copyright string) and add the default curve to those (which you can toggle with the tone curve checkbox).

Adobe's DNG converter may add a "baseline exposure" to the DNG file. Adobe's DCP are designed to work with that baseline exposure and then produce a default output which is about the same brightness and contrast as the camera's own JPEGs. RawTherapee can honor this baseline exposure (NOT YET IMPLEMENTED), but this is of course only available when opening a DNG file converted by Adobe's DNG converter. If you instead open a native raw file there will be no baseline exposure and Adobe's DCP may then make a too bright or dark rendering. You can simply adjust with the exposure slider of course.

The DCP format also has a black render tag. This indicates if the raw converter should do "automatic" black subtraction or not. RawTherapee ignores this tag, you can do manual black subtraction with the black slider. As many of Adobe's profiles indicate auto black subtraction and Adobe Camera Raw / Lightroom does it, RawTherapee will in comparison in those cases render a bit lower contrast and brighter shadows.

Third-party ICC support

RawTherapee has specific support for ICC profiles bundled with Capture One and Nikon NX2, so those should work well. Older ICC profiles are not likely to work well though (typically the image becomes extremely dark with unsupported ICC profiles).

Some ICC profiles apply a tone-curve and desaturate bright highlights for a more "film-like" look. Those profiles may not work well together with Highlight Reconstruction. If you see a radical change in contrast when you apply your ICC profile it has applied a tone-curve and then you should not use it together with Highlight Reconstruction.

Unlike DCP profiles, ICC profile processing may cause clipping of extremely saturated colors during conversion. In practice this is rarely if ever a problem, but still DCP should be considered the primary choice if available.

Note on using Capture One ICC profiles: RawTherapee applies the ICC before exposure adjustments, as the intention is that camera profiles should only be used to make the camera more accurate, not really to apply a look (you design the look using the tools instead). Phase One's ICC profiles contain a subjective look though, which means that they typically contain "hue twists", for example saturation in the shadows are increased a bit extra. This means that if you have an underexposed file and push it a few stops those hue twists have been applied on the dark image before exposure adjustment and will thus be in the wrong places after pushing, that is you don't get the same look as in Phase One's Capture One. Therefore it's recommended to have the right exposure out of the camera when using Phase One ICC profiles. You should also apply a suitable RGB "film curve" for example by using the curve tool, as those ICC profiles are designed to be used together with that.

We are aware that LUT ICCs should typically be applied after exposure (just as DCP Looktables are applied), and that would support for example Capture One profiles better. This may be fixed in a future version.

DCP Illuminant

Some of RawTherapee's profiles are single-illuminant (Daylight/D50), while others are double-illuminant (Daylight/D50 and Tungsten/StdA). If a dual-illuminant profile is loaded the "DCP Illuminant" setting will be enabled and you can choose which illuminant to use. The actual DCP standard (part of the DNG standard) does not provide this choice, but instead an interpolation between the two illuminants is calculated based on the chosen white balance (there will only be an interpolation if the white balance is inbetween both illuminants, otherwise the closest is picked). This "interpolated" mode is the default setting of "DCP Illuminant" and for any normal use you do not need to change this.

You can however choose to base the color rendering on one of the specific illuminants. In some cases this might produce more pleasing color. It can also be interesting for diagnostic purposes to see how large (or small) a difference there is in color rendering between the illuminants, but, as said, for general use this setting should be untouched.

Use DCP's Tone Curve

Some DCPs contain a tone curve which may be used to add contrast and brightness to provide a film-like look. This is mainly used for profiles simulating camera maker settings. The tone curve checkbox will be disabled for profiles which do not contain a tone curve.

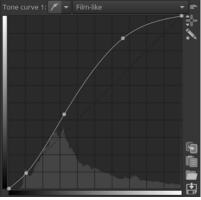
The curve mode used by the DCP tone curve is the same as the Exposure tool's "film-like" mode, meaning you can reproduce the effect using the Exposure tool's tone curves in film-like mode. When contrast is applied with a film-like curve the appearance of the colors will change and overall saturation is increased, except for bright colors which instead are de-saturated. Some profiles which have curves embedded are pre-corrected for this color appearance change and will thus not provide the intended look without the curve applied. Most will however work well without the tone curve applied especially if you add a similar curve yourself using the Exposure tool's curves, but if you want to see exactly how the profile designer intended the colors to look you should enable the tone curve.

While the input color profile is applied at the first stages of the toolchain pipeline, the DCP tone curve is applied later in the pipeline at some point after the Exposure tool.

The curve used in some of RawTherapee's bundled DCP profiles is the same as the default curve used by Adobe Camera Raw. You can get the curve from this .rtc file. In the Exposure tool set either Tone Curve 1 or Tone Curve 1 to "Custom", set the curve mode to "Film-Like", and apply this file: File:DCP tone curve.rtc

Use DCP's base table

This enables the DCP "HueSatMap" lookup table which is used to add non-linear corrections on top of the basic matrix. This is an advanced user setting and unless you want only the pure matrix result should leave it on. It's grayed out if the loaded profile lacks a HueSatMap table.



DCP (ACR) tone curve, DCP tone curve.rtc

Use DCP's look table

This enables the DCP "LookTable" lookup table which is intended to add a subjective look on top generally together with an embedded tone curve. That is if you disable the DCP curve and looktable you may get a neutral "colorimetric" profile, if the DCP was designed that way which is not always the case (if the DCP has both a look table and a base table it's likely that it is, but if it only has a look table it will probably not work well with it disabled). Disabling individual DCP elements are considered advanced user settings, normally you would leave this on.

Use DCP's baseline exposure offset

The DCP may indicate an exposure offset that corresponds to an offset of the exposure slider. The purpose of this is typically to make the brightness of the image match the brightness of the camera's own JPEGs, which can be useful if you're shooting with auto-exposure. Currently this offset is applied "under the surface" so you don't see it on the exposure slider.

Note that if you are using Adobe's proprietary profiles those are expecting that the DNG's "baseline exposure" tag is applied too (the profile's offset is added on top). Currently there is no support for the DNG tag so you need to find that out on your own (using exiftool for example) and then set that offset using the exposure slider if you want to get the exact same brightness as in Adobe Camera Raw.

Save Reference Image for Profiling

Clicking this button saves a linear TIFF image before the input profile is applied. This file can then be used for profiling, i.e. creating a new ICC camera profile. There are various commercial software out there to make ICC profiles, but you can also use the free and open-source Argyll. For DNG profiles there is DCamProf as an open-source alternative.

Cropping, resizing and transform (rotate) will be applied so you can use that to make the output more managable by the receiving software. Argyll is very picky for example and want no more than the test target visible in the image.

You can also choose if you want to export with white balance applied or not. For ICC profiles you should export with white balance applied, but if you intend to make a DNG Profile ColorMatrix (or a DCRAW style color matrix) you should export without.

Working Profile

The default working profile is ProPhoto and should not be changed for normal use.

The working profile specifies the working color space, which is the color space used for internal calculations, for instance for calculating saturation, RGB brightness/contrast and tone curve adjustments, chrominance, etc.

When RawTherapee was based on integer math it was wise to not use working space larger than absolutely needed to get the best precision in the calculations. However, RawTherapee had switched to floating-point processing in 2011, and since version 4.0.12 the default working profile is ProPhoto, which has a very large gamut.

The choice of working profile has an influence on the effect of the curves in all modes except for perceptual - in that mode, changing the working profile will not alter the effect of the curve. If you have trouble fitting colors within the output gamut you can experiment with changing the working profile when using curves in any mode but perceptual.

Note that the working profile will only specify the red, green and blue primaries, gamma will not change as RawTherapee's processing pipeline is floating point with no gamma encoding (that is gamma = 1.0). Some tools (like curves and histograms) will still display with a gamma (usually sRGB gamma) which is hard-coded for the tool and stays the same regardless of working profile.

Adding Custom Working Profiles

RawTherapee 5.5 allows you to specify custom working profiles through a JSON (https://en.wikipedia.org/wiki/JSON) file. The file should be named workingspaces.json and it can reside in:

- the ICC profile folder as set in Preferences > Color Management > Directory containing color profiles,
- or in RawTherapee's own ICC profile folder:
 - Windows: <rt-install-folder>\iccprofiles
 - Linux:
 - When installed using your package manager or self-compiled with BUILD_BUNDLE=OFF: /usr/share/rawtherapee/iccprofiles
 - When self-compiled with BUILD BUNDLE=ON: <rt-install-folder>/iccprofiles
 - macOS: /library/ColorSync/Profiles/Displays

The format of workingspaces. json is as follows:

If "matrix" is present, "file" is ignored. If only "file" is present, the matrix is extracted from the ICC profile. For this, RawTherapee looks only at the R, G, and B matrix columns and the white point set in the profile. Bradford adaptation is used to convert the matrix to D50. Anything else in the profile (LUT, TRC, etc.) is ignored.

It is the user's responsibility to ensure that the profile is suitable to be used as a working space.

Output Profile

Specify the output color profile; the saved image will be transformed into this color space and the profile will be embedded in the metadata. The effects the output profile has on the image cannot be seen in the preview.

RawTherapee lets you specify "input" (e.g. you're camera's profile), "display" and "output" (i.e. printer) device class profiles with an RGB color space, because RawTherapee saves only RGB images. Profiles listed in this combobox are those which come bundled with RawTherapee and those located in the folder set in Preferences > Color Management.

The soft-proofing feature is dedicated to simulating printer rendering. It lets you preview what your image will look like when printed, assuming you use a printer profile which correctly simulates your printer and paper combination. For best printout quality, after you have tweaked your photo using soft-proofing, you should select your printer profile as the output profile and save the image using it. This ensures that the image is encoded using your printer's color space directly from RawTherapee's internal high quality floating-point representation, instead of being saved to an 8-bit image in sRGB for example and then having to be subsequently converted to the printer profile, which would be quite lossy.

The main histogram, navigator and clipping indicators will use either the working or the output profile, depending on your setting in Preference > General.

RawTherapee comes bundled with a number of custom-made high quality output profiles:

```
RT_sRGB
Similar to sRGB
Gamma close to sRGB: g=2.40, slope=12.92
RT_sRGB_gBT709
Similar to sRGB
Gamma BT709: g=2.22, slope=4.5
RT_sRGB_g10
Similar to sRGB
Linear gamma g=1.0, slope=0
RT_Medium_gsRGB
Similar to AdobeRGB1998
Gamma close to sRGB: g=2.40, slope=12.92
RT_Large_gsRGB
```

Similar to ProPhoto Gamma close to sRGB g=2.40, slope=12.92 (close to "Melissa" used by Lightroom)

RT_Large_gBT709

Similar to ProPhoto

Gamma BT709: g=2.22, slop=4.5

RT_Large_g10

Similar to ProPhoto

Linear gamma g=1.0, slope=0

Rec2020

 $\label{thm:conditional} Wide \ gamut, \ larger \ than \ AdobeRGB \ but \ smaller \ than \ ProPhoto$

Gamma BT709: g=2.22, slope=4.5

The recommended output profile when you're saving to an 8-bit format and/or publishing to the web is RT_sRGB. If no profile is selected, none will be embedded, which means that "sRGB" is implied, though it is safer to embed RT_sRGB in terms of getting your image displayed properly in various applications.

RT_sRGB is a **higher quality** version of the standard sRGB profile, which surprisingly is inconsistent between implementations. RT_sRGB was custom-made for RawTherapee by Jacques Desmis and has 4096 LUT points, as opposed to the lower quality 1024 point sRGB profiles. Applications that aren't color managed and won't take advantage of RT_sRGB will fall back on sRGB.

Wide-gamut output profiles such as RT_Large_gsRGB are generally used if you export to a 16-bit or higher bit-depth format for further editing in another application. If you will be sending your image for printing, a wide-gamut output profile is also recommended, since some printers may have wide gamuts (at least in certain colors).

You should have a wide-gamut monitor if you want to work with wide-gamut profiles, otherwise you're flying in the dark.

Color Management addon

Colorimetry concepts

This document aims at several objectives about colorimetry:

- 1. To summarize the raw file processing pipeline from a colorimetry point of view and to highlight key points and gaps
- 2. To explain the main steps principles in order to allow the user to be able to weight what is at stake.
- 3. To explain the use of some functions that may look obscure to newcomers.

Note: This document only deals with RawTherapee 4 and computing with real numbers, and not RawTherapee 3. More over it does not take into account possible malfunctioning (bugs) still unsolved!

Warning

- This document doesn't aim to deal with the whole colorimetry aspects that aren't specific to RawTherapee, like for instance:
 - the printing,
 - screens calibration.
- However, it is recommended to calibrate the screens with one of the many products in the market: colorimetric probe plud software. The created profile is for the monitor only and must never, in any case, be used either as input profile or output profile.
- Under Windows, MacOS or Linux, a software like DispalGUI from Argyll, combined with a probe of quality, even old (e.g. the DTP94 probe I own, and for which it doesn't exist any more [??] drivers for Windows) give very good results; the processing time is quite long (about an hour).
- RawTherapee automatically detects the system profile, however you can enter the screen icc filename in « Preference/ Color Management / Monitor Profile »
- The reader shall preferably have minimal knowledges about color management: matrix triplet, RGB, XYZ, Lab, color space, primaries and colorimetric profile... wise readers can read the B. Lindbloom's website [1] (http://www.brucelindbloom.com)
- The histogram display as well as the browser can be configured: by default, the displayed values take into account the "Output profile". You can change this behaviour going into "Preferences" and checking "Use working profile for main histogram and navigator"
- Of course, this document is far to be exhaustive, the subject is complex.

Initial raw processing before RGB conversion

Read the raw file and use its data

- Le first step that is essentially carried out on a Dcraw basis (thanks to D.Coffin) consists in reading all types of raw files, with their proprietary coding: data depth, 12 or 14 bits, sensor saturation, white balance, ... and of course rggb or rgbg data;
- the default white balance is the one chosen by the user on his/her camera for the shooting!
- the interpolation (AMaZE, AHD, DCB...) takes place after by rgb data modifying (the image can be visually evaluated on a screen): the interpolation shouldn't modify the colorimetry (or a very few) and that is the case for all the existing interpolations in RawTherapee (deltaE94 due to interpolation is approximatively 1, so negligible), in the other hand, at the limits (upper highlights ...) artefacts may appear for some of them.
- in order to give them a more realistic look, these rgb data are modified by either a color matrix (from Adobe), or a DCP profile (from Adobe or RawTherapee), or an ICC input profile, we will below come back about these ICC profiles, their elaboration and use;

- rgb values are without any color space this point is essential (the choice sRGB or AdobeRGB for example suggested by the camera is for JPG files sonly)
- the color management is processed: a) partially with LCMS2, that made great improvements, and can now work using floating point avoiding colors cliping; b) directly by computing (conversion of XYZ, Lab, RGB, gamma, etc.)

ICC input profiles: elaboration, use, gaps

- these profiles apply either like external profiles after RGB conversion (as does Capture NX2) to the rgb data, so theoretically without rendering intent (https://en.wikipedia.org/wiki/Color_management#Rendering _intent) (relative, absolute, perceptual, saturation). They modify Lab values but not rgb ones histograms remain the same, or like internal profiles as does RawTherapee (theoretically without rendering intent (https://en.wikipedia.org/wiki/Color_management#Rendering_intent));
- they try to reduce the gap between their original values (the sensor ones) and a target value, theoretically perfect;
- they are theoretically matching: a given illuminant (D50, C, shadow,...), test pattern shooting conditions, lens shooting...
- nevertheless, we can without any problem, use them as long as we stay approximatively in the same environment, e.g. flash instead of daylight, D55 instead of de D50...
- elaboration:
- 1. shoot a test pattern in ideal conditions matching the planned use (outside, studio, ...);
- 2. the wider the test pattern gamut is, the better the result is, e.g. the ColorChecker24 is close to sRGB, even if it gives good results in usual cases, how could it efficiently evaluate real colors that are out sRGB range (flowers, artificial colors ...)?
- 3. the higher is the test pattern cells number, better the result is (better profile guiding)
- 4. The 468 colors test pattern, designed with a colleague "Rouli", is over WideGamut for some colors, and owns low luminance values.
- 5. for your information, about my D200, deltaE94 results gotten from my 468 colors test pattern shoot (NEF) on which I apply the input profile or matrix: a) original color matrix (Dcraw) (or results obtained with Camera Raw 6.6 and DCP) profile): average=4.37, standard deviation=1.82, maximum=13.75; b) ICC profile, stored in the "Iccprofile" folder, elaborated from the ColorChecker24 close to sRGB: average==3.66, standard deviation=2.08, maximum=11.28; c) ICC profile, elaborated by myself from the 468 colors test pattern with very wide gamut, close to WideGamutRCB: average==2.05, standard deviation=1.4



Test pattern 468 colors, for calibration

- close to WideGamutRGB: average==2.05, standard deviation=1.44, maximum=8.8; d) of course, most of the time, the profile made with a ColorChecker24 (ICC or DCP) wil be enough!
- 6. shoot the test pattern at about 12 o'clock in direct sunlight (a), or in cloudy weather (b), or in the shadow (c), or with tungsten light (d), or with a Solux studio light (e) that has a spectre very close to the daylight, or (f) another light that match your needs:
 - 1. set the camera in manual mode (exposure...);
 - 2. set the white balance on a) 5000K (or similar to "sunlight"); b) 6000K; c) 8000K; d) tungsten 2850K; e) Sun 4700K; f)...;
 - 3. ensure there is no reflect;
 - 4. take several photos each 1/3 EV;
 - 5. ensure the most perfect possible exposure equality between the test pattern centre and four sides.
- 7. save the raw files in a folder called RawTherapee

- 8. open the raw with the "neutral" pp3 profile et choose "Prophoto" as "working profile", in "Input Profile", choose "No profile"
- 9. evaluate the exposure from one of the test pattern grey cell which luminance lie between L=40 et L=60 and study the exposure difference between the test pattern dark sides; choose the shoot with the best compromise; do the shoot again if necessary.
- 10. set the exposure with "raw white-black point" "white point: linear corr.factor" ensuring that "White point HL preserving corr(EV)" is set to zero, in such a manner that the L value of the grey reference cell being the closest possible of the reference.
- 11. tweak the white balance with SpotWB, choosing a grey cell (20 < L < 80) of which the "a" and "b" values are the closest to zero (for this operation the "a" and "b" values must be inferior to 0,5; otherwise, if the cell has "a" and "b" values close to 1, tweak with the "temperature" and "hue" sliders (see below the notes about white balance) in such a manner to arrive to the reference cell Lab values.
- 12. then click on "Save reference image for profiling"
- 13. use your "profiler" which will use, according to the maker, either the spectral values or the Lab or XYZ values: generate a reproduction type profile for the illuminant matching the shooting (D50, D65, Solux, etc.).
- 14. of course, some « profilers » (Profilemaker5, etc.) allows to elaborate profiles that are not « reproduction » profiles that minimise the deltaE94, but profiles that will give a specific rendering (portrait, landscape, etc.) by acting on the contrast curve and on the differentiated chromaticity between pastel and saturated tones. These profiles are working, but in my point of view, are out of RawTherapee's mind by introducing from the processes beginning, hue and contrasts gaps that can't be corrected any more by the various algorithms. Nevertheless, this choice is possible.

White balance

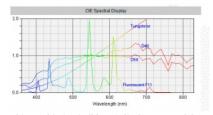
White balance gaps

- The white balance is really operational only with the use of "SpotWB" on a perfect grey ("a" and "b" Lab values set to zero), but as it is quasi impossible to put a grey chart on each shot (otherwise to do it like Alfred Hitchcock), if you want to tweak the camera white balance, you have generally to use the "Temperature" and "Tint" sliders;
- But the sliders have a range from 1500K to 25000K, and the used computing basis matching the D
 illuminant (Daylight) is not valid under 4000K (computing is only an extrapolation)
- Moreover, data become false if the illuminant is different from the day light (D), for example the illuminant "Blackbody" or "Fluorescent"
- Then, show caution, great caution, when you are out of the D illuminant (daylight) and for temperatures under 4000K

Principles

Shooting: When shooting, two basis options are available to the user: a) To work in raw mode, in this case, mistakes are authorized and a retouches are possible with a raw processing software, for example RawTherapee; b) to work in JPEG mode, in this case, if the white balance choice on the camera is different that the real light conditions, the retouch is more difficult. We will give more importance to the raw mode study.

However a camera owns (essential for JPEG, useful for raw) several white balance settings. Of course theses functionalities vary from a model and maker to one another, but we often find:



This graphic (B.Lindbloom) displays spectral data for 4 illuminants I chose arbitrarily: D50 (5000K), D40 (4000K), A (tungsten - 2850K) and F11 (fluorescent). It shows the obvious difficulty to extrapolate the D illuminant...

- An "auto" mode: the camera electronic decide which value is right from home-made algorithms. This mode generally works pretty well, except when there is strong dominant colors.
- A "manual" mode where the user, for some brands, enters a temperature value, e.g. 7000K. This choice is made by a confirmed user, on his/her experience basis.
- A "pre-set" mode where the user can choose amongst several situations predetermined in "the factory": "sun", "shadow", "cloudy", « flash », « incandescent », « fluorescent »,...
- Note: each brand and model has its own specificities, e.g. for the fluorescent mode: a) Canon equips its cameras with only one mode; b)Fuji gives 3; c) Pentax gives 3; d) Nikon gives up to 7 (D3S, D3oo,...); e) etc.
- Note also: the "flash" value differs from one brand to another one, e.g. for: a) Nikon D300, the flash illuminant matches roughly 6400K; b) Leica R9, the flash illuminant matches roughly 5500K; c) Sony A900, the flash illuminant matches roughly "shade", that means about 7000K...
- of course in most of the cases, the choice of pre-set is pretty obvious, but in some other cases the user won't know what to chose... Indeed, during an exhibition, a visit to the museum, etc. what is the lighting in use?
- All these settings act on the channels multipliers;
- At the end, during a raw processing, this choice will be displayed in RawTherapee as "Camera".

Raw processing

RawTherapee allows 5 possibilities:

- « Camera » : the software uses when existing the shooting EXIF data
- « Auto » : the software values the white balance at the "average" data on a theoretical neutral grey basis:
- « pre-set settings »: like Shade, Daylight, Fluorescent, etc.
- « Custom »: the user can choose Temperature and Tint (see the section White balance gaps)
- « Spot WB »: the user choose a neutral-grey area as a reference. This almost means a neutral-grey chart must be present when shooting.

Colored object - illuminant - observer

Simplifying, colorimetry can be synthesised in three kinds of data:

- Colored object: characterized by its pigments (red, blue,...) whatever the lighting source is. It can be evaluated with a spectrometer, that will give an unique representation of it
- Illuminant or Lighting source: this data characterize the nature of the illumination source (sun at noon, shade, flash, tungsten, fluorescent...). It is valued by two data: its spectral power distribution and its correlated temperature; caution! two lighting sources that have the same temperature are not identical because their spectral data are different.
- Observer: Do the following test, observe the color of a painted wall, then the one of the sample that was used to make the choice, you will notice that the color feeling is different. CIE (Commission Internationale de l'Éclairage) defined two "observers", 2° and 10°, characterized by the angle under which the color is observed.

To compute the values observed from a colored object, under the X illuminant, it is necessary to take into account in the computing the object, illuminant and observer spectral data. These values XYZ are specific to the illuminant temperature (e.g. 6500K); at a software level, (Photoshop, RawTherapee...) the standard is D50. So, it is imperative to proceed to a chromatic adaptation from the XYZ values (illuminant) to the XYZ (D50). For that, several methods exist from the less to the most efficient: Von Kries, Bradford, CIECAM02.

The Daylight illuminant

This illuminant has been object of many studies by Judd, MacAdam and Wyszecki, from several hundreds of examples. In short, the "D" illuminant is the sum of three parts: S(lamda)=So(lamda) + M1*S1(lamda)+M2*S2(lamda)

- 1. one "fix" part, So that is the average of all the tested examples;
- 2. one first "variable" part, S1 that matches the "blue/yellow" variations due to the presence or not of clouds or to the position and the intensity of the direct sun;
- 3. one second "variable" part, S2 that matches the "pink/green" variations due to the presence of humidity in the form of vapour or fog ...;
- 4. in practical terms, that can be written in a "simple" formula that determines two values x_D and y_D on a temperature and an illuminant basis
- $x_D=0.244063 + 0.0991*103/T + 2.9678*106/T2 4.6070*109/T3$ for 4000K < T < 7000K
- $x_D=0.237040 + 0.24748*103/T + 1.9018*106/T2 2.0064*109/T3$ for 7000K < T < 25000K
- $y_D = -3.0 \times x_D + 2.87 \times D 0.275$
- these formulas are used to calculate the parameters M1(x_D,y_D) and M2(x_D,y_D) of S(lamda)=So(lamda) + M1*S1(lamda)+M2*S2(lamda) it was not the case so far in RawTherapee. Previously, in RawTherapee, the xD and yD values were directly used to calculate the channels multipliers, now, these values are used to determine the illuminant spectral values at the T temperature. It is only later on that second values are calculated that allow to determine the multipliers.

Two things can be immediately noticed:

- it does not exist any "Daylight" reference under 4000K, the previous formula used in RawTherapee (from 1200K to 4000K) has been "invented" for Ufraw;
- nothing prevents to modify RawTherapee to switch the maximum value from 12000K to 25000K (done since 2013)

"Blackbody" illuminant and "A: tungsten"

The A illuminant from CIE is used to represent the typical light of a domestic bulb tungsten filament. Its relative spectral power distribution is the one of a Planck radiator at an approximative temperature of 2856K. The A illuminant from CIE can be used in any colorimetry application involving the use of an incandescent light, unless there is specific reasons to use another illuminant.

The "Blackbody" illuminant can be caculated with the Planck formula that is the generalisation of the illuminant on a T basis:

■ S(lamda)= c1 * pow(wavelength, -5.0)) / (exp(c2 / (wavelength * blackbody Temp)) – 1.0);

where the 2 values c1 and c2 match to: a) c1=2*Pi*h*c2 h=Planck constant c=light velocity; b) c2=h*c/k k=Boltzmann constant

This illuminant correctly joins up at roughly 4000K with the "Daylight" illuminant with minimal gaps. I accordingly chose for RawTherapee to use the "Blackbody" illuminant under 4000K and up to 1500K (it seems that "ACR" did the same choice).

The Fluorescent illuminants

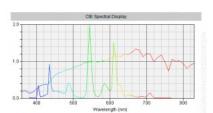
Standardisation has foreseen 12 illuminants of his type, matching the lighting tubes sold in the shops:

- F1: Fluorescent daylight 6430K
- F2: White fluorescent cold 4230K
- F3: White fluorescent 3450K
- F4: White fluorescent warm 2940K
- F5 : Fluorescent daylight 6350K
- F6: White luminous fluorescent 4150K
- F7: D65 simulator 6500K
- F8: D50 simulator 5000K
- F9: White cold deluxe 4150K
- F10: Philips T85 5000K
- F11: Philips T84 4000K
- F12: Philips T83 3000K

These illuminants (see white balance gaps) have a very different spectral power distribution between them and between the "daylight" and "blackbody" illuminants. So it is not advised to replace - when a lighting is fluorescent - the illuminant in question (e.g. F11 4000K) by an equivalent "Daylight 4000K".

Actually, the white balance calculates on the spectral data basis, two coefficients xD, yD that modify the channels multipliers: this calculation acts as an averaged integral calculus. Indeed, the "average" white balance will be exact, but the spectral data peaks or gaps, compared with a theoretical ideal (blackbody or daylight) will locally bring, for some colors, hue gaps more or less important.

It exists a concept called "CRI = Color Rendering Index" that translates the lighting source quality. This "CRI" is a number equal to 100 for a perfect source. We consider that values over 90 give good results. e.g.:



On this graphic, we can see that for a same temperature, the Daylight 4000 illuminant (continuous spectre) is very different from Fluorescent F11 (4000K)., so, it will give a different colors rendering.

- Fluo F4 "warm white": CRI=51
- Clear Mercury Vapor: CRI=17
- several LED with "CRI" between 50 and 96
- Solux 4700: CRI=92
- etc.

This concept is embedded (not used yet) in RawTherapee with the following choices:

- 20 reference colors of which, 8 "standard" from the Colorchecker24, 4 skin hues, 4 greys (white-black), 3 blues
- use of CIECAMo2 for the chromatic adaptation
- use of CIE Lab for deltaE calculation

To remedy it (partially) it is enough to create an input profile with the desired lighting source and the matching spectral data.

Other illuminants

Other illuminants exist, stemmed from the A illuminant and close to "Daylight"

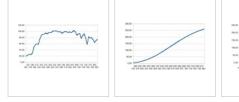
- B and C that I did not embedded in RawTherapee, but it is possible to do it
- an illuminant with equal energy: "E"
- The studio bulbs illuminants (movies, stage lighting, museums, photo studio, etc.) that are named HMI, GTI, Solux 4700K, JudgeIII, Solix4100K, Solux3500K, etc., they are embedded in

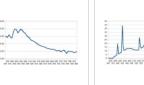
RawTherapee.

- The LED illuminant, these lamps often have large gaps in the blues. Some of them own very satisfactory features
- "Proprietary" flashes illuminants (Canon, Nikon, Pentax...) and studio flashes illuminants, they are generally very close to daylight, but each one at different temperatures. I did several grouping at 5500K, 6000K et 6500K, about studio flashes, it should be useful to have their features: a) theoretically it should be useful to have the spectral data for each flash (I don't have them), moreover, these data vary depending on the flash power ...: b) so, I preferred to use the equivalent "daylight".

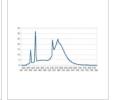
As we can see, the situation is not simple and pose many problems to the raw processing software, of which RawTherapee.











Illuminant D50 (CRI=100 Sigma=o)

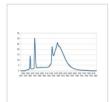
Illuminant A 2856K (CRI=100 Sigma=o)

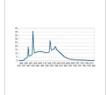
-15000K (CRI=100 Sigma=o)

Illuminant D₁₅₀ Illuminant Fluo F1 "daylight" 6430K (CRI=77 Sigma=10)

Illuminant Fluo "coolwite"4230 K (CRI=64 Sigma=12)

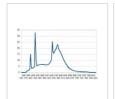
Illuminant Fluo F3 "white" 3450K (CRI=60 Sigma=12)



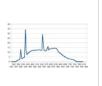


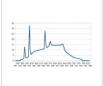
Illuminant Fluo F4 "warm white" 2940K (CRI=54 Sigma=11)

Illuminant Fluo F5 "daylight" 6350K (CRI=74 Sigma=12)

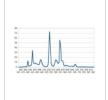












Illuminant Fluo F6 "Lite white" 4150K (CRI=61 Sigma=13)

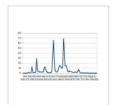
Illuminant Fluo F7 "D65 simulator" 6500K (CRI=90 Sigma=2)

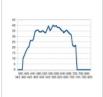
Illuminant Fluo F8 "D50 sylvania F40" 5000K (CRI=94 Sigma=1.4)

Illuminant Fluo F9 "Cool white delux" 4150K -4330K (CRI=89 Sigma=2)

Illuminant Fluo F10 "Philips TL85" 5000K (CRI=72 Sigma=11)

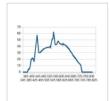
Illuminant Fluo F11 "Philips TL84" 4150K -4000K (CRI=77 Sigma=9)



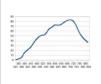


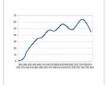
Illuminant Fluo F12 "Philips TL853" 3000K (CRI=72 Sigma=8)

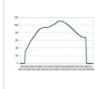
Illuminant Lamp HMI 4800K (CRI=97 Sigma=1)

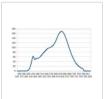












Illuminant Lamp GTI 5000K (CRI=90 Sigma=2)

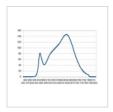
Illuminant Lamp Judge III 5000K (CRI=92 Sigma=2)

Illuminant Lamp Solux 3500K (CRI=95 Sigma=2)

Illuminant Lamp Solux 4100K (CRI=92 Sigma=2)

Illuminant Lamp SoluxNG 4700K - 4480K (CRI=97 Sigma=1)

Illuminant Lamp LED LSI Lumelex 2040 -3000K(CRI=90 Sigma=2)



Illuminant Lamp LED CRS SP12 WWMR16 -3050K(CRI=94

Sigma=3)

Algorithm

I use the basic algorithm "Daylight": a) calculation of the x_D and y_D values that are given as parameters to M1 and M2 (S(lamda) = So(lamda) + M1*S1(lamda) + M2*S2(lamda) from which we differentiate Xi,Yi,Zi by matrix calculation [XiYiZi]=[observ2°xyz][S(lambda)], then we calculate the channels modifiers modifications by a simple matrix calculation [mulrgb]=[sRGBd65_xyz]*[XiYiZi] (lebarhon: no b)!)

■ I used works of John Walker (public domain), and B.Lindbloom - increasing the accuracy and the spectral range - mostly the "Spectrum_to_xyz" function, as known as "CIE_color_match" that converts spectral data (350 – 830nm) of a color or an illuminant into xBar, yBar, zBar values (via the Observer 2° data). We get the output values x and y.

About the blackbody, I use the Planck formula: the seem between the two formulas is very good with a very slight gap of the xD and yD values at 4000K (that we can make out with the histogram between 3995K and 40005K): a) daylight 4000K: xD=0.382 yD=0.383 (for your information for 4500K: xD=0.362 yD=0.370, for 7000K xD=0.30 yD=0.32, for 25000K xD=0.25 yD=0.25); b) blackbody 4000K: xD=0.381 yD=0.377;

For the other illuminants, the works I previously did on the calibration (test pattern 468 colors) led me to search for (and find out) the illuminants spectral data I selected (tungsten, fluorescents, HMI, GTI, Solux, etc.)

rgb ==> RGB conversion - Working space "Working Profile"

This "conversion" converts rgb data (without color space) into the working space chosen by the user.

These working spaces are 8 of them (which sounds more than enough, or even too much...): sRGB, AdobeRGB, Prophoto, Widegamut, BruceRGB, BetaRGB, BetaRGB, Rec2020. amongst these 8 profiles, 5 have a wide gamut: BetaRGB (origin B.Lindbllom), BestRGB, Rec2020, WideGamut and Prophoto.

During this conversion, an internal gamma is set by RawTherapee that always is "gamma sRGB" i.e. "gamma=2.4 and slope=12.92" (like Lightroom, see later on)

Note that other software made other choices:

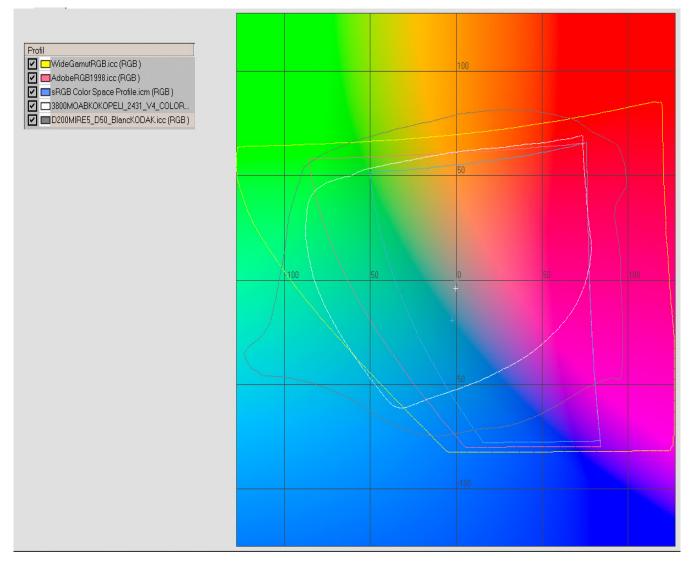
- Adobe with ACR suggests 4 choices (AdobeRGB, ColorMatch, Prophoto and SRGB)
- Adobe with Lightroom: no choice, but a modified Prophoto space with a sRGB gamma (Melissa)
- DxO: no choice, but a Adobe space
- NX2: choice in the available output spaces

Which one to choose? Wide ranging debate where small spaces supporters conflict with wide spaces ones... between data lost and false or imaginary data. Indeed, the widest space (Prophoto) houses, by construction, invisible or even imaginary colors. Moreover, in the blues area, it may under some circumstances, generate artefacts.

A good color space would have a shape taking into account the gamut that minimize the waste of space... which means the space should be chosen depending on each image, a too wide space could be generating unsaturated colors in extremes cases (in the working profile, but which will be reproduced during the output conversion ...). Nevertheless, theoretically, the color management must make this choice transparent.

My answer is pragmatic, choose the space that match the best! but, on what basis?

• Do you make essentially prints with a printer that use a CMYK driver? In this case it is not really useful to choose a wide gamut profile.



- Do you make prints with a high quality ink-jet printer? In this case, it is better to choose Prophoto (this kind of printers have a gamut, that for some colors, is wider than WidegamutRGB), as working profile, but also as output profile, and of course to choose the right printer profile... on this graphic, we can see the gamut of a) 3 usual color spaces (sRGB in blue, AdobeRGB in pink, WideGamut in yellow), b) the ICC profile ffor my D200 in grey, c) the Epson printer profile and its "3800MOABKOKOPELI_2431_V4" wide gamut paper.
- Do you have a very high quality monitor which has a gamut close to AdobeRGB or WideGamutRGB, in this case take a wide gamut profile.

A pretty simple way to evaluate the minimum profile is to use the statistics given by "vibrance" in debug mode (with verbose=true). In the RawTherapee.exe window you will see a message:

- Gamut: G1negat=x iter G165535= y iter G2negat= z iter G265535= w iter
- if a value (x or y) over 0 appears for one of the two G1, that means the initial image (with the controls made before: contrast, exposure,...) exceed the chosen space gamut in the working profile.

- if a value (z or w) over o appears for one of the two G2, that means the saturation set by "vibrance" exceeded the gamut.
- It is up to you to choose if you want to keep these values (see above) or to put them into the gamut (vibrance do it for you, as well as "avoid color cliping" + "enable saturation limiter"), but you 'lose' colors!
- the conversion rgb==>RGB allows when working in "float" to preserve the essential negatives and over 65535 values.

Work in the color space chosen by the user (working space "Working Profile")

RawTherapee made the good choice to work in Lab mode (or its flavours Luv or Lch) or in CIECAMo2 mode and with real numbers. This allows to preserve the best colors and gamut.

Functions available in "Exposure" does not modify the hue, except "Saturation" that make it by dysfunction of the Lab mode (see about that the Munsell correction section

Same thing for "Lab adjustments", for the "Saturation" slider, along with the "a" and "b" curves and all the curves that control the chromaticity.

"Channel Mixer" and "HSV equalizer" modify strongly the colorimetry, to be used with all the knowledge of the consequences on the colorimetry. The functions: contrast, brightness, tone curve,...may quite strongly modify the gamut (see above the control with the statistics given by "Vibrance")

What happens when we are in a "Working Profile" or when we change of "Working Profile" or the settings?

When we are in a narrow profile like sRGB, it may seem obvious that the colors will be limited to this profile limits! So, if an initial color (the sensor one, slightly modified by interpolation) is inside sRGB, what happens when we tweak sliders and curves?:

- in RGB mode (Exposure), if saturation or lightness or contrast is/are modified, we work in rgb mode or its linear derived "hsv". This means that the resulting effects will depend on the working space ((sRGB, AdobeRGB, Prophoto), the results will be different when we move from "sRGB" to "Prophoto", even if we stay inside the gamut limits.
- in Lab mode, if lightness, chromaticity, contrast or curves are modified, we modify directly the L, a, b, (or C,h) values. This means that if we stay inside the gamut limits, the image will be the same (within the limits of the color management)

When exceeding the gamut - that is, the original image - either acting the sliders or the curves, what happens?

- let's take an example, an original color is L=27 a=2 b=-75; this color is in the Prophoto space and is worth R=42 G=52 B=158, and in sRGB R=-85 G=69 B=184 (negative R means out of gamut). If we change the working profile, it is obvious this color can't be restored; the XYZ conversion will give L=32 a=21 b=-67 and R=0 G=69 B=184. The color will have a different aspect because it is completely impossible to reproduce it in a smaller color space.
- second example, a color inside sRGB: L=40 a=42 b=-44, that is in RGB (sRGB) R=133 G=66 B=166, and in RGB(Prophoto) R=102 G=64 B=140: a) if we apply saturation (+30) in exposure, the values become Prophoto L=36 a=49 b=-49 and sRGB L=38 a=47 b=-47, because we act on the RGB values; b) if we apply chromaticity in Lab adjustments; Lab values become L=40 a=55 b=-58 as well in sRGB than Prophoto, because we stay inside the sRGB gamut; c) if we choose a color, close enough

to the sRGB gamut limits, but inside the gamut: L=40 a=63 b=37 and we apply chromaticity +30, Lab values become - Prophoto L=40 a=81 b=49 - we can notice that the hue is preserved (arctg(b,a)), and in sRGB L=44 a=69 b=50 - the hue isn't preserved.

What does "Avoid color shift"?

If we apply to c) above, "Avoid color shift", the system will try to preserve the hue, by applying a relative colorimetry Lab values become L=40~a=65~b=38 (which leads to the RGB channel G=0).

"Avoid color shift" realises two things:

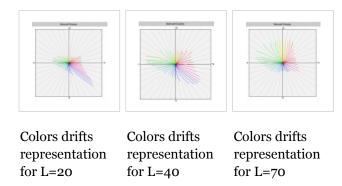
- try to put in the working profile gamut data out of this gamut, giving priority to a relative colorimetry.
 RGB negatives values are detected and the chromaticity (and luminance as well) are modified to reach the value o.
- apply a "Munsell" correction.

The "Munsell" correction

Goal: to give the user possibility to automatically correct the colors that are sometimes false in Lab mode when modifying the saturation, mostly in the blues-crimsons, reds-yellows and the greens (Munsell type correction)

This goal requires the creation of 160 look-up tables with floating-point precision (LUTf):

- These LUTf give to each color (in Munsell sense), each luminance, hue values on a chromaticity basis: 2, 3 or 4 points are set for chromaticities [0..180] from 5, 45, 85, 125 or even 139 when possible. Intermediate values are linearly interpolated. These LUT cause a minimal memory occupation: each one has between 45 and 140 entries, that means about 16000 entries at the total...
- The LUTf are realised for four critical areas, where the drift compared with the Lab mode is important (blue-crimson, red-yellow, green, red-crimson), for the other areas the gaps are low, in any cases very much under the matrix and ICC profiles possibilities.
- These LUTf have the C illuminant as basis (slightly different of D50 or D65), but due to the fact that we work on a gap basis and not on absolute values, the calculation error is very low (less than 1% of the correction, in any cases very much under the matrix and ICC profiles possibilities).
- Theses corrections are quick, nevertheless they are increasing the "vibrance" or "Lab adjustments" processing time of about 10%.
- If "verbose" option is enabled, we can see appear for each correction type, the concerned number of pixels and a rough idea of the correction size (in radians). For correction values in radians, the color gap in deltaE94 is the following:
- 1. correction=0.4 rad deltaE94=12 for blue-crimson
- 2. correction=0.2 rad deltaE94=8 for red-yellow
- 3. correction=0,1 rad deltaE94=3,7 for green
- 4. correction=0.05 rad deltaE94=2 for red-crimson
- 5. the maximum common values of the drift can reach about from 0.05 radian to 0.15 radian for average modifications of saturation, values from 0.2 to 0.25 radian are not exceptionnal.
- 6. for information, a deltaE94 less than 1 is negligible, it is noticeable around 2 or 3 and very important at 8 or 11.



Output space "Output Profile"

Space choice

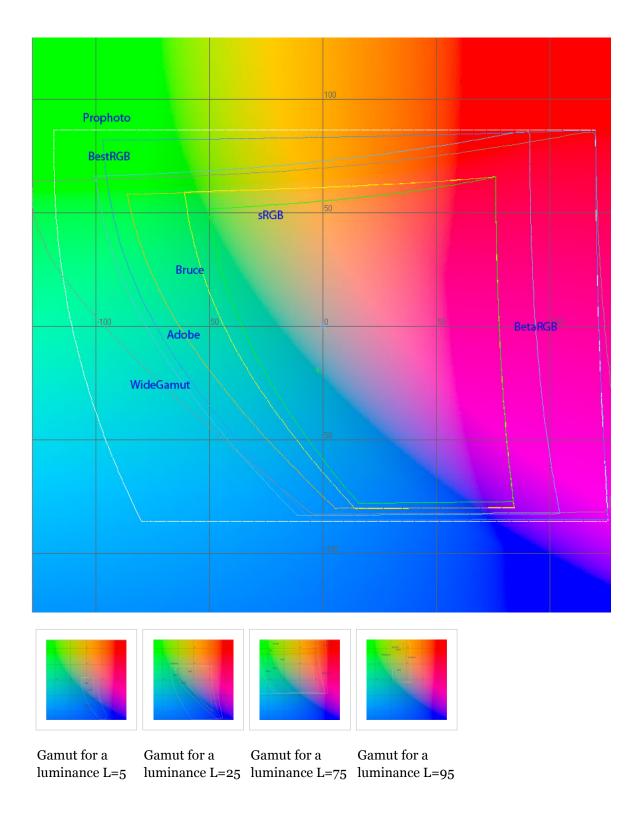
The first thing to examine is: what are the output profiles that are installed on your computer? This depends of: a) the operating system (on the face of it, Linux does not install any profile); b) the other graphical software that are installed (Capture NX2, PhotoShop CS, DxO, etc.), each one install proprietary profiles, for example NX2 installs some NKsRGB.icm NkAdobe.icm, etc. that are copyrighted...; c) profiles that you can have downloaded on the web, for example on the Adobe or B.Lindbloom websites.

On principle I would recommend to check the installation or to install the output profiles matching the working profiles - that are *.icm or *.icc files physically present on your computer and that have nothing to do with the calculation matrix of "iccmatrices.h". If these files are missing, TIFF or JPEG output cannot be done towards these profiles, but will be by default (if the RT_sRGB.icm file is present) done towards the SRGB output space.

These profiles have the following names (we can find others that have the same features or close features), they are generally copyrighted and consequently cannot be shipped with an Open Source software without authorization. They are available on the Web: ProPhoto.icm; SRB Color Space profile.icm; AdobeRGB1998.icc; BestRGB.icm; BetaRGB.icc; Bruce.icm; WideGamutRGB.icc. Of course you can install some others like CIE.icc; Colormatch.icc; etc.

They must be installed in the RawTherapee "Iccprofiles/output" folder or in \windows\system32\spool\drivers\color for Windows and \usr/share/color/icc for the other systems.

When you choose an output profile, e.g. AdobeRGB1998 and the working profile Prophoto, LCMS2 will convert with a Colorimetric intent (chosen by default in the RawTherapee options: relative, perceptual, ...) the RGB data from the working space to the output space.



Of courses, remarks about the choice of the output color space are similar to the ones about the working space (print, screen,...).

It you want to print with a high quality ink-jet printer (remind: RawTherapee has no print module so far), you have to use a third party software (Photoshop...), in this case I strongly recommend an output profile of Prophoto or WideGamut type.

Take care however, JPG outputs, therefore 8 bits, are almost incompatible - important risk of posterization - with wide gamut spaces (Prophoto, WideGamut...).

Shipped files

Because of the copyrights, I shipped specific files with more detailed LUT that should bring less posterization in the shadows. These files are a by-product of Output Gamma (see below)

- RT_sRGB.icm: similar (primary) to sRGB.icm standard with internal gamma close to sRGB: g=2.40 slope=12.92
- RT_sRGB_gBT709.icm: similar (primary) to sRGB.icm standard with internal gamma BT709: g=2.22 slope=4.5
- RT_sRGB_g10.icm: similar (primary) to sRGB.icm standard with internal gamma linear: g=1.0 slope=0
- RT_Middle_gsRGB.icc: similar (primary) to AdobeRGB1998.icc standard with internal gamma close to sRGB: g=2.40 slope=12.92
- RT_Large_gsRGB.icc: similar (primary) to ProPhoto.icm standard with internal gamma close to sRGB: g=2.40 slope=12.92 (close to "Melissa" used by Lightroom)
- RT_Large_gBT709.icc: similar (primary) to ProPhoto.icm standard with internal gamma BT709: g=2.22 slope=4.5
- RT_Large_g10.icc: similar (primary) to ProPhoto.icm standard with internal gamma linear: g=1.0 slope=0
- Rec2020.icm: new primaries large gamut with internal gamma BT709: g=2.22 slope=4.5

RawTherapee gaps

The user can easily note that the output is slightly different from preview. This is not an output default but is due to the curves elaboration that badly take TRC notion into account (ICC profiles embedding to modify from the interior the tone rendering).

It is one of the reasons that led me to give possibility of "adjustable" outputs:

- 1. either by choosing an output profile with another gamma;
- 2. either with an adjustable gamma (gamma and slope);
- 3. either possibly to realise a linear output and adjust it in Photoshop...
- 4. (see further the histograms sketches)

Output Gamma

From my point of view, Output Gamma is one of the key points of a successful TIFF or JPEG output, for several reasons: since the adding of icc/icm profiles above - that are directly derived and elaborate by Output gamma - this option shows a less important interest because the use of these "new" pseudoprofiles Prophoto and SRGB brings similar advantages to Output gamma when selecting the pseudo sRGB and Prophoto! (see above). This option allows to partially compensate the RawTherapee gap (difference between output/preview)

The ideal would have been to put Output gamma in first process, before the Exposure, Highlights reconstructions, Shadows/highlights, etc. tabs; but, and it is what I thing is a RawTherapee gap, this modification turned out to be impossible without bringing important artefacts: the different RawTherapee pipelines are overlapping and colorimetry in the processing initial part, looks more like do it yourself than professional processing...

I accordingly chose to implant this process in final stage, what is not totally incongruous (even if I think it would be better in initial stage) Output Gamma is going to allow:

- an image evaluation for software without color management
- an image modification if youy print in CMYK (with of course a third party software).

[Lebarhon: this section is indeed not clear, the writer says firstly "Output gamma is a key point for several reasons", and then "this option shows a less important interest" and where are the several reasons?

Some thoughts

The gamma acts in a pretty similar way of a combination between "Exposure curves" + "black point" + "tone curves" present in RawTherapee, but modify in a more radically way the contrast, the histogram repartition, mostly between shades and highlights, by modifying at the same time (what previous functions can't do), the TRC curves of the file headers similar to an input ICC profile. A friend of mine, photographer told me recently: "At the beginning, I thought we could simulate gamma with contrasts and tonal curves... but the result is different".

Why indeed has Adobe, with Lightroom, designed "Melissa" that is a Prophoto color space with a sRGB gamma, i.e. a linear part until r=12.92 and then a gamma of 2,4?

Why has D.Coffin embedded in Dcraw for a long time ago, a linear gamma and a varying gamma as well?

Theoretically, if the color management is perfect, whatever the output gamma is (standard, varying...) the image should be identical, because the color management use Lab (or XYZ) data and what we call PCS (Profil Connexion Space in D50); in practice, the image seems identical, but is we watch in the shades, we can see there are gaps, admittedly pretty low but enough for RawTherapee users could have said: the output image is different from "preview".

In the other hand, for software that does not manage colors like many web browsers (Chrome,...) the output image will depend on gamma; it is then important RawTherapee allow to view what the output file will be (softproofing). Today, several factors must be taken into account:

- working space ("working profile") conversion towards output space (Output profile): it may seem
 obvious that if the image has colors out of the gamut for one or two of the both spaces, if the spaces
 range is different, then the image rendering will be different;
- color management existing or not: in the case of RawTherapee or of an editor that manage colors, the gaps (with identical output spaces) will be low but nevertheless significant; in the case of software that do not manage colors, gamma impact will be very important.
- possibility to configure in the softproofing, the intent and the black point.
- Later on, it must be quite easy possible to simulate a print by converting the output towards the printer profile; note that in this case, a view of the printable colors should be an important plus.

When RawTherapee will have a function similar to Photoshop CS one "proof format-proof colors", the viewing issue of Output gamma will be solved! this function allows the "soft proofing", e.g. to simulate the look a file is going to have in a website or on a CMYK printer.

Various Output Gamma

In the drop down list you have 7 pre-defined gamma:

- BT709: slope=4.5 gamma=2.2
- sRGB: slope=12.92 gamma=2.4

- linear: gamma=1.0
- standard:slope=0 gamma=1.8
- standard: slope=0 gamma=2.2
- High: slope=3.35 gamma=1.3
- Low: slope=6.9 gamma=2.6

BT709 "is better to process shades (they will be less grey) than sRGB and all the more so, 2.2 or 1.8

"Low" will increase the contrast of the quite poor images, and allow a better post processing for over exposed images

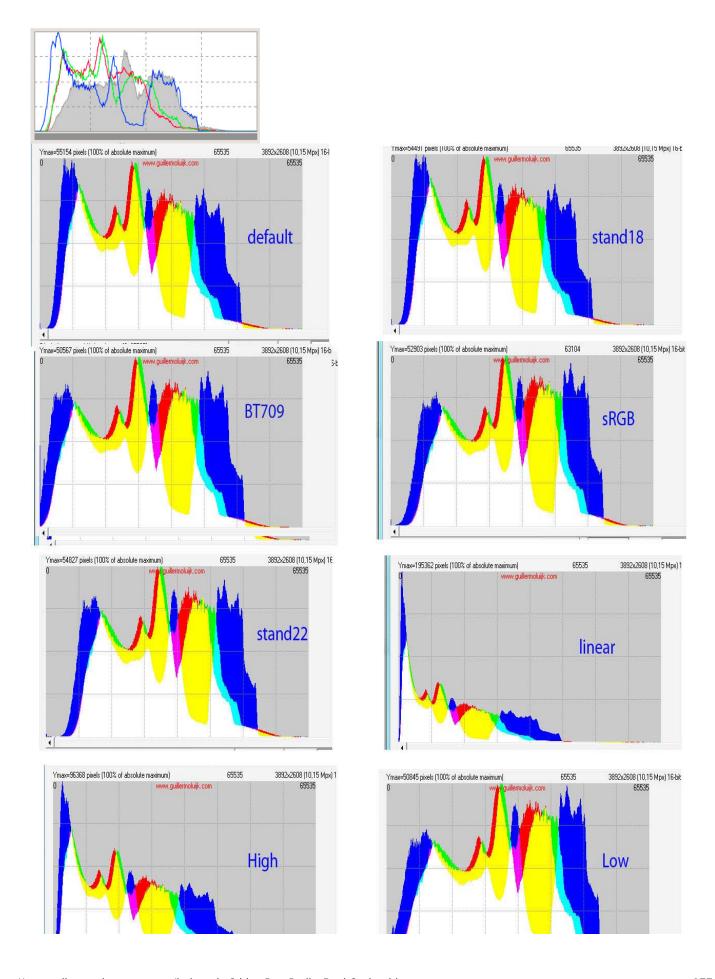
"High" at the opposite will decrease the contrast...

Linear: allows a process in Photoshop of very high dynamic images by adjusting RGB curves in Photoshop (difficult exercice...)

Moreover, you have possibility to use "free gamma" that allows to connect any slope and gamma values to a given output profile, so, you can if desired go out of the new outputs connected to the new added profiles icc/icm:

- sRGB with standard gamma 1.8
- WideGamut with gamma BT709
- sRGB with gamma: 2.2 and slope 6.5

As examples here is with the same NEF file various histograms with various gamma and as a reference the RawTherapee histogram (preview) with the same settings (working spacece Prophoto, profile "neutral", output profile = Prophoto and its varying gamma).

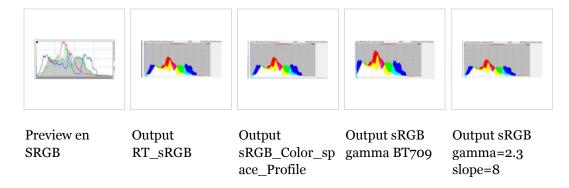




The more the histogram is shifted to the left the more the image appears dark...

Moreover, it is not because the histograms are strictly identical that the images rendering will be identical. Indeed, the "TRC" notion take place as well for the preview as for the output file. This "TRC" acts on the file headers (ICC profile) and modify the tone data. If the "TRC" value of the output file is known for certain, because it is determined by the output file features (Prophoto.icm, RT_srgb.icc,...), I think that it is not the same thing about "preview"... (see further notes about sRGB output).

Some people could be worried, is "RT_SRGB" identical to "sRGB_Color_Space_Profile" and different from the preview. Moreover, what are the effects of another gamma?



How to use it?

To satisfy a desire for simplicity, the Output profile will be a derivative of the Working profile, the box Output profile appear in grey. That means Output profile = Working profile.

For instance, you select Working profile = Prophoto and Free gamma = 2.1 and slope =4.0.

Then you validate a TIF output or towards the editor and you generate an output TIF file, with Prophoto profile and gamma 2.1 / 4.0. To open the file in an external editor (e.g. Photoshop CS), it will appear "Prefer the embedded profile: sRGB IEC61966-2.1 (RTH gamma BT709 similar to HP sRGB)" that match to the RT_sRGB_gBT709 profile but with a modification that we will examine further.

Other example, you select Working profile = sRGB and Free gamma=2.3 and slope=10.0, you are going to generate a TIFF with sRGB and gamma 2.3 and slope=10 output. To open the file in an external editor (e.g. Photoshop CS), it will appear "Prefer the embedded profile: sRGB IEC61966-2.1 (RTH gamma BT709 similar to HP sRGB)" that match to the RT_sRGB_gBT709 profile but with a modification that we will examine further.

If you enable the option (Photoshop CS): "Delete the embedded profile", the TIFF file will appear with the new RGB values due to the new gamma and slope values, but the image look will be different (file headers missing).

The algorithm uses the LCMS function "CMSToneCurve":

- the output spaces are computed from their primary (red, green, blue), e.g. for Prophoto: p1=0.7347; p2=0.2653; p3=0.1596; p4=0.8404; p5=0.0366; p6=0.0001;
- the gamma parameters are computed with the "calcgamma" function that will, accordingly with the gamma and the slope, determine 5 parameters to give to the LCMS2 right function.

So, we create a pseudo-profile, kind of RGB "Prophoto" and with a gamma matching the selected one.

But we find here, a LCMS2 gap, which creating this profile, does not write the matching profile in the file header, because it works with RGB values and not with LUT / Lab. Theoretically, it should need as much profiles with an adapted gamma and not only one. In practice, I brought a large modification to Output Gamma and worked around the LCMS2 gap, by applying - after the RGB conversion, a profile *.icc that has the same features that the *.icc or *.icm profiles used by Output Gamma but where the rTRC, gTRC, bTRC tags are computed with "calcgamma".

To improve the understanding of the TIF processing in linear mode, you can read the Dcraw tutorial from Guillermo Luijk http://www.guillermoluijk.com/tutorial/dcraw/index_en.htm (http://www.guillermoluijk.com/tutorial/dcraw/index_en.htm)

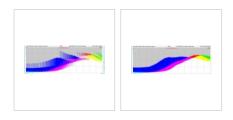
From which the necessity to have in the "Iccdirectory" folder, the "*.icc" and "*.icm" files: BestRGB.icm; BetaRGB.icc; Bruce.icm; WideGamutRGB.icc, (and thus the icc/icm files added for the pseudo-Prophoto, Adobe, SRGB in "Iccprofile/output").

RawTherapee profiles quality

The user can ask himself with reason what is the RawTherapee profiles validity (RT_sRGB, RT_Large,...)?

These profiles have the same features that the "original" ones (AdobeRGB1998, Prophoto, SrGB Color Space Profile), there is only small differences at the primaries and/or white point level. They do not have any incidence on output quality and level.

In the other hand, TRC have more detailed LUT going from 1024 points to 4096 points. This has for consequence - in the sRGB case that is the more frequent output - a histogram with much less fish bones that may bring posterization in the shades. Here is for comparison with the same image, a 16 bits histogram enlargement in low lights, between sRGB Color Space profile and RT_sRGB



partial 16 bits partial 16 bits histogram histogram sRGB_color_sp RT_sRGB ace profile

Template:Languages

Advanced Tab

Retinex

Note: This translation of the "Retinex" French page is a working document.

Generalities

While the eyes are able to see correctly the colors through a poor lighting, a colored surrounding or a veil of fog, cameras badly manage in theses conditions. It is by copying the eyes biological mechanisms to adapt itself to these conditions that the MSR algorithm (MultiScale Retinex) has been created. In addition to the digital photography, the Retinex algorithm (Retinex is the contraction for Retina + Cortex) is used in astronomy to show up information laying in the astronomical photographs, in medicine to detect not much visible structures in radiography and tomodensitometry. Numerous theories and algorithms have been working out for more than 20 years. The first experimentation has been proposed by Rahman in 1996. The approach follows the human visual perception and the Edwin Land's retinal function. In a way, this approach is pretty similar to CIECAM. This function and more particularly its general form is similar to a DOG (Difference Of Gaussian). The idea consists of characterizing the luminous information of a point from its intensity and the intensity of its neighbours. This said, this approach has no scientific basis and is only lying on experience and various empirical constants. Over the years many improvements have been added, but from my own point of view, no one is fully satisfactory. I relied on two documents:

- "Automatic Image Haze Removal Based on Luminance Component" (Fan Guo, Zixing Cai, Bin Xie, Jin Tang"[2] (http://www.mcs.csueastbay.edu/~grewe/pubs/DistSensorNetworkBook2011/Atmosph ere/HazeREmoval_SingleImage_Luminance.pdf)
- "Retinex Algorithm on Changing Scales for Haze Removal with Depth Map" (Weixing Wang, Lian Xu)"[3] (http://www.sersc.org/journals/IJHIT/vol7_no4_2014/30.pdf)
- and from some programming tricks inspired from "2003 Fabien Pelisson
 Fabien.Pelisson@inrialpes.fr>"

The use of Retinex might be beneficial to images processing:

- that are hazy, misty or having a veil of fog
- with important luminance gaps
- where the user looks for special effects (tone mapping ...)

Retinex at the beginning of the processing

It is the algorithm described in this page, with its limitations, advantages and drawbacks.

Retinex in "Wavelets"

I installed 2 possibilities to enable Retinex in the Wavelet process: Wavelet levels/fr. Some of the quoted limitations are no more there!

Quick comparison between the both versions "Retinex in Wavelets" and "Retinex at the beginning of the processing":

Wavelet_levels/fr#Avantages_.28.2B.29_et_inconv.C3.A9nients_.28-.29_de_Retinex.2Ffr_par_rapport _.C3.A0_Retinex_in_wavelet

Imposed limitations by RawTherapee

The basic algorithm impose a reference to the whole image, and not to a crop or a reduction like in the RawTherapee processing. This limitation imposed by the gaussian function causes several consequences:

- The processing has been shifted from its dedicated place, that should have been close to "Lab Adjustments" near the beginning of the RawTherapee processing that is necessarily (except if somebody has an idea to do differently) a raw process. As a consequence, non raw files (TIFF, JPG, ...) can't be processed with this algorithm? This problem should be solved soon (November 2015 ??).
- The second consequence of this position is that the characteristics of the raw data situated just after demosaicing are very different that the ones situated downstream: no gamma, no gamut limitation, no white balance, no RGB conversion... so we must expect artefacts and poor luminance and color rendering.
- The third consequence is the system response time which imposes to re-actualize the whole process after each change in the settings.
- The fourth, sometimes, will need a white point modification, in order to avoid colors distortions (for example: a magenta sky).
- But advantage, this process being situated before "Denoise", the noise reduction will be here fully effective.

Nevertheless, despite these handicaps, as we will see it further, results are more than satisfactory as well about processing time that in contrast and colors rendering, particularly by joining Retinex and Wavelets actions.

Algorithm principle

For more information, see the "pdf" documents given in the "Generalities" section.

To elaborate a "Transmission Map" file obtained by

- making the difference, for luminance only, between each pixel logarithm of the input image and the neighbouring pixels logarithm of the matching gaussian image. The standard deviation (gaussian function) used here is very high usually in RawTherapee sigma values from 0.5 to 5 are common here the values range from 10 to 280 (Single Scale Retinex), it is up to the user to choose.
- modifying the input image luminance distribution by
 - applying a gamma before the "Transmission Map" file creation
 - applying an inverse gamma to restore the input image characteristics

This modification of the distribution allows:

- to change the image tonality
- to modify the "Transmission Map" file to take into account for example the under or over exposed areas.
- applying several time (Scale) 3 times, not modifiable by the user this algorithm with an addition using empirical coefficients (Multi Scale Retinex). (Low scale values increase the apparent contrast, but give a perspective look to the image, high scale values make the image more natural but they have a tendency to increase the noise).

- a variation of MSR by the user according to the desired effect (Uniform, Low, High, Highlight)
 - Uniform: Strive to process low and high intensities in an equilibrated way.
 - Low: Improve low intensities areas
 - High: Improve the rendering of the more exposed areas
 - Highlight: Improve the rendering of the highlight areas that may become magenta, but can bring strong artefacts.
- choosing the colorspace
 - logarithmic L*a*b*
 - logarithmic HSL
 - linear HSL this version doesn't comply with Retinex algorithm but in some cases it allows a more suitable image processing.

So, we get a logarithmic distribution "Transmission Map" – except in linear mode! - that we will be able to "subtract" to the input image, either to get an image theoretically free from mist and from veil fog or to get special effects. This distribution owns approximatively a gaussian distribution with a minima (minT) about, according to the images – higher for images with veil – from -10 to -40, an average close to from -1 to +2, a standard deviation often about 2 to 6 and a maxima (maxT) about 10 to 40. The negative values mean low intensity and positive values mean high intensity. Caution, these values are logarithmic coefficients that stand for very high values (log 1000 = 6.9) (exp 10 # 22000)...(exp 20) # 500.000.000. In theory we should use high "Scale" and gaussian values in the areas closest to the lens (where the veil effect is low) and low values in the distance (where the veil effect is important).

To process the mask resulting from the gaussian process

The "Transmission Map" file correspond to a recursive process of:

- The source image Input image which went through various modifications by the histogram equalizer and the gamma.
- The "mask" files directly resulting from the gaussian process (scale, radius, method low., high...)

"Mask" display type

"Mask" files processing (the idea came to me by studying the Rusell Cottrell's plugin) will allow to decrease flares and artefacts. I added a combo-box to choose the display type. This is both an educational system and also an aid to find the "right" settings.

- Process:
 - Standard: it is the default setting
 - Mask: display the mask obtained by the Retinex algorithm resulted from the gaussian process. We will see here the impacts:
 - of the various settings upstream: method (low, uniform, high and highlight), radius, gamma, histogram equalizer. In the other hand, the other settings downstream have no effect on this display: contrast, gain, brightness, threshold and the curve "Transmission Map".
 - of the various settings of the "Mask equalizer".
 - you can export these settings (TIFF/JPG) to use this mask in external software (Gimp, Photoshop).
 - Unsharp mask: you can use this option to subtract the mask to the input image. In this case, an action on "strength" will allow to equilibrate between the input image and the mask. So we get the possibility to have images with very high radius values.
 - Transmission: display an "image" of the "Transmission Map" file.
 - this "image" doesn't match the reality that matches to a logarithmic scale which values are mostly situated between -30 and +30.

- (fixed) in order to make this values "visible", I applied two arbitrary coefficients I avoided the automatic computing in order to allow comparisons which shift and increase the basic values.
- (auto) I used here the values displayed in the Preview (TM Min Max Mean Sigma) to elaborate a file with values ranged from 0 to 32700. Of course these data become visible with the maximum range but don't match the reality. This configuration can be, depending on the images, preferable to "fixed".
- the displayed image takes into account each Retinex processing, except: a) "median filter"; b) "Gain"; c) "Brightness" (offset); d) and of course "Strength".
- you can act on the different settings, particularly on the "Mask equalizer" ones, and also the "Transmission Map" curve, and then see directly the actions on the haze, the artefacts annot contrasts.

Action methods on the masks

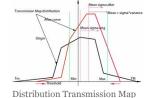
I set 3 methods:

- a diagonal curve which acts directly on the recursive contrast of the "mask" files take care, this curve is very sensitive and can lead to important artefacts. In the other hand, it hasn't, or a few, impact on the processing time. You can use this curve alone or combined with the other methods.
- a gaussian process is applied to the masks. You can act on the shadows, the highlights and the radius. The increase of the processing time is moderate.
- a "wavelet" process (Sharp mask) is applied to the masks. You can act on the shadows and the highlights. The increase of the processing time is important or very important according to the selected method (partial or total).

"Transmission Map" file processing

The main basis processing consist of to apply, using the "Transmission Map" average and standard-deviation, a transformation of the type:

- newT = (oldT mini) / maxi-mini.
- with mini = average k * standard-deviation
- with maxi = average + k * standard-deviation



The choice of k - contrast (variance): modifiable by the user – is decisive in the image rendering: low values will increase the seeming contrast, hight values will make the image more natural with less artefacts and hazes.

I added three components to this basis processing:

■ a threshold - modifiable by the user – that will allow to reduce the maximum (maxT) and minimum (minT) values by integrating the clipped values to the new distribution.

Example: if the chosen threshold is 10, all the values lower than -10 will became -10, all the values over +10 will become +10. This action is going to increase the contrasts, but can lead to artefacts.

- a curve (flatcurve) modifiable by the user to directly act on the distribution, by replacing theory by practical... It is by acting on the distribution that the user will be able for example to reduce the negative values and increase the positive ones (or the opposite) to improve the rendering in real time.
- a median 3x3 modifiable by the user that acts on the distribution irregularities, in order to reduce the artefacts.

Elaborating the "Image Haze-free" file

This file is obtained by "difference" between the input image and the "Transmission Map" file.

We are going to use here 2 extra parameters, Gain and Brightness (Offset): These two parameters allow to include the "Transmission Map" file to an exploitable image. Indeed, without precautions, the restored image could be totally out of the luminance usual limits that are between 0 and 32768.

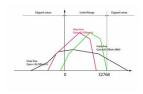


Image Haze Free

- "Gain" will act on the total amplitude by increasing or decreasing it and for example reduce the numbe of totally black pixels.
- "Brightness" (Offset) will allow to reset the luminance after "Gain" acted.

In order to give some help to the user, I added a display for the "haze-free" image minimum values to the GUI. We can act on "Gain" and "Brightness" (Offset), but also on each other parameters, particularly "contrast" (variance) and the "Transmission Map" curve. The ideal is to set Min values the closest possible to 0 and the Max values the closest possible to 32768. For example, an "original" image with min=-44230 and max=76000 could become min=32, max=32500. Of course no other values are forbidden, but in all cases these values will be clipped and brought back to 0-32768, the "Transmission Map" file wealth is lost!!

Keys for acronyms and labels used in the GUI to ease the use:

- Restored haze-free Min=-5681 Max=34568 : display image maximum and minimum values "without veil" before clipping.
- TM Min=-3.8 Max=4.1 Mean=0.1 Sigma=3.3: Tansmission Map -- Min and Max match to the following computing Min=Mean Sigma * Variance / 100; Max=Mean + Sigma * Variance / 100 *TM Tm=-6.5 TM=11.8: Transmission Map -- Tm minimum value of the distribution; TM maximum value of the distribution.

It is obvious that that a display in the GUI:

- of the "Transmission Map" distribution and also
- the "Haze-free" image data

should make easier both the system understanding and ergonomics... but this is still to be fixed ???

Retinex versus Tone-mapping

There is several ways (according to specialized sources in this field) to adapt the Retinex algorithm to a "Tone-mapping" rendering. The study of the existing code, my understanding of the problem... led me to the following solution:

- Proceed to several "Retinex" specific code iterations. I name specific code the part excluding
 - the RGB==>L*a*b* or RGB==>HSL conversions (and vice versa).
 - the histogram equalizer.
 - the differential gamma (low, middle, high, free).
 - in the same way, aren't concerned the choices done with the "Low", "Uniform", "High", "Highlight" method
- These iterations obviously increase the time processing, up to 2.5 times for 5 iterations.

You can choose between several ways to carry on the iterations by acting on 3 factors I called "Gradient":

- "Gaussian gradient" that act on: a) the gaussien blur sigma (Radius or Neighboring pixels); b) the "scale" factor (set by default to 3).
- "Transmission gradient" that act on: a) the contrast (variance); b) the threshold that limits the "Transmission Map" file.
- "Strength gradient" that act on "Strength", that means the combination of the "Haze-free" image with the original one.

If you set these 3 gradients to 0, all the iterations will be identical to the basis processes.

Impact of each gradient:

- "Gaussian gradient": according to the "it" iteration, the Radius value ("Neighboring pixels") is divided by "grad" and "scale" value become:
- Gg 1:
 - it=1 grad=1; it=2 grad=1.25; it=3 grad=1.5; it=4 grad=1.75; it=5 grad=2
 - it=1 scal=4; it=2 scal=3; it=3 scal=3; it=4 scal=3; it=5 scal=2
- Gg 2:
 - it=1 grad=1; it=2 grad=1.50; it=3 grad=2.0; it=4 grad=2.50; it=5 grad=3
 - it=1 scal=5; it=2 scal=4; it=3 scal=3; it=4 scal=2; it=5 sca=2
- Gg 3:
 - it=1 grad=1; it=2 grad=1.66; it=3 grad=2.3; it=4 grad=3.00; it=5 grad=3.6
 - it=1 scal=5; it=2 scal=4; it=3 scal=3; it=4 scal=2; it=5 scal=2
- Gg 4:
 - it=1 grad=1; it=2 grad=1.80; it=3 grad=2.6; it=4 grad=3.40; it=5 grad=4.2
 - it=1 scal=5; it=2 scal=4; it=3 scal=3; it=4 scal=2; it=5 scal=2
- Gg 5 :
 - it=1 grad=1; it=2 grad=3.50;it=3 grad=6.0;it=4 grad=8.5;it=5 grad=11.0
 - it=1 scal=5; it=2 scal=4; it=3 scal=3; it=4 scal=2; it=5 scal=2
- Gg 6:
 - it=1 grad=1; it=2 grad=6.00; it=3 grad=11.0; it=4 grad=16.0; it=5 grad=21.0
 - it=1 scal=5; it=2 scal=4; it=3 scal=3; it=4 scal=2; it=5 scal=2
- Gg-1:
 - it=1 grad=1; it=2 grad=0.875; it=3 grad=0.75; it=4 grad=0.50; it=5 grad=0.37
 - it=1 scal=3; it=2 scal=3; it=3 scal=3; it=4 scal=3; it=5 scal=3

For Gg 5 and Gg 6, the selected method is taken into account, if it is "highlight", 'grad' values are boosted according to "Highlight Threshold"

- "Transmission gradient" : according to the "it" iteration, the contrast (variance) and threshold are multiplicated by "var":
 - Tg 1: it=1 var=1; it=2 var=0.875; it=3 var=0.75; it=4 var=0.625; it=5 var=0.5
 - Tg 2: it=1 var=1; it=2 var=0.800; it=3 var=0.60; it=4 var=0.400; it=5 var=0.2
 - Tg-1: it=1 var=1; it=2 var=1.125; it=3 var=1.25; it=4 var=1.375; it=5 var=1.5
 - Tg-2: it=1 var=1; it=2 var=1.400; it=3 var=1.80; it=4 var=2.200; it=5 var=2.6
- "Strength gradient": according to the "it" iteration, "strength" is multiplicated by "s":
 - Sg 1: it=1 s=1.30; it=2 s=1; it=3 s=0.70; it=4 s=0.5; it=5 s=0.5
 - Sg 2: it=1 s=1.60; it=2 s=1; it=3 s=0.40; it=4 s=0.3; it=5 s=0.3
 - Sg-1: it=1 s=0.80; it=2 s=1; it=3 s=1.20; it=4 s=1.2; it=5 s=1.2
 - Sg-2: it=1 s=0.60; it=2 s=1; it=3 s=1.40; it=4 s=1.5; it=5 s=1.5

Note: "Hue equalizer" and "Transmission median filter" are applied only once, whatever the number of iterations is.

The various adjustments

Adjustments always accessible

There are 5 adjustments at your disposal:

- Retinex method: with "Low" "Uniform" "High" "Highlight"
 - Low: improve low intensity areas
 - Uniform: Strive to process low and high intensities in an equilibrated way.
 - High: Improve the rendering of the more exposed areas
 - Highlight: Improve the rendering of the highlight areas that may become magenta, but can bring strong artefacts.

These 4 methods act directly on the "Transmission Map" file.

- Space
 - L*a*b*
 - HSL Log
 - HSL Lin

The choice between these 3 spaces depends on the image... The 2 firsts "meet" the Retinex algorithms.

- Strength: Combine "Haze-free" image with the original image. Strength=0 brings no Retinex processing to the image. Strength=100 brings a total processing, "Haze-free" image only is displayed. Values under 50 are recommended.
- Chroma (patch): act on the color component, using the Retinex algorithm, in percentage of the "strength" value relative to the luminance. The slider runs a Retinex process dedicated to the color component by "simplifying" the process, luminance specific components - gamma, luminance, gaussian mask, ...- aren't executed. Caution, processing times are doubled and the extra memory calls important.
- Radius (Neigboring pixels): take into account the neighbouring pixels using the "Difference gaussian" algorithm, higher are the values, more affected is the foreground. Lower are the values, more affected are the distant areas, act directly on the "Transmission Map" file.
- Highlight threshold: only accessible with the "Highlight" method; set the "Highlight" method threshold, a value of 1 will give about the same effects that the "High" method, high values should allow to reduce the magenta colors in the overexposed areas. Be careful, it will probably be needed to:
 - change some adjustments, in particular "Radius" (Neigboring pixels)
 - increase the "Raw white point" value in case of magenta hues in the overexposed areas
 - use "Hue equalizer" in case of magenta hues in the overexposed areas
 - act on "Gamma Retinex"
- Contrast (Variance): This coefficient equilibrates "Transmission Map" average and standard-deviation. It is a key factor for the image rendering. It acts on the low and high contrast thresholds of the "Transmission Map" file, the lower are the values the more contrasted is the image, the higher are the values the more natural is the image and the artefacts reduced. This Contrast (Variance) doesn't act on the "Transmission Map" file but allow the Haze-free image computing.

Adjustments accessible with "Setting"

■ Gamma Retinex: with 4 choices "Low", "Middle", "High" and "Free", it allows to choose the gamma and the slope for a gamma applied before the Retinex algorithm and for an inverse gamma applied immediately after rebuilding the image. High gamma values associated with low slope values can lead to artefacts. The using depends on images, La*b* or HSL space, used settings... and the wished effects.

- Histogram Equalizer: This curve is a lesser evil... because it aims alone to restore gamma, RGB conversion, icc input profile, white balance ... which gap is due to the Retinex position in the processing, that is placed in Raw in order to work on the whole image. Without this curve, "Transmission Map" data are generally bad and bring important artefacts. To be noticed that the interface has 2 different curves for L*a*b* and HSL. This curve is without action by default, but in L*a*b* mode it is quite mandatory to reduce the first quarter action (to be adjusted by the user) and to reduce the high values in order to reduce artefacts and haze. It is experience that will guide the user.
- Hue equalizer: This curve will allow:
 - to adjust Retinex action according to the hue, for example to weaken a sky, to boost leaves ...
 - to reduce the action in the overexposed areas of magenta hue
 - it should be noted that when the "highlight" method is selected, this curve also acts on the final image chroma.

These 3 adjustments (Gamma Retinex, Histogram Equalizer, Hue equalizer), don't act directly on the "Transmission Map" file, but modify its original value.

- Gain and Brightness (Offset): used together, they allow to correctly locate the "haze-free" image. You can find help in the information given in the GUI, and place Min close to 0 and Max close to 32768. Gain and offset act on the "haze-free" image and have no action on the "Transmission Map" file.
- Threshold: this coefficient act on the "Transmission Map" maximum and minimum values.
- Transmission Map: this curve is interactive with the "Transmission Map" file, the X axis of the curve represents the "Transmission Map" distribution basis (MinT, mean, MaxT). To drag down the left part will reduce the minimum values (the ones relative to the foreground) and in the same way, to drag down the right part will reduce the maximum values (the ones relative to the background). This curve is a key point for controling the image seeming quality.
- Transmission median filter: reduce the artefacts due to the too strong local variations of "Transmission Map".

I added a "combobox" that allow to select the display type. It is both a pedagogic system and also an help to find the "right" settings.

Process

- standard: the default setting
- mask: display the mask obtained by the Retinex algorithm, resulting from the gaussian processing. We will se here the results of:
 - the different adjustments placed upstream: method (low, uniform, high and highlight), radius, gamma, histogram equalizer. In the other hand, the other adjustments placed downstream have no impact on the display: contrast, gain, brightness, threshold and the "Transmission map" curve.
 - the various "mask equalizer" adjustments.
 - you can export these adjustments (TIFF/JPG) to use this mask in external software (Gimp, Photoshop)
- Unsharp mask: you can use this option in order to subtract the mask to the original image. In this
 case, the action on "Strength" will allow to equilibrate the combination between the original
 image and the mask. Doing so give the possibility to have image with very high radus values.
- Transmission: display an "image" of the "Transmission Map" file
 - this "image" doesn't match the reality, that one corresponding to a logarithmic scale which values are generally ranging between -30 to +30.
 - (fixed): in order to make these values "visible", I applied 2 arbitrary coefficients I avoided automatic computing to allow comparisons that shift and boost the basis values.
 - (auto): I use here the values displayed in the Preview (TM Min Max Mean Sigma) to elaborate a file which values are ranging between 0 and 32700. Of course these data become

- visible with the maximum range but doesn't match the reality. This configuration may, according to the image, be preferable to "(fixed)".
- the displayed image take into account the totality of the Ratinex processings, except a) the "median filter"; b) the "Gain"; c) "Brightness" (offset); d) and of course "Strength".
- you can act on the various adjustments mostly the "Mask Equalizer" ones and also the "Transmission Map" curve and thus directly see actions on the haze, the artefacts and the contrast.

Practical synthesis of the various adjustments

Synthesis of the effects of the main adjustments on the contrast

Adjustments	Main feature	Complementary feature - comments
Strength	Combine the action of the Retinex image with the original one.	No impact on the Retinex processing
Method : Low,Uniforme, High, highlight	Differentiate the action according to the luminance intensity (low = low intensity, highlight = may reduce colored artefacts)	Higher are the values, more impacted is the foreground (high favours the foreground)
Radius	High values for the foreground, low values for the background	High values may create artefacts
Contrast	Low values increase the global contrast, favour the foreground, create artefacts	High values reduce artefacts
Threshold	Low values increase the global contrast, increase transition artefacts	Low values reduce background artefacts
Transmission-map curve	Left part acts on foreground, right part acts on background	Moving the median point acts on contrast and artefacts
Gain and Brightness	Locate Retinex according to the source image (see Settings - restored haze-free)	May give illusion to act on contrast and lightness

Other adjustments synthesis

Adjustments	Main feature	Complementary feature - comments
Gaussian-mask and Sharp- mask	Reduce the artefacts resulting from too high local contrasts	High values reduce Retinex action
Method L*a*b* - HSL	Change the luminance initial repartition, allow to reduce artefacts	L*a*b* seems more adapted for images with high expositions
Gamma and Histogram equalizer	Change the luminance initial repartition to reduce artefacts	Gamma is compensated by its inverse at the end of Retinex processing

Use in combination with wavelets and Lab Adjustments

In case of atmospheric veil, the contrast restoration with Retinex algorithm:

- is sometimes not good enough or leads to artefacts if "Strength" too high,
- leads to a chroma and sharpness insufficiency

Several solutions can be tried, particularly:

- to increase the chroma using L*a*b* adjustments (Chroma slider or C/C curve or * = f(a*) and b* = f(b*) curves;
- to adjust the global contrast using La*b* adjustments (L curve)
- to increase the chroma using the residual image of wavelet.
- it is of course possible to use other RawTherapee controls, particularly in the "Exposure" tab.

But I think that to give priority to the wavelet tool is a good choice, mostly for images with an atmospheric veil. We will be, in addition to the adjustments given above, to:

- increase the residual image contrast,
- use the "contrast" compression method by sliding down "compression strength" and adjusting "gamma",
- lightly increase "Edge sharpness" with "edge detection"
- use "Final touchup", in particular: "Final local contrast", "Final contrast curve" and if necessary ("Tone-Mapping" effect), "Contrast balance method".

A possible example – Dehaze – some recommendations

In order to better explain the processing using Retinex, I will take the example of a difficult image, very hazy and tainted with an atmospheric veil Madeira (http://rawtherapee.com/shared/test_images/madeira.dng)

Be careful, this processing is for pedagogic goal only and has no claim to obtain a perfect image! (by the way, what does it mean, it is subjective). The user has total freedom to proceed to changes, modifications, to get the rendering he wishes.

First Retinex adjustments

- At the start [following not understandable]
- Enable "Retinex" and that's all

To get a synthesis of the different adjustments see # Practical synthesis of the various adjustments

With this kind of image, where the veil is present both in the foreground and in the background and having an uniform histogram, I choose the options:

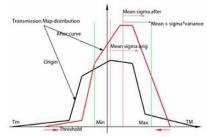
- Method: Uniform
- L*a*b*: because it is the most efficient of the 3 (Lab, HSL Log and HSL Lin)

Then I choose, for "Radius", "Contrast" and "Strength", the adjustments by default, that means respectively 80, 200 and 20.

Transmission map adjustment – Iterations with Radius, contrast, scale

The next step consist to study in the tool "settings" the indicator that give partial information about Retinex optimization, the "Transmission-map" file, I quoted "Restore haze-free". You can also, iif you want, view the "Transmission-map" file, (set Process to the "Transmission fixed" option).

We can see the values: Min=-5799 Max=14069., and also other values: TM min =-9.8, Max=7.7 Mean=8.8 Sigma=4.4 Tm=-13.1 TM=20.7



Distribution Transmission Map

Let's have a look to the both values Min=-5799 and Max=14069. That means that the image corresponding to these values and which will be subtracted to the original image, doesn't correspond at all to the usual luminance values that, as a reminder, are between 0 and 32768 in RawTherapee using $L^*a^*b^*$. The first thing to do is to use the "Transmission-map" curve to get more relevant values. Be careful, the goal isn't to get Min = 0 and Max = 32768, but more reasonable values. For example, it is possible to lower a little bit the left part of the curve and to lift a little bit the central part. We end up for example to the following adjustments: :Media:madeira-first.pp3

The Haze-free image

Just now I don't modify the values that could have an influence on "Transmission-map" like "radius", "contrast".

Then, in an iterative way between "Gain transmission" + "offset" and "radius", "contrast" and "scale", I will try:

- to act on the "Transmission-map" curve and "offset" in order to get Min and Max values close to 0 and 32768, but -400 and +39000 may be a good choice too.
- and then to act on "radius", "contrast", "scale" to obtain a looking fine image, contrasted enough, but without too much effect of depth... it is a matter of taste.
- work by iterations

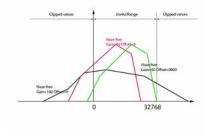


Image Haze Free

I arbitrary took the following values given in the pp3 Media:madeira-second.pp3

To apply a Gaussian mask and a gamma

These two tools allow to modify the distribution of the action between high and low lights.

- choose "Mask Method" = Gaussian mask and increase hihlight = 23 to lighten the sky, you can use the "mask" option to view the changes
- put "Gamma" in "settings" on "Low"

Again, all is matter of taste.

We end up, for example to the following adjustments: Media:madeira-third.pp3

Act on the chroma

Two cases:

- In the version "Retinex at the beginning of the processing" there is no Retinex action on the chroma. You can use the "Lab adjustments" tool and act either on the "Chromacity" slider or on the two curves a*=f(a*) and b*=f(b*), or on the curve C=f© more intuitive.
- In the version "Retinex in Wavelets", act on the Chroma slider in Wavelet.

We end up to the following adjustments: Media:madeira-four.pp3

To boost the contrast and to reduce the atmospheric veil even more – Wavelet

- Slightly increase the contrast with "contrast levels"
- In Residual image" increase contrast and chroma (contrast = 25, chroma = 22), then choose "Compression method" = "Tone mapping" and reduce compression strength = -0.29
- In "Final Touchup", choose "Contrast balance" select "slider" and adjust "Contrast balance d/v-h" to
- Enable "chroma balance"

You end up to the following adjustments: Media:madeira-five.pp3

Final adjustment

You can, if necessary:

- Adjust "Strength" in Retinex
 Act on "Final contrast" and/or "after contrast curve" in wavelet
 And also on all the other usual adjustments.

You end up to the following adjustments: Media:madeira-six.pp3

CIECAMo2

By J.Desmis

About CIECAMo2

Introduction - history

Since many years now, men tries to model colors, its perception by peoples. Lots of work has been done through out the years since the Middle-Age, but it's only starting from the 19th century, then in the 20th one that has been made the main discoveries.

I'm not a specialist of the physiology of the human visual system, neither a researcher in the complex domain of colorimetry. I've I took some minimum information that seem essential to understanding, up to the interested reader to expand it thanks to the Web and the elements I've joined.

Commonly in photography, we use (more than) 50 years old models: RGB and its derivative (HSV, HSL, CMYK,...), XYZ, and Lab and its derivative (Luv, Lch). I won't comeback on the RGB model, known by everyone, it is dependent of the peripheral and doesn't take into account any CAM (color appearance model). The CIE's definition of XYZ (1931) was the first step of the « Commission Internationale de l'Éclairage » (CIE - International Commission on Illumination) towards a human description of the colors faithful to the human vision. To summarize, a color can be characterized by its 3 X, Y and Z values, obtained by a combination of « tristrimulus values », « CIE standard observer » and the base color's « spectral power distribution ». This model has been taken up in RawTherapee, particularly in terms of white balance... This model doesn't take into account any CAM, but it's an extraordinary leap forward, because we now can model a color in cognitive terms.

The Lab model has been designed in 1976 by the CIE by derivating it from the XYZ model, it characterize a color with an intensity parameter corresponding to the luminance and two chrominance parameters that describe the color. It has been specifically studied so that the computed distance between colors correspond to the differences perceived by the human eye. The Lab model is well established in RawTherapee, it is used as a basis for most of the tools: sharpening, denoising, tone mapping, Lab adjustments, etc.. It integrate some characteristics of a CAM, but the benefits are sketchy. The CIECAM02 model, derived from CIECAM97 and using G.Hunt's work, is the first commonly usable model in photography, because it is invertible... and relatively « simple », it can take into account other than purely cognitive aspects and is based on the work of many researchers on the basis of sample of persons who evaluate different parameters, like:

- 1. simultaneous contrast: variation of the colored appearance of an object depending on the colorimetric characteristics of its close environment. For example, the same color will be perceived differently on a white or dark background. The darker the background will be, the more we'll have to boost colors...
- 2. the Hunt's effect: increased seen coloration (colorfulness) with luminance. An object appears more vivid and contrasty in full light than in shade.
- 3. the Stevens' effect: augmentation of the perceived contrast with the luminance. When the luminance increases, the dark colors looks like even darker and the luminous colors appears even more luminous.
- 4. Helmholtz-Kohlrausch's effect: dependence of the brightness in relation to the luminance and chromaticity. Colored objects appear lighter than the achromatic objects with the same luminance. The most saturated colors appear brighter.

5. Chromatic adaptation: adjustment by the human vision system of some color stimuli. The chromatic adaptation let us interpret a color depending on its spatiotemporal environment. It's an essential effect to be taken care of by a CAM.

The chromatic adaptation is the human visual system's ability to adjust to changing illuminant conditions. In other words, we adapt to the color of the light source to better preserve the color of objects. For example, under incandescent light, a white paper appears yellow. However, we have the ability to automatically model the yellowish light so we see as white paper. The world around us would be very complicated if the objects changed color whenever the light source changes even slightly. Since the dawn of time, we must be able to know whether a fruit is ripe, would it be the morning, afternoon or evening. The chromatic adaptation makes it possible. But it can also be the source of many optical illusions. I think the majority of RT users know, at least by name, the previous model of chromatic adaptation, called "Bradford". etc.

Note: there will be no question here of "Munsell Correction" because CIECAMO2 is, by principle, built around Munsell's tables... so this correction is taken into account, even if the model has shortcomings!

My first thoughts about CIECAMo2 dates back to 2007, and the development of a spreadsheet, for best results in the development of ICC "input profiles". Early 2012, I addressed a request from users: "can we have reference colors - color palette - (skin, sky, ...) which would allow a better white balance through a comparison/iteration process". I also worked on the concept of CRI (Color rendering index) which reflects the difference of illuminants compared to a base illuminant... The lower the CRI is, worse the rendering will be with an identical color temperature see: Color Management/fr)

Based on CIECAMo2, the patch contains the necessary basic elements to work these two points, but it lacks an essential element, not easy to develop: a pipette. I have long considered CIECAMo2, not as a gimmick, but as something difficult to implement... and with a quite small bonus compared to Lab. The request of Michael Ezra who surprised me at first, led me to re-open the file; the plug-in for Photoshop was a discovery for me by the example, of CIECAMo2. I am now convinced that even if the model is not perfect (for some pictures, its use is almost impossible!), it is today the most undeniable (effective) in terms of color management. The module I am proposing is an "initiation". From the data of CIECAMo2, it is possible to develop a series of features similar to those already developed in RT (Lab adjustments with various curves, tone-mapping, etc.). Probably with significant advancements in terms of quality.

The lack of effective documentation adds to the complexity... Some points of view are personal (can be tainted by errors?). If a specialist reads these lines, I'll be happy to change my text and my algorithms!

Some definitions

1. Brightness [brilliance] (CIECAMo2):

The amount of perceived light from a stimulus = indicator that a stimulus appears as more or less bright, light.

2. Lightness [luminance] (Lab, CIECAMo2):

The clarity of a stimulus relative to the brightness of a stimulus that appears white under similar viewing conditions.

Note that in RT, the "brightness" term applies to "Lightness"! You will need to make a patch to rename "brightness" to "lightness" in the "exposure", "Lab adjustments", etc... modules.

3. Hue and hue angle (partly in Lab, CIECAMo2):

The degree to which a stimulus can be described as similar to a color described as red, green, blue and yellow.

4. Colorfulness (CIECAMo2):

The perceived amount of color relative to gray = indicator that a stimulus appears to be more or less colored.

5. Chroma (Lab, CIECAMo2):

The "coloration" of a stimulus relative to the brightness of a stimulus that appears white under identical conditions.

6. Saturation (CIECAMo2):

Coloration of a stimulus relative to its own brightness.

To summarize:

- 1. Chroma = (Colorfulness) / (Brightness of White)
- 2. Saturation= (Colorfulness) / (Brightness)
- 3. Lightness= (Brightness) / (Brightness of White)
- 4. Saturation= (Chroma) / (Lightness)
 - = [(Colorfulness) / (Brightness of White)] x [(Brightness of White) / (Brightness)]
 - = (Colorfulness) / (Brightness)

CIECAMo2 develops and uses several types of correlated variables that allow the use of these concepts:

J: lightness or clarity, close to L (Lab)

C: Chroma, close to C (Lab)

h: hue angle, close to H (Lab)

H: hue. A color can be described by the composition of 2 base colors between 4 (red, yellow, green, blue), e.g. 30B70G or 40R60Y.

Q: brightness

M: colorfulness

ac, bc : close to a and b (Lab)

Why the saturation in addition to other close variables? Here is a quote from a text by Robert Hunt (2001):

"Of the three basic color perceptions, hue, brightness and colorfulness, hue has no relative version, but brightness has lightness, and colorfulness has chroma and saturation. Correlates of chroma are widely used in color difference formulae, but saturation currently plays little part in color science and technology. This is perhaps because in many industries, flat samples are viewed in uniform lighting for the evaluation of color differences, and in this case chroma is the appropriate contributor for samples of small angular subtense. For samples of large angular subtense, however, a correlate of saturation may be more appropriate to use. In the real world, it is common for solid objects to be seen in directional lighting; in these circumstances, saturation is a more useful percept than chroma because saturation remains constant in

shadows. In imaging, artists and computer-graphics operators make extensive use of series of colors of constant saturation. In optical imaging, saturation can be an important percept in large dark areas. Recent experimental work has provided a much improved correlate of saturation."

The 3 processes

Three processes allow the use of CIECAM, their names depends on each designer. I've made a synthesis (reminder: this document is not a course, or a thesis on CIECAM... but an aid to its understanding and use).

Process 1

Names like "origin", "forward", "input", "source" are generally used... I finally chose "source", which corresponds to shooting conditions and how to bring back the conditions and data to a "normal" area. By "normal", i mean medium or standard conditions and data, i.e. without taking into account CIECAM corrections, e.g. "surround = average", D50 white balance!

Process 2

It corresponds to the treatment of correlated variables (J C h H Q M s a b) for various purposes: action on lightness (J), brightness (Q), chroma (C), saturation (s), color level (colorfullness M), the hue angle h, as well as ac et bc. It is quite possible to build an images editing software around those variables...

In the case of this patch for RT, I arbitrarily selected 4 groups of algorithms:

- 1. JC by adding an contrast function;
- 2. Js, as above
- 3. QM
- 4. All: all parameters including h.

These modules are simple, more for pedagogical purpose than trying to solve the problems of colorimetry, even if the results are in my opinion excellent.

I completed this process by:

- 1. double "curves", acting on contrasts J (lightness) or Q (brightness), whose principle is similar to the double curves of "exposure";
- 2. a choice for color curves between chroma, saturation and color level (colorfulness).

We could add other algorithms based on the Fourier transform, or replacing equivalent functions of RT...

Process 3

Names like "inverse", "reverse", "output" or "viewing conditions" are generally used. I've chose "viewing conditions", which reflects the media on which the final image will be viewed (monitor, TV, projector, ...), as well as its environment. This process will take the data from the process 2 and "bring them" to the support so that the viewing conditions and environment are taken into account.

Note: we find here the explanation of the rendering difference between a printed photo and a picture viewed on a monitor - even if the printer is a high-end and well calibrated one: the viewing (observation) conditions! A printed photo will often viewed in an album, often on a black background, in low light... and

often tungsten lighting. The original will be seen on a monitor with a light background, and a D50 illuminant... There is no question of changing the "print" output, but to adapt the "monitor, TV..." output.

That is to say, but here stops the comparison, that we realize something like soft-proofing, but it's not the case because it's the purpose of CIECAM. We takes into account the settings specified in "Preferences" (white point of the output device [screen TV, projector...] and its average luminance [% gray]. We also takes into account the luminance of the room in which the observation is made, as well as the relative luminance of the visualization device's environment (more or less black).

Simplified synthesis of what RT allows with the current patch:

- 1. general case of the user who uses RT to see his development... that should represent 95% of the cases. In this case, "viewing conditions" corresponds to the RT work environment, e.g.:
 - monitor's white point set to 6000K
 - calibrated monitor, so Yb=18
 - but according to:
 - the selected "theme" in "Preferences / General" (almost black or gray), you have to change "surround"
 - the monitor's location (on a neutral or dark background), you have to change "surround"
 - lighting of the room, that will change with time, you'll have to change "adaptation luminosity viewing La": e.g. at night without lighting up, "La" will be close to 0 or 1, and on the contrary by day in a bright room, "La" will be close to 1000
- 2. less common case, but possible, because I've already done it, I use RT and the family TV to show pictures as well as RT's possibilities. The "viewing conditions" will be different and to be adapted to each case; you'll have to review each of the points above with possibly different settings: TV's white point, TV's Yb (empirical?), a different "surround" because we generally look at the TV with a dull background, and with a reduced room's illumination.
- 3. you want to prepare a series of photographs for an exhibition: in this case, in a professional manner, we will "see" viewing conditions on site and ask for data like the projector's white point, calibration (?), the room's brightness on the day of the exposition, etc... In RT, the user will set "viewing conditions" to suit the exposition's conditions, and produce X corresponding jpeg (or TIFF)
- 4. etc.
- 5. that's why I put most of the settings for the process # 3 (Output Device) in "Preference" is not an error, but appears similar setting as the monitor profile that depends on the monitor...

Data

Which data are taken into account and which simplifications I (arbitrarily?) made? How to adjust them?:

- **Yb**: Yb is the relative luminance of the background! With that, we're much advanced! Specifically it is expressed in % of gray. A 18% gray corresponds to a background luminance expressed in CIE L of 50%.
 - for process #3, if your monitor is calibrated, you can easily have a value of Yb close to 18 or 20. If your TV or projector, that seems difficult to calibrate, seems dark or light, you can adjust this value empirically. It depends on the visualization support and can be considered as constant for a set of photos and in a condition of observation. If you want to change this value, go to "Preferences / Color Management / Yb luminance output device (%)"
 - for process #1, it's much more complex because:
 - an image has rarely a constant exposure and small luminance variations
 - I placed the CIECAM module at the end of the Lab process, just before the RGB conversion and the sends to the output device, so we can assume that the user has used various tools of RT to make the image have an "average" histogram

So I arbitrarily made Yb inaccessible by calculating it from the average luminance of the image. Of course, if in the future RT integrates pipettes to separate the image areas (dark, normal, bright...) then it would be possible to enter several Yb values. For example on an image we could see three areas:

- standard, which corresponds to the average luminance of the image with a Yb set to 20%;
- underexposed (approximate contours delineated by the pipette...), where the luminance would be calculated and would e.g. give as a result Yb=5%;
- overexposed, where Yb would be as high as 70%...
- La: La is the adaptation field's absolute luminance! Again, we're much advanced now!
 - In process #1, it corresponds to the luminance when shooting. E.g. if you make a photo in the shade, "La" will be close to 2000cd/m2; if you make interior shots, "La" will vary depending on the lighting from 20 to 300cd/m2... In reproduction, these values may be even lower
 - "Scene luminosity" and the "Auto" checkbox (process #1):
 - If Checkbox enabled, La is calculated with Exif data (shutter speed, ISO speed, F number, Camera Exposure comprensation) and also Raw White Point and Exposure compensation slider
 - In process #3, it corresponds to the luminance of the place in which is made the observation. When you calibrate your monitor, you are asked for this value... or you are offered the choice of using a probe. Orders of magnitude from 15 to 100 should resolve most cases. But e.g. for a theatrical projection in the dark, it can lead to lower values (1-10)
 - These 2 values of "La" are adjustable in RT, in the "CIE Color Appearance Model 2002" tool

Surround

Again, I have made simplifications...

- for process #1, this data reflects some shooting conditions, such as photos in a museum with a dark background, or portrait shooting on a black background. Usually RT's user will have corrected the deviations from its perception thanks to its numerous tools. However, I've added a checkbox "Surround (scene) dark", which can be activated if necessary. Its use will lighten the image (recall: this process bring the data "back to normal")
- this data reflects the surroundings of the image when viewing. The darker the surrounding will be, the more you'll have to increase the contrast of the image. The "surround" variable is not acting as a D-lighting or tone curve, it also changes the colors in the red-green and blue-yellow axis. If the environment's luminance is greater than 20%, choose "average", otherwise adapt to your conditions, e.g. RT's settings (Preferences / General / Select theme) will affect the final rendering. This setting is accessible by "Surround (viewing)". The darker the surrounding will be, the higher will be the image's simultaneous contrast.

White-points model

- "WB RT + output": here we trust RT's white balance for the process #1; CIECAM uses D50 as a reference: RT's white balance bring back the conditions to a D50 equivalent, while for process #3, it will be necessary as needed to set the white point of the output device. Go to "Preferences / Color Management / Settings white output device (monitor, TV, projector)" and select an illuminant in the list (is it sufficient? I have no idea about the characteristics of projectors, light, temperature...)
- WB RT+CATO2 + output": for process #3, we are in the same situation than above; for process #1, a mix is made between RT's white balance and CATO2 that is using its settings, which makes a solution where the two effects (RT and CATO2) are combined. You can modulate the action of CATO2 by acting on the "CATO2 adaptation" slider. You'll probably have to change RT's white balance settings to benefit from the "mix" advantages, otherwise the effects does add themselves.

■ CATo2 is a chromatic adaptation, it converts the XYZ values of an image whose white points is Xwo, Ywo, Zwo, in new XYZ values whose white point becomes XW1, Yw1, ZW1; the algorithm used is similar to the one from Von Kries, therefore different from RT's correction that takes into account the channels multipliers!

"CATo2 adaptation" and the "Auto" checkbox

- see above for the usage of "WB CATo2 + output"
- however, even when "white point model" is set to "equal", this slider can be useful. Usually, the "auto" checkbox must be checked and CIECAM calculates itself an internal "D" coefficient that is used for other purpose than the chromatic adaptation. The result is a value greater than 0.65 (65%); you can uncheck the box that will alter process #1, the effects can be unexpected...

Algorithm

You can choose between JC, JS, QM (of course there are other possible combinations!), or "All" which lists all the possible settings (I arbitrarily excluded "h" as well as "ac" and "bc" from the 3 algorithms JC, JS, QM). The most common use (if one can use that term for CIECAM) is JC, then act on the sliders to get the desired rendering... which I recall depends on the display device, its environment, its settings and the brightness of the room.

■ "JC" algorithm

- J simulates the lightness close to L (Lab) and C simulates the chroma, near the c (Lch) chromaticity. But, important difference, J and C take into account the "effects" (simultaneous contrast, Hunt, Stevens, etc...) which is not the case of Lab and even less RGB.
- J varies in the [0-100] range and corresponds to a relative value of the brightness (likewise L, or Value...) and theoretically C in the [0-180] range (it can be higher)
- The two cursors using J and C may vary from -100 to +100 with actions similar to the "Brightness" (to be renamed "Lightness") and "Chromaticity" of "Lab adjustments" sliders.
- with the "JC" algorithm, a skin tones control is possible, the action is similar to the similar cursor from "Lab adjustments"
- The "contrast" cursor modulates the action of "J" with an "S" curve, which takes into account the histogram 's average brightness "J".

■ "Js" algorithm

It is similar to JC, but:

- the chroma is replaced by the saturation (CIECAM). But for which purpose? I'm quoting again an excerpt of the text from G.Hunt For samples of large angular subtense, however, a correlate of saturation may be more appropriate to use. In the real world, it is common for solid objects to be seen in directional lighting; in these circumstances saturation is a more useful percept than chroma because saturation remains constant in shadows. In imaging, artists and computer-graphics operators make extensive use of series of colors of constant saturation. In optical imaging, saturation can be an important percept in large dark areas. Recent experimental work has provided a much improved correlate of saturation.
 - Skin tones control is less "fine" than with "JC", globally wider in the reds

• "OM" algorithm

- here we use 2 variables Q ("brightness") and M ("Colorfulness") that are not relative data, but absolute. We takes into account the white's brightness. It is easy to realize that the same white "J=100" will appear brighter in the sun than in a dark room...
- the white's brightness takes into account the following parameters (scene): "adaptation luminosity La", "CATO2 adaptation" and "Yb" (currently not adjustable)

- in common use, the control is more difficult than with "JC", however it provides opportunities for high contrast images and opens the door for HDR processing
- Skin tones control is less "fine" than with "JC", globally wider in the reds
- The "contrast" obviously acts different, since it takes into account Q differently than J.

"All" algorithm

you can control all the CIECAM variables: J, Q, C chroma, saturation s, M color level, J contrast,
 Q contrast, hue angle h, skin tones protection (acts on C only)

Tone curves and color

Curves

- you have likewise in the "exposure" module a set of 2 tone curves, which acts on the "J" lightness and "Q" brightness. You can use one curve only or both, by eventually mixing "lightness" and "brightness". Beware, "brightness" curves can easily lead to out of bounds results! "Brightness" is an absolute scale, while "Lightness" is a relative scale, the same "J" white will appear whiter directly illuminated by the sun than in the shade, which is taken into account by "brightness" (Q). Thus shadows and highlights will be rendered differently by the "lightness" and "brightness" curves.
- you also have a set of "chroma" curve with 3 choices: chroma (the most common), saturation and colorfulness. These 3 curves are used to adjust the chosen parameter according to itself, e.g. modulate the saturation to avoid that the already saturated colors goes out of gamut. For these 3 curves, the "red and skin tones protection" cursor is operational, it is more suitable for skin tones in the "chroma" mode. I recommend using the "parametric" mode that allow to differentiates according to the colors's saturation level. Note: All "chroma curves" combinations (chroma, saturation, colorfulness) and sliders (chroma, saturation, colorfulness) are not possible without overly complicating the code, that's why in few cases some sliders are grayed out.

Histograms in Tone Curves

The tone curve histograms in the *CIE Color Appearance Model 2002* tool can show values before or after CIECAMo2 is applied. To view the values after CIECAMo2 adjustments, enable the "*Show CIECAMo2 output histograms in curves*" option. If disabled, the histograms show values before CIECAMo2.

Histograms in the Color Curve

The histogram in the Color curve shows the distribution of chroma (saturation/colorfulness) according to the intensity of the chroma (saturation/colorfulness) or chromaticity in Lab mode. The more the histogram is shifted to the right, the more saturated colors are close to the limits of the gamut. The more the histogram is shifted to the left, the more the colors are dull.

The abscissa represents the value of the chroma (saturation/colorfulness) or chromaticity (in Lab mode). The abscissa scale is "open".

As usual, the ordinate represents the number of pixels involved.

Gamut control (Lab + CIECAM)

- this checkbox will constrain the data into the workspace. I could perform this action in CIECAM mode (process #3), but it would have considerably slowed down the system
- The used algorithm is the same than in "Lab adjustements", it works in relative colorimetry. I think that differences with what could be produced in CIECAM mode are minimal.
- Some adjustments of the CIECAM code are made when "Gamut control" is enabled..

Code, calculation accuracy and processing time

The code for processes #1 and #3 is strictly the one of CIECAMo2 (M.Fairchild, Billy Biggs, ...) that I've adapted and optimized to RT, as well as improvements to the gamut correction by Changjun Li, Esther Perales, M Ronnier Luo and Francisco Martínez-Verdú.

The processes #1 and #3 are symmetrical and stacks many floating point calculations. The use of "double" is mandatory, hence important processing time of about 1 second per Mpix. After extensive testing, we have insights that using float instead of double racing acceleration allowed for processing without sacrificing image quality You can change this setting in "Preferences / Color Management"

In terms of accuracy, I wanted to make some checks by comparing a series of data before and after CIECAM; the differences are very small, e.g. an XYZ input value of 6432.456 has an output value of 6432.388, which is correct.

Limitations of CIECAM02

This model is not perfect, and the following limitations are identified. They can lead for certain images to process them correctly:

- we have already seen this for the Yb settings;
- CIECAMo2 is not a workspace as sRGB or Prophoto, or even Lab. So it is difficult to control the gamut. CIECAM is even known for its problems of narrow gamut, that's why unexpected results may occur to the limits if you're pushing up too much the sliders (J, C, s ...); this may lead in critical situations (highlights, ...) to black or white spots in these areas. Do not hesitate to use RT's tools (highlight recovery, highlight reconstruction, impulse noise reduction, ...), or burned or black areas (raw white and black points, avoid color shift, ...)
- large workspaces (widegamut, Prophoto, ...) can lead, in some cases, to black areas while they won't appear in sRGB (narrowness of CIECAM's gamut).
- The noisy images will influence CIECAM, which will think that the colored dots are realities; that's why I placed CIECAM after "denoise"
- the CIECAM model favors "cones" and takes sparsely into account the "sticks", which means that peripheral vision is sparsely taken into account.
- So, with CIECAM, do not expect to find a cure for "difficult" pictures (overexposure, sensor's saturation, etc...). But for "normal" images (which is the majority), advances that seem more than significant.
- Etc...

Maybe we'll see one day CIECAMxx appearing that could overcome the lacks of CIECAMo2?

The 12 principles of CAM by R.Hunt

- The model should be as comprehensive as possible, so that it can be used in a variety of applications; but at this stage, only static states of adaptation should be included, because of the great complexity of dynamic effects.
- 2. The model should cover a wide range of stimulus intensities, from very dark object colors to very bright self-luminous color. This means that the dynamic response function must have a maximum, and cannot be a simple logarithmic or power function.
- 3. The model should cover a wide range of adapting intensities, from very low scotopic levels, such as occur in starlight, to very high photopic levels, such as occur in sunlight. This means that rod vision should be included in the model; but because many applications will be such that rod vision is negligible, the model should be usable in a mode that does not include rod vision.

- 4. The model should cover a wide range of viewing conditions, including backgrounds of different luminance factors, and dark, dim, and average surrounds. It is necessary to cover the different surrounds because of their widespread use in projected and self-luminous displays.
- 5. For ease of use, the spectral sensitivities of the cones should be a linear transformation of the CIE x , y , z or x 10 , y 10 , z 10 functions, and the V'() function should be used for the spectral sensitivity of the rods. Because scotopic photometric data is often unknown, methods of providing approximate scotopic values should be provided.
- 6. The model should be able to provide for any degree of adaptation between complete and none, for cognitive factors, and for the Helson- Judd effect, as options.
- 7. The model should give predictions of hue (both as hue-angle, and as hue-quadrature), brightness, lightness, saturation, chroma, and colorfulness.
- 8. The model should be capable of being operated in a reverse mode.
- 9. The model should be no more complicated than is necessary to meet the above requirements.
- 10. Any simplified version of the model, intended for particular applications, should give the same predictions as the complete model for some specified set of conditions.
- 11. The model should give predictions of color appearance that are not appreciably worse than those given by the model that is best in each application.
- 12. A version of the model should be available for application to unrelated colors (those seen in dark surrounds in isolation from other colors).

Some links

CIECAMo2 Wikipedia [4] (http://en.wikipedia.org/wiki/CIECAMo2)

Color Appearance Model - Fairchild [5] (http://www.cis.rit.edu/fairchild/PDFs/AppearanceLec.pdf)

Mémoire Laborie ENS Louis Lumière [6] (http://www.ens-louis-lumiere.fr/fileadmin/recherche/Laborie-photo-2007-me m.pdf)

Wavelets

The effects of this tool are visible at any zoom level. However, due to the nature of the algorithm, only the 1:1 (or more) preview zoom will match the saved image perfectly. The size of the image has a direct impact on the perceived sharpness. In practice, this means that you should use this tool at a zoom level similar to that of the desired final size of the image in terms of viewing distance, screen or paper size, etc. So if you intend to print a high resolution image on a 90x60cm canvas and view it from a distance of 30cm, then it is reasonable to zoom in to 100% and adjust the "O (Finest)" contrast slider. However, in reality such large prints when hung on the wall are observed from a good distance off - say, from a couch some meters away - and at this distance adjustments to the finest detail level cannot be distinguished; the eyes cannot perceive such fine details at this distance. The same applies if your image will subsequently be resized, e.g. for email attachments to friends or customers, for image sharing websites, for social platforms, etc. Not only will the image resolutions be lowered but the images may be viewed on devices with a lower resolution, not even fullscreen, on laptops, tablets or smartphones. Keep this in mind, so you do not waste time adjusting detail levels which will be lost anyhow. Usually only Contrast sliders 2 and above provide a useful effect.

The wavelet tool is close to being finished (today being 2015-02-11). The information contained herein should be mostly accurate for the wavelet tool as available in version 4.2.217. Some information still needs to be written or updated.

What is a Wavelet?

A wavelet (http://en.wikipedia.org/wiki/Wavelet), or more specifically a wavelet transform (http://en.wikipedia.org/wiki/Wavelet_transform), is a mathematical function very useful in image processing as it enables one to split the signal into various levels of detail so that you may work on the levels which are of interest to you. The English term "wavelet" was introduced in the early 1980s by French physicists Jean Morlet and Alex Grossman. They used the French word "ondelette", which means "small wave". Afterwards, this word was transferred to English by translating "onde" into "wave" giving "wavelet". The wavelet transform is fairly similar to a Fourier transform, the general difference being that wavelets are localized in both time and frequency whereas the standard Fourier transform is only localized in frequency.

Decomposition in RawTherapee currently uses the L*a*b* space. This decomposition produces two kinds of fundamentally different products:

- 1. Several levels of detail. Choosing the number of levels depends on your needs. The first level corresponds to an area of 2x2 pixels, the tenth level corresponds to a region of 1024x1024 pixels. The more levels, the more processing time and memory is required. Each level contains only the "variations" (positive or negative, for example the contrast that gives data type 0+D1, 0-D2). If a picture is absolutely uniform in luminance and color, each level will contain no information. Thus, one can act separately on each of the levels, but with specific algorithms that take this feature into account. The "quality" of information present in each level is a function of the mother wavelet used: either it is simple as is the case with the "Mexican hat wavelet (http://en.wikipedia.org/wiki/Mexican_hat_wavelet)" as used in dcraw (https://en.wikipedia.org/wiki/Dcraw) and GIMP (http://en.wikipedia.org/wiki/GIMP), or it is more complex and powerful as is the case with the "Daubechies wavelet (http://en.wikipedia.org/wiki/Daubechies_wavelet)" as used in RawTherapee.
- 2. A residual image. This image corresponds to the difference between the source image and the combined details of all level scales. The residual image has the same characteristics in terms of its treatment as the source image, and lacks the details which are present in the other levels. Modifying

the characteristics (contrast, chroma, etc.) of a level has no effect on the residual image, and vice versa.

Depending on our needs, we will need to look either at an individual level of detail, at several levels of detail, at the residual image, or at all of these at once. RawTherapee makes this so easy it's ridiculous.

After decomposition it is possible to use wavelets for different purposes:

- Image compression,
- Noise processing,
- Secret watermarking (http://www.intechopen.com/books/discrete-wavelet-transforms-algorithms-a nd-applications/application-of-discrete-wavelet-transform-in-watermarking),
- Specific treatment for astronomy, separating the background (residual image) from the planets and stars, or separating stars from dust clouds.
- etc.

Wavelet decomposition is used to separate each luminance level (contrast), the chromaticity of each of the a* and b* channels, and the residual image. This action enables you to apply different treatments to the luminance detail levels, the chrominance detail levels, and the residual image which depends to the selected number of luminance detail levels. It is possible (for software developers) to extend the existing processes to others. For example, each level corresponds in fact to the equivalent of a layer; by having suitable GUI tools, it is possible to edit the details (stains, spots, etc.) only in the levels they appear in, and therefore preserve better image quality. Similarly, it is theoretically possible to apply treatments available in other RawTherapee tools to the residual image (L*a*b* Adjustments, Color Toning, etc.). Their implementation may require increased memory requirements and processing time.

You can change the type of mother wavelet used by going to Preferences > Performance & Quality:

- Check the option "Use Wavelet Daubechies D6 instead of D4" to increase the number of coefficients of the wavelet.
- This action has no or little effect on the residual image or high levels (5..10)
- ((can't make sense of this)) On the other hand it (it what? it=D4 or it=D6?) can reduce unnecessary stress effects ((what is this?)) for low detail levels (0..1) under undesirable areas that D6 is more discriminating than D4 (skies, high values of initial contrast, etc.), for example when using high values of "Edge Sharpness".
- There is virtually no difference in processing time or memory requirements between D4 and D6.

Once the "treatments" have been applied (to the levels and residual image), the image is re-composed.

RawTherapee has a tool called Contrast by Detail Levels - it "looks" like this Wavelet Levels tool but has several differences:

- CbDL has fewer levels,
- ((did I get this right?))CbDL lets you only work on the levels' luminance, while Wavelet Levels lets you also modify the chromaticity,
- CbDL has no residual layer.

Having said that, there is nothing preventing you from using both Contrast by Detail Levels and Wavelet Levels together. CbDL is located upstream in the toolchain pipeline, ((what does this mean?)) and can interfere on levels o-5.

Utilities

Total Levels

This slider allows you to select the number of detail levels the image should be decomposed to. The greater the number, the more processing time and memory will be required.

Choosing a low number of levels, 5 for example, will limit the system to details of a scale of 32 pixels. The residual image will contain all the detail the image has except for that contained in levels 0-5.

Choosing a high number of levels, 9 for example (the 10th level is called "Extra"), will allow you to change the details of a scale of 512 pixels (level 8) or 1024 pixels (level "extra", whose effects are only visible if your preview size is large enough). The residual image will differ significantly from the original image as levels 0-10 contain all of the details, so what is left in the residual image is a blurry mask.

The residual image is what is left after all the detail levels have been removed, as many as you have chosen to use via the Max Levels slider. For example if you decompose the image into 7 detail levels (set Max Levels to 7), but you only use the contrast slider on the first three levels (on levels 0-2), the residual image is still whatever is left after the details of levels 0-6 have been removed from it, not just levels 0-2.

Important! RawTherapee uses as many levels as it can unless the size of the level is larger than the size of the image. That is to say, if you are working on a photo larger than 1024x1024px, then changes made even to the 1024px level will be visible in the saved image. If the photo is smaller than that, then the 1024px level will be skipped. This is important to keep in mind, as the same principle applies to the preview: if the photo is larger than 1024x1024px but the preview area is smaller than 1024x1024px, then the 1024px level will be skipped in the preview, though it will be used when saving the full image. This may lead to significant differences between the preview and saved image, so it is worth committing this caveat to memory. You may expand the size of the preview area by shrinking or hiding the left and right panels (shortcut "m") and you gain a few more pixels by making the RawTherapee window fullscreen (shortcut "F11"). The interface informs you how many levels are used in the preview under the contrast sliders, "Preview maximum possible levels=9" for instance.

Full Image vs Tiles

A selection box allows you to choose between:

- Full image,
- Big tiles,
- Little tiles.

It is always preferable to use "full image" as it avoids quality issues in the transition zones between tiles. If you do not have enough RAM or/because you are processing huge images then tiles may need to be used.

Memory requirements

Memory requirements, in bytes, for "full image" mode, assuming 9 levels are used, with examples given for two cameras:

- 1. Pentax K10D (10.2MP) 3888 x 2608
- 2. Nikon D810 (36.3MP) 7360 × 4912

Baseline

Minimal requirement for opening an image in RawTherapee, when all tools are off/neutral width*height*12

K10D: 116MiB

D810: 414MiB

When using level contrast, chromaticity or hue targetting/protection

levels*width*height*3 + width*height*7

K10D: 329MiB D810: 1172MiB

Additionally when using Avoid Color Shift

width*height*4 K10D: 39MiB D810: 138MiB

Total

K10D: 483MiB D810: 1724MiB

Show wavelet levels

This panel has three combobox menus which allow you to refine what you see in the main preview. It serves to teach you, to help you understand how to work with the wavelet system. Changes made to this panel affect only the main preview, not the processed image.

The first combobox lets you choose to preview:

- 1. One wavelet level,
- 2. Below or equal to the level,
- 3. Above the level + residual,
- 4. All levels & directions.

The second combobox lets you specify the reference level for the above options, from "level o" to "level 8".

The third combobox lets you choose to see the following wavelet directions for the first three preview options:

- 1. Vertical,
- 2. Horizontal,
- 3. Diagonal,
- 4. All.

For example you can choose 'one wavelet level' mode to find the wavelet which contains the scale of details you want to work on (e.g. the level containing the scale of skin blemishes but not skin texture) and to see the effects of your adjustments on that level.

Another example, you could choose to preview "above the level + residual" with "level 8" selected to see the residual image and the influence upon it of the tweaks you make in the "Residual image" panel, explained further on.

Contrast

Contrast Levels

By default, seven levels appear in this menu (starting from level o). You can reduce or increase this number by changing the Total Levels slider.

Reminder that the fewer levels you use, the less RAM and processing time is required, and the more the residual image matches the source image because fewer detail levels are removed from it.

The "Contrast -" and "Contrast +" buttons provide an easy way of tapering off the level values in a gentle slop. Of course, these buttons aid in adjustments where small levels are most significant, but often one may want to adjust the high levels too, for instance so that the small levels increase contrast while the large levels decrease it.

Each level corresponds to a scale of details. Action on a given level is only possible if that level contains contrasting areas to begin with. If a level is uniform in contrast, the slider will have no effect.

Levels:

- 1. Level 0: 2x2 pixels
- 2. Level 1: 4x4 pixels
- 3. Level 2: 8x8 pixels
- 4. Level 3: 16x16 pixels
- 5. Level 4: 32x32 pixels
- 6. Level 5: 64x64 pixels
- 7. Level 6: 128x128 pixels
- 8. Level 7: 256x256 pixels
- 9. Level 8: 512x512 pixels
- 10. Extra: 1024x1024 pixels

Remember what you read in the "Total levels" section above: the preview area must be at least as big as the level in order for it to be visible in the preview.

The residual image is not listed above because it is not a level - it is whatever is left over from the original image after the detail levels are taken away from it.

Changes to a level can be applied in a straightforward manner, or only to highlights and shadows. The same is true for changes to chromaticity. Read on.

Apply to

This module allows you to focus the sliders' action on a luminance range, for example increasing the contrast of fine details for high luminance, and reducing the contrast of coarse details for low luminance.

All Luminance

If you select "All Luminance", each slider has the same action regardless of the luminance of the area.

Shadows/Highlights

If you select "Shadows/Highlights" more controls appear in order to customize the action:

- A luminance range for shadows,
- A slider which allows you to select how many levels should only have a darkening effect, always starting from "level 8", even when you use less levels than that.
- A luminance range for highlights,
- A slider which allows you to select how many levels should only have a lightening effect, always starting from "level o".

For example if you only use 7 levels, which makes the highest level be "level 6" (the count starts from 0), and you want "level 6" to only have a darkening effect and levels 0 and 1 to have a lightening effect, then set the Shadow Levels slider to 3 and the Highlight Levels slider to 2.

Changes to levels which are not targetted by these two adjusters apply evenly to their whole luminance range.

Chromaticity

The construction of this module is similar to that of the contrast module. You have the following choices:

- 1. All chroma: in this case, the whole chromaticity range is affected by each level's change, independent of contrast levels.
- 2. Pastel/Saturated: here you can choose two limiting ranges for pastels and saturated colors, independent of contrast levels.
- 3. Link contrast levels: here chromaticity changes are directly related to those of contrast.

All Chroma

If you select "All Chroma", the whole chromaticity range of each affected wavelet level is modified, regardless of how saturated it is.

Chromaticity = sqrt(a*a + b*b) is a wide open, but I limited it to 0..140, which should satisfy almost all cases.

The chromaticity curve:

- The *x*-axis represents the wavelet levels, from the finest on the left end to the coarsest on the right end. Its effect is therefore not continuous, but discrete: the vertical values are sampled in 9 points equally spaced out along the *x*-axis, starting at the very left edge of the curve, and ending at the very right edge. If you use less than 9 wavelet levels, say 7 points, the rightmost 2 sampled points would be ignored.
- The *y*-axis represents the force of the action for each level. If the curve is above the equator, chromaticity for the affected wavelet levels is increased. If below decreased.

The same remark applies here as that for the contrast levels: action on a given level is only possible if that level contains varying chromaticity to begin with. If a level's chromaticity is uniform (does not vary), the curve will have no effect for that level.

Pastel/Saturated

You can treat the chromaticity in three ways:

- 1. All chroma Evenly apply the change in chromaticity (*y*-axis) regardless of what chromaticity the levels already have.
- 2. Pastel/Saturated Apply the change in chromaticity with respect to the pastel/saturated threshold value, default at 5, so the limiting range of saturated tones will be effective for the first 5 levels and the limited range of pastel tones for the remaining levels.
- 3. Link contrast levels Here the action on the chromaticity insofar as it varies in the original image will be directly proportional to the slider settings for contrast levels. This proportionality is achieved using the "chroma link" slider: o has no effect on the chromaticity, 100 has maximum effect.

Split Toning

This "partial color shift" (toning) is perhaps not the best name, but it must do...

We can not act directly on the color (hue) in the levels because the a* and b* channels are decomposed and it is near impossible to chart an accurate mathematical relationship between the chosen hue and such decomposed channels.

Nevertheless, we propose two curves built on the same principle as the chromaticity curve, where the x-axis represents the (discrete) number of levels (o-8), and the y-axis the force of the action with a positive effect above the equator and a negative one below it:

- The a* curve, which acts on the red-green tones,
- The b* curve, which act on the blue-yellow tones.

Note: if you apply identical adjustments to both a* ad b* curves, the end result will be identical to that which you would get had you used the chromaticity curve.

Edge Sharpness

This control is an addition to RawTherapee's existing arsenal of sharpening tools, including Unsharp Mask, RL Deconvolution, Edges, etc. Its principle is naturally based on the use of the wavelet decomposition. From a distance it vaguely resembles the Unsharp Mask to the extent that the wavelet decomposition creates a residual image which can be likened to a mask. But there the comparison ends. The wavelet decomposition of levels using the "Daubechies" method is a kind of direct recognition of contours as it scans each level of detail. The first level, level 0, uses every second pixel to act on details of a "radius" of about 2 pixels; the second details of a "radius" of 4 pixels, etc.

If the algorithm is interested only in the smallest details ((??? we can select the radius in the UI, so this "smallest details" doesnt really make sense)), and if a value is placed in reference((???)) (I called "Radius") and compare it to the current level, we are able to adjust the increase in contrast only for these levels, with a significantly higher increase values to those obtainable with the 10 sliders "Contrast". ((this whole paragraph needs rewriting, I can't make interesting/useful sense of it))

There are 3 sliders:((this section is a direct copy, i did not rewrite it because its meaning is unclear to me))

- Radius may need to call it otherwise?((why?)) Its size will allow to increase the effect of perspective((perspective?)). The increase is not linear. Its action is not zero if the cursor is at zero. Discriminant action is maximum (for low wavelet levels), for low radius values, without being able to predict the level of action.((???))
- 2. Value controls the strength((then why dont we call it strength?)). Setting it to 0 effectively turns sharpening off.
- 3. Threshold (Threshold) allows different breakdown levels higher (distribution) for low levels (0, 1..2) when moving the cursor to the right.(???))

The edge sharpness tool is very sensitive to Maximum Levels setting because in principle it is interested only in the smallest levels, which are always decomposed.

Gamut - Controls

Reduce Artifacts in Blue Sky

Digital images often have some mottle noise in the blue sky. The wavelet process, because it increases local contrast, can generate or amplify small artifacts. This checkbox introduces a median filter to reduce these artifacts.

Do not forget to use the Noise Reduction tool which is called upstream in the toolchain pipeline.

Hue Targetting/Protection

This control is linked to the Contrast and Chromaticity wavelet controls. Toning is not affected. Moving the slider to the left will target the effects of these linked tools at the selected hue range. Moving it to the right will target the effects at all hues other than the chosen range.

The threshold selector widget below the slider lets you specify which hues to target (or protect). By default, a range of skin hues is selected. The selection is in the L*a*b* space: 100*PI, + 100*PI.

Avoid Color Shift

Applies an action on the chrominance and luminance of the recomposed image in order to take into account the working colorspace and deliver the data via a relative intent within the gamut of the working space.

Residual Image

As a reminder:

- The residual image uses the same data format as the original image (source) before decomposition.
 ((what does this mean really, of what use to the user is this?))
- It corresponds to the difference between the original image and the sum of the decomposed levels (the action on each level has no effect on the residual image).
- The more levels are used, the more the residual image deviates from the original image.

Modifying the residual image is one of the key points of the wavelet processing. It allows one to:

- 1. Work on the shadows and highlights independently of details,
- 2. Reduce contrast (and chroma) overall, to better perceive micro-contrast,
- 3. Change chrominance to reduce artifacts which result from excessive action of the levels (e.g. to the skies),
- 4. etc.

Shadows/Highlights

Moving the Shadows and Highlights sliders to the right increases luminance in these areas; to the left - decreases. This action is comparable to that of the Shadows/Highlights tool in the Exposure tab. Beware though that it does not invoke the principle of highlight reconstruction.

The action of these sliders is influenced by the two threshold sliders:

- 1. Shadows Threshold, it is set to 30 by default (on the $o-100 L^*$ scale), which means that only luminance values around 30 and below will be influenced by the Shadows slider.
- 2. Highlights Threshold, it is set to 70 by default (on the 0-100 L* scale), which means that only luminance values around 70 and above will be influenced by the Highlights slider.

Read the L* value from the Navigator panel to help you decide the correct threshold values.

You could use the Shadows/Highlights adjusters to:

- 1. add bloom to bright objects,
- 2. prevent highlights from clipping,

- 3. lift shadows,
- 4. etc.

Contrast

This is one of the key points of the treatment made possible by wavelets. It allows you to change the contrast of the featureless residual image (remember, the residual image has no details, as those are handled by the contrast sliders) while letting you adjust the contrast of the detail levels using the contrast sliders above. For example by moderately reducing the contrast of the residual image, this leaves more room to accentuate that of each level, and thus increase the perception of depth and relief.

Chromaticity

The same principle as for contrast slider above applies to the use of the chromaticity slider, but in addition you can act selectively on colors by use of the Hue slider in conjunction with the Hue Targetting/Protection slider. By default a typical blue sky-colored hue range is selected, but of course you can change it to anything you like, for example to match the colors of a flower, or skin. The Hue Targetting/Protection slider lets you choose how you want the chromaticity change to affect the selected hue range:

- 1. By moving it to the left, only the chosen hues will be affected by your chromaticity change, e.g. select green grass hues, increase the Chromaticity slider, and move this Hue Targetting/Protection slider to the left to increase the chromaticity of only grass-colored in the residual image.
- 2. By moving it to the right, your chromaticity change will apply too all hues except those chosen by the Hue selector. So in the example in the previous point, if you were to move the slider to the right, the chromaticity of all hues except the grass-colored ones would increase.

This control is very useful in the prevention of over-saturation of human skin, a carrot look which our eyes are used to noticing.

Transform Tab

Crop

Cropping in RawTherapee does not discard the cropped-off area, it just hides it by drawing a crop frame and zooming in to fit the frame to the available screen space. You can still see the cropped-off area in the preview if you zoom out and set the background color of the preview to "theme-based". The cropped-off area will of course not appear in the saved image.

When you set the background color of the preview to anything other than "theme-based" then the cropped-off area will be completely hidden from the preview, but if you set it to "theme-based" then by default the cropped-off area will be covered by a dark but semi-transparent mask. You can define the color and transparency of this mask in "Preferences > General > Crop mask color/transparency".

Activate crop-placing mode by clicking the "Select Crop" button in the tool panel, the the button in the Editor's top toolbar, or the appropriate keyboard shortcut, then create the crop by clicking and dragging over the preview with your mouse. Use the **Shift** key to move (pan) the crop over the image. Resize a crop by placing the mouse on one of the sides or corners. To clear the crop, activate crop-placing mode again (via the keyboard shortcut or either of the buttons mentioned above), and click anywhere in the preview without dragging. To see only the cropped area, use the "Fit cropped area to screen" keyboard shortcut.

Use *Guide Type* to select popular guides to help you in composition while cropping, and a horizontal (landscape) or vertical (portrait) orientation. By default, as of version 4.2.214, RawTherapee automatically detects and uses the same crop orientation as the orientation of your image - the "As Image" option.

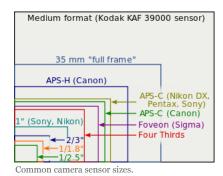
The PPI value (pixels per inch) does not change any physical property of the image, it only serves to help you see what physical size the current crop would print at using that PPI value.

Aspect Ratios

Custom crop ratio

You can make a custom crop ratio as of RawTherapee 5.1.

- 1. Disable "Lock ratio",
- 2. Make a crop, use "Width" and "Height" values to define the ratio, e.g. to make a crop using a 5:4 ratio set width=5 height=4, or width=1280 height=1024.
- 3. To resize the crop while maintaining the custom ratio, hold the Shift key while dragging the edge or corner of the crop frame with the mouse.



Fixed crop ratio

Use "Lock ratio" to set the crop to a fixed ratio.

3:2

Classic negatives have this ratio, as do APS-C (https://en.wikipedia.org/wiki/APS-C) DSLR (https://en.wikipedia.org/wiki/Digital_single-lens_reflex_camera) cameras.

4:3

The Four Thirds System (https://en.wikipedia.org/wiki/Four_Thirds_system).

16:9

The 1080p (https://en.wikipedia.org/wiki/1080p) and 720p (https://en.wikipedia.org/wiki/720p) high-definition (https://en.wikipedia.org/wiki/High-definition_video) video format, and due to this the most common computer monitor aspect ratio since 2010.

16:10

The most popular computer monitor aspect ratio between 2005-2009. Still popular in tablets.

24:65 XPan

Hasselblad's medium-format cameras.

1.414 DIN EN ISO 216

1.414 DIN EN ISO 216 (https://en.wikipedia.org/wiki/ISO_216) is the standard paper size ratio such as A4, B5, etc.

8.5:11

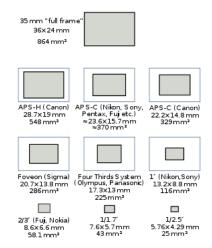
The US Letter (https://en.wikipedia.org/wiki/Letter_(paper_size)) size.

11:17 - Tabloid

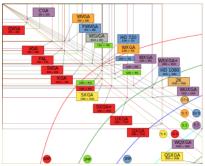
A common tabloid newspaper format (https://en.wikipedia.org/wiki/Tabloid_(newspaper_format)).

45:35 - ePassport

Guides to help you crop a portrait for a biometric passport (http s://en.wikipedia.org/wiki/Biometric_passport). Official measurements do not specify exact ratios, just min/max measurements within which the eyes and chin-crown distance must lie. The guides represent the averages of those distances. The first horizontal guide is for the crown, the second is roughly for the nostrils, the third is for the chin. "On the photo, the face must be between 29mm and 34mm from the bottom of the chin to the crown (the top of the head, not the top of the hair)." [7] (http://www.homeoffice.gov.uk/agencies-public-bodies/ips/passports/information-p hotographers/).



Common camera sensor sizes.



Video standards.

Resize

Resize

Resizing is one of the last things to happen when saving an image - this tool runs after most other tools and transformations. As downscaling an image involves a certain loss of detail, this tool includes a "Post-Resize Sharpening" component which you can use to make the downscaled image crisp.

The effects of the Resize tool will not be shown in the preview. The saved image will of course be resized.

Size limits:

- The downscaling limit is 32x32px.
- The upscaling limit is 16-times the image's size as of RawTherapee version 5.5. In versions prior to and including 5.4, it was 4-times the image's size.

Method

Resizing can be performed using various algorithms, each with its own strengths and weaknesses. We narrowed the choice down to the two most significant:

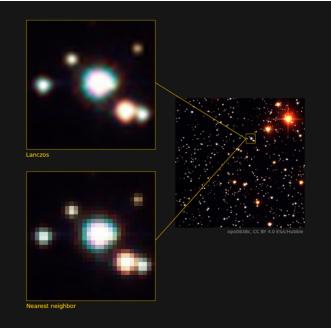
Nearest

The "nearest neighbor" method of interpolation is meant to be used when you want to upscale an image in order to make some detail larger without introducing smooth blending between pixels. It preserves the pixels as they are. It is usually used when upscaling by squares of 2: 2x, 4x, 8x, 16x, etc.. The result will be very sharp, but also very blocky. This method is not meant to be used for everyday photography.

Lanczos

Lanczos is the default method. It is meant to be used for everyday photography and in all cases except for the one described above. It results in a smooth yet sharp and high quality image. Use it to resize by any amount.

Screenshot of the whole Resize tool in RawTherapee 5.4-789-ga7b7f2ob8.



Demonstration of the effects of the resizing methods "nearest neighbor" vs "Lanczos"

Specify, Width and Height

You can specify what dimensions to scale to or by:

Scale

Use this if you want to resize the image by a specific amount. e.g. to make it a quarter of its current size, set scale to 0.25, or it make it 4x larger, set scale to 4.

Width

Use this to specify the desired width in pixels, and have the height adapt while keeping proportions.

Height

As for width.

Bounding Box

Use this if you want the image to fit inside a certain width and height.
e.g. if you want the longest dimension of your image to be 1000px regardless whether the image is landscape or portrait, set a bounding box of 1000x1000px.
e.g. if your image has dimensions 3000x2000 and you want it to fit on a 1080p/Full HD (https://en.wiki pedia.org/wiki/1080p) screen, set the bounding box to 1920x1080 and your image will be resized proportionally in such a way that the width will not exceed 1920 and the height will not exceed 1080.

Applies To

You will typically want the width and height to apply to the cropped area, however you can also have it apply to the full image, though even then of course only the cropped area gets saved.

Allow Upscaling

This setting controls whether the resize tool will also upscale your images, or only downscale. If it is disabled and the combination of image dimensions, crop dimensions and resize dimensions is such that the image would need to be upscaled to meet your target, this setting will prevent it from being upscaled it will remain at its current size.

This setting was introduced in RawTherapee 5.5. If in RawTherapee 5.5 you use a sidecar file from an older version of RawTherapee, this setting will automatically be disabled.

Post-Resize Sharpening

Resizing an image often leads to a loss of sharpness, so it is common practice to sharpen the image again after having resized it. With the Post-Resize Sharpening tool you can save crystal-clear images straight away with no further hassle. Because this tool works on the image after it is resized, you cannot use the preview to see what it will do, though this is not a problem because the procedure for finding the right values is straightforward.

The default values work great, but if you want to change them, here's how to preview what the image will look like:

- 1. Tweak your image as you usually would and enable the Resize tool, e.g. downscale using the Lanczos method to a 900x900px bounding box.
- 2. Save the image to a lossless format such as TIFF.
- 3. Open that saved TIFF in RawTherapee, apply the Neutral processing profile if that wasn't done automatically, and enable the Sharpening tool in the Detail tab.

- 4. Zoom to 100% (1:1) and tweak the Sharpening tool's parameters until you get a result that satisfies you. These are the values you should use in the Post-Resize Sharpening tool.
- 5. Go back to the raw image, enable the Post-Resize Sharpening tool and set it up with the values from the previous step.

The Post-Resize Sharpening tool is only available when you use the "Lanczos" resizing method.

As the Post-Resize Sharpening tool works identically to the standard Sharpening tool (except that it takes place right at the end of processing), refer to the Sharpening article to learn more about how the sharpening tools work.

Lens/Geometry

To keep the preview fast, RawTherapee uses the preview image of the current zoom level when applying these transformations. Because of this, the preview image can become soft. Lets assume you are editing a Nikon D700 image: 4256×2832px (that's 12.1 megapixels), and the preview image's size is 600x400px. Rotating it 5° will not be the same as rotating the full 12.1Mp image and then scaling it down to 600x400px. The former will be softer than the latter, though rotating the former will take less time than rotating the latter, which is why RawTherapee does that. But don't worry, when saving the image RawTherapee uses the full resolution image, so it will be sharp. If you zoom the preview in, then RawTherapee will use this higher resolution preview image when calculating the transformation, so to see what the saved file will look like, just zoom in to 100% .

Lens Hood

Do not confuse vignetting with a blurry lens hood being visible in the corners of your image. Some cameras, typically small ones - compacts, bridge-type and even mirrorless - will capture parts of the lens hood or lens mechanism in the corners of the frame. Typically the same cameras have lenses which suffer from strong distortion. The way these cameras deal with these two problems is by correcting the distortion, as a result of which the corners of the image get "pushed out" beyond the frame, thereby hiding the corners darkened by the lens or lens hood. When you view a JPEG image from these cameras the distortion has already been corrected in-camera, so owners are often unaware that the problem existed in the first place and are surprised to find that the raw image shows these dark corners.

It is not possible to fix the dark corner issue by using vignetting correction - there is no information about the scene in those corners, the scene is occluded by the lens mechanism/lens hood. Do as the camera does: correct the distortion, and the dark corners disappear

Auto-Fill

This option will upscale or downscale the photo to the extent that the whole image fits within the image boundaries with no black borders visible.

When correcting images that suffer from barrel distortion (https://en.wikipe dia.org/wiki/Distortion_(optic s)), "Auto-fill" will perform downscaling to fit as much of the reprojected image as possible into the image boundaries, so that you don't unnecessarily lose any parts of the image. Conversely if the image



Barrel distortion correction leaves black empty space in image periphery.

suffers from pincushion distortion, "Auto-Fill" will upscale the corrected image to fill the frame without black borders around the periphery.



"Auto-fill" enlarges the image to fill in the empty space.

Auto-Crop

"th Auto-Crop" is available when "Auto-fill" is disabled. When activated, it will not cause image interpolation, but instead will crop away the empty space left by the distortion correction or image rotation.



"Auto-Crop" applied after distortion correction.



"Auto-Crop" applied after image rotation.

Rotate



Example of image rotation.

Rotate the image between -45° and +45°. Use the "Select Straight Line" button to set either a vertical or a horizontal image alignment. Use the mouse to draw this line - click and hold mouse to start, move to draw a new vertical or horizontal axis and release to engage image rotation.

Perspective

Horizontal



Example of horizontal perspective correction.



Example of vertical perspective correction.

When your picture was taken while you were slightly off-center of the object, you can correct this (within certain limits) with the horizontal slider.

Vertical

Very useful to correct 'falling lines', e.g. when photographing architecture. Higher values for both sliders produce heavy distortion, so use with care. Or don't care at all and have fun!

Profiled Lens Correction

As of RawTherapee 5.3, lens distortion, vignetting and chromatic aberration can be corrected using one of two systems: Lensfun or Adobe LCP. Both rely on a profile which contains the parameters for correcting lens distortion, vignetting and chromatic aberration.

Lensfun

Lensfun is an open-source project. A database of lens parameters is maintained by the developers. You can find more information in the Lensfun FAQ (http://lensfun.sourceforge.net/faq/).

The options "Auto-matched correction parameters" and "Manual correction parameters" both use the Lensfun system. Both the "Camera" and "Lens" comboboxes will contain long lists of cameras and lenses (if they do not, your RawTherapee build is buggy - please report that). A profile matching your camera and lens will be auto-detected if your photo contains metadata which matches an entry in the Lensfun database.

Common causes of failure to auto-detect a profile include:

- A profile for your camera/lens combination does not exist in your version of the Lensfun database.
- A profile for your camera/lens combination does exist, but it uses names and parameters which differ enough to confuse the matching algorithm, for example "Pentax" vs "Ricoh Pentax", or "F4.0" vs "f/4".

In both cases you could try to find and run the executable <code>lensfun-update-data</code> to download the latest version of the Lensfun database. If that does not help, you could copy and paste the relevant section from the Lensfun database (which could be taken from one of the files in <code>/usr/share/lensfun/</code>) into a new file <code>\$HOME/.local/share/lensfun/myLensfun</code> and modify the relevant parameter to match the metadata from your photos. You can find the camera and lens name and parameters contained in your photos by viewing the Exif quick info panel in RawTherapee, shortcut key "i" in the Editor tab. If parameters for your camera and lens simply don't exist, view the Lensfun documentation (http://lensfun.sour ceforge.net/calibration/) to find out how to measure them and contribute them for everyone's benefit.

Note that while editing the Lensfun database in /usr/share/lensfun/ directly may be possible, this is not recommended because you could lose your changes during an update.

LCP

Adobe provides and offers tools to create and share what are called Lens Correction Profiles. These are text files which describe the distortion, vignetting and chromatic aberrations (CA) of a lens, so that simply loading this file in LCP-capable software such as RawTherapee will correct these issues. Select an Adobe LCP (http://www.adobe.com/products/photoshop/extend.displayTab2.html#resources) file (read the guide on how to get LCP profiles) to automatically correct geometric distortion, vignetting and lateral chromatic aberrations.

The Lens Correction Profile tool's distortion correction feature can be used together with the manual Distortion Correction tool, and the vignetting correction feature can be used together with the manual Vignetting Correction tool. This lets you use the manual controls in addition to the LCP profile for artistic reasons or if the LCP fails to sufficiently correct a parameter (which happens on some extreme distortion occasions, like with some heavily distorting compact cameras). Be careful that you don't overdo the distortion and vignetting correction by forgetting to turn the manual tools off if you use the LCP equivalents. The vignetting correction feature however is linked to the Flat Field tool, so that when you select a flat-field image then the LCP's vignetting correction will have no effect.

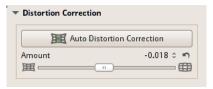
The following restrictions apply:

- Distortion, vignetting and chromatic aberration correction are all supported in raw files, but only distortion correction is supported in non-raw files.
- While distortion correction is visible in the full image preview, chromatic aberration and distortion correction are not reflected in the detail crop windows, only in the fully processed result image. Autofilling is also not supported.
- Chromatic aberration correction is only supported if the Exif information contains the focus distance (e.g. in DNGs from Nikon files).
- Auto-Fill is disabled when an LCP with distortion correction is enabled, otherwise the preview may become distorted see bug 1791 (https://github.com/Beep6581/RawTherapee/issues/1791).
- To keep the preview fast and responsive, the main preview image is used to show the effects of the LCP. As this image is small (exactly the size you see), fixing the distortions will make it appear a little blurry. This has no effect on your saved image, which will be sharp, and so will the zoomed-to-100% preview. Only the zoomed-out preview that will look soft. See feature request 2186 (https://code.google.com/p/rawtherapee/issues/detail?id=2186).

As with any other tool, you can apply an LCP to multiple images either by including it in the processing profile (see Creating processing profiles for general use), or by selecting multiple images where the same lens was used (you can use the Metadata Filter in the File Browser tab to make this easier) and applying the LCP from the File Browser tab.

Distortion Correction

Corrects lens distortion. A negative number corrects barrel distortion, a positive value will correct pincushion distortion. You can place a grid over the image by activating Crop (without cropping) and using "Guide Type > Grid". This may serve as a guide to correct lens distortion.



This button has been renamed to "Automatic" in RawTherapee 5.

The "Automatic" button only works if your camera corrected the distortion of the JPEG image embedded in the raw file (most cameras embed a JPEG image in every raw file, and some cameras correct the distortion of that image too). What this feature does is it looks at the JPEG image and, if it was corrected, tries to fix distortion in the raw image by making it match the JPEG image. There are two limitations:

- If the JPEG image was not distortion-corrected by your camera, this button will have no effect.
- If the JPEG image is insufficiently corrected or over-corrected, so will the results be, but as the computed correction will be shown on the *Amount* slider, you can further refine it manually.

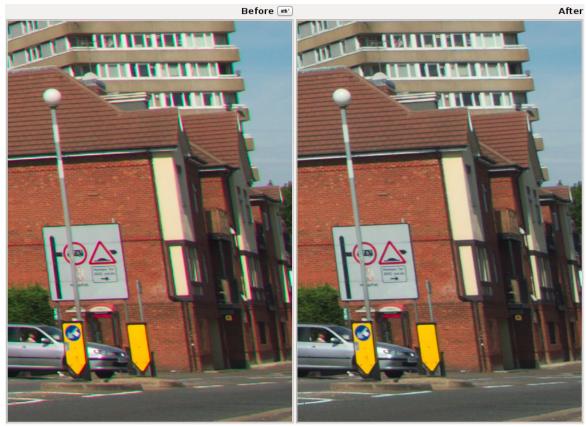
Chromatic Aberration Correction

This "Chromatic Aberration Correction" tool in the *Transform* tab works on the image **after** demosaicing. The Chromatic Aberration tool in the *Raw* tab works on the image **before** demosaicing.

Chromatic aberration can be corrected by using the "Red" and "Blue" sliders. Normally you won't see any chromatic aberration in the fit-to-screen preview, therefore it is highly recommended to open a detail window " or to zoom the main preview in to 100% or more when you attempt this kind of correction. As in other software tools, this algorithm eliminates moderate chromatic aberration quite well. Do not expect miracles with images having extremely high chromatic aberration - garbage in, garbage out.



400% zoom showing correction of CA.



CA easily visible even at 100% corrected.

Vignetting Correction

Vignetting means light fall-off around the periphery of an image as compared to the center. One of the differences between a cheap lens and an expensive one is that the former is likely to produce stronger vignetting than the latter. The "Vignetting Correction" tool is meant to correct vignetting caused by the lens. This tool is not intended for artistic vignetting; use the Vignetting Filter tool for that.

Amount

Setting the "Amount" slider to a positive value brightens the four edges of the images to correct the classical vignetting. Setting it to a negative value darkens them.

Radius

Influences how much of the image beginning from the edges will be brightened or darkened. Lower values: area of darkening is bigger; higher values: area of darkening is smaller.

Strength

Amplifies the settings of the "Amount" and "Radius" sliders. Set "Amount" to -100, "Radius" to 50 and move "Strength" from 1 to 100 to see how this works.

Raw Tab

Demosaicing

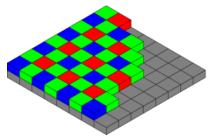
The effects of this tool are only visible at a preview scale of 1:1 or more. Use a detail window (click on the icon under the main preview panel) to inspect a part of the image, or zoom the main preview to 100% (also called 1:1) .

Most digital cameras today use a color filter array (https://en.wikipedia.or g/wiki/Color filter array) over their sensors. To display such raw files their data needs to be demosaiced (http://www.cambridgeincolour.com/tutori als/camera-sensors.htm). Cameras with a Foveon X3 sensor (https://en.wikip edia.org/wiki/Foveon_X3_sensor) (Sigma) do not have color filter arrays and so do not need demosaicing. Demosaicing is why opening a raw file always takes a bit longer than opening a JPEG (https://en.wikipedia.or g/wiki/JPEG) or TIFF (https://en.wikipedia.org/wiki/TIFF) file, where the data is already 'display-ready'. RawTherapee offers several demosaicing algorithms, each with its own characteristics. The differences between them are often very subtle - one might need to zoom in to 200-400% to discern them - but since the program works on a pixel-by-pixel basis and demosaicing is the basis upon which all other tools work, the choice of demosaicing algorithm can have a visually significant effect when combined with other tools, such as the sharpening ones. The choice of a certain algorithm influences, among other things, the quality of very fine details in the image, whether false maze patterns will appear, and decides how well colored edges are rendered.

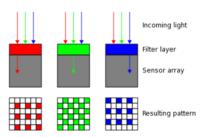
RawTherapee supports demosaicing images from sensors with Bayer filters (https://en.wikipedia.org/wiki/Bayer_filter) and Fujifilm X-Trans filters (https://en.wikipedia.org/wiki/Bayer_filter#Fujifilm_.22X-Trans.22_filter). If you take a look inside the "Raw" tab, you will notice there are two tools: "Sensor with Bayer Matrix" and "Sensor with X-Trans Matrix". The settings in one of these two tools have no influence over the settings in the other - if you open a raw image from a Bayer-type sensor, only the settings from the "Sensor with Bayer Matrix" tool will be used, the settings from the "Sensor with X-Trans Matrix" tool will be ignored, and vice versa if you open a raw image from an X-Trans type sensor. For simplicity, we will describe both here.



Cutaway illustration of a camera showing the light sensor with a Bayer filter.



Bayer pattern on sensor.



Profile/cross-section of sensor.

Method

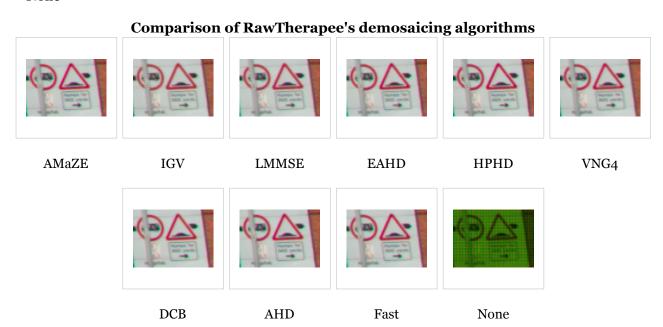
The following demosaicing algorithms are available for raw files from Bayer sensors:

- AMaZE
- AMaZE+VNG4
- RCD
- RCD+VNG4
- DCB

- DCB+VNG4
- LMMSE
- IGV
- AHD
- EAHD
- HPHD
- VNG4
- Fast
- Mono
- Pixel Shift
- None

The following demosaicing algorithms are available for raw files from X-Trans sensors:

- 3-pass+fast
- 3-pass (Markesteijn)
- 1-Pass+fast
- 1-Pass (Markesteijn)
- Fast
- Mono
- None



Which demosaicing method should you use?

This page aims to tell you as much about the various algorithms as is relevant to a photographer, but there is no much to say as the explanations would quickly become technical and of a programmatical and mathematical nature. After reading though this article you should know that LMMSE and IGV are to be used on high ISO photos and that for the majority of other cases you should stick with the default AMaZE method, but of course you are free to explore for yourself and to test each method out on your own raw files. It is of no use reading an article somewhere on the internet which claims that some specific method is best because in their test that method was the sharpest, as the performance of each method depends specifically on the sensor in your camera and even on ISO, so keeping our suggestions in mind, run your own test and make up your own mind!

RawTherapee versions 4.2.91 and newer always use the demosaicing method you choose regardless of zoom level. RawTherapee versions older than 4.2.91 use the Fast algorithm to initially open the image for editing. After this, the selected demosaicing method is applied when the image is zoomed to 100% magnification or when the detail window is opened. The selected method is also used for batch processing. It is not recommended to select the Fast method for the final conversion, as it is a low-quality algorithm for display purposes.

AMaZE (Aliasing Minimization and Zipper Elimination) is the default demosaicing method, as it yields the best results in most cases. In RawTherapee versions 2.4 and older VNG4 used to be the preferred algorithm for Olympus cameras, as AMaZE didn't exist yet and VNG4 eliminated certain maze pattern artifacts that might have been created by the other methods, but with the introduction of the AMaZE method in RawTherapee version 3.0 Olympus users might prefer that instead.

RCD does an excellent job for round edges, for example stars in astrophotography, while preserving almost the same level of detail as AMaZE.

DCB produces similar results to AMaZE. AMaZE can often be slightly superior in recovering details, while DCB can be better at avoiding false colors especially in images from cameras without anti-aliasing filters.

When working with very noisy, high ISO images in conjunction with the Noise Reduction tool, it is recommended to use the **LMMSE** or **IGV** demosaicing methods. They will prevent false maze patterns from appearing, and prevent the image from looking washed-out due to heavy noise reduction.

If you use a medium format technical camera with near-symmetrical wide angle lenses such as the Schneider Digitar 28mm or 35mm it's likely that your file will contain some crosstalk, especially if the lens is shifted (due to the low angle of incoming light from these lenses some light leaks over to the next pixel on the sensor), and in this case you can get mazing artifacts with AMaZE and DCB because of green channel separation caused by the crosstalk. If you, via adapters, combine a mirrorless camera with a wide angle lens designed for film, you may also get crosstalk. It can then be better to use the more robust **VNG4** algorithm (Variable Number of Gradients) which handles this situation well, at the cost of some fine detail. An alternative is to enable green equilibration to even-out the green channel differences.

AHD (Adaptive Homogeneity-Directed), **EAHD** (Horvath's AHD) and **HPHD** (Heterogeneity-Projection Hard-Decision) are old methods which are generally slow and inferior to the other methods.

None means no demosaicing is performed. This can be useful for diagnostics, but you would not use it for photography.

Mono is only useful for users of either monochrome cameras, or cameras with the color filter array removed.

Fast is a very fast but simple and low quality demosaicing method, not recommended.

- **3-Pass** is a demosaicing method for cameras with X-Trans sensors (Fuji). It runs three passes over the image which leads to sharper results though you can only see this on low ISO photos. It is slower than 1-Pass.
- **1-Pass** is a demosaicing method for cameras with X-Trans sensors (Fuji). It is faster than the 3-Pass method though a bit inferior in quality, though this difference is only visible in low ISO shots, so if speed is an issue you can use this method on high ISO shots with no visual difference in quality.

False Color Suppression Steps

Sets the number of median filter passes applied to suppress demosaicing artifacts when applying the demosaicing algorithm. False colors (speckles) could be introduced during the demosaicing phase where very fine detail is resolved. False color suppression is similar to color smoothing. The luminance channel is not affected by this suppression.

False colors are generally more apparent in images from cameras without anti-aliasing filters. Note that it is foremost the chosen demosaicing algorithm which is the deciding factor in how prominent will be the false color problem with which you will have to deal. In some situations it may be better to change the demosaicing algorithm than to enable false color suppression, as the latter reduces color resolution.

How to Find the Best Demosaicing Method

Zoom in to at least 100% (1:1) and try all the demosaicing algorithms, see which works best for you. Try them on sharp photos with fine detail and tiny patterns, such as the wavy and repetitive fabric of a sweater (watch for maze pattern artifacts), a distant brick wall, a distant round road sign (watch for aliasing along the round edges), and test with both low and high ISO shots. Use photos from your own camera; what's best for Nikon raw images may not be what's best for Olympus ones.



A good image to test the demosaicing algorithms on. Zoomed in to 800%, you can clearly see that VNG4 is not a good choice for this Pentax K10D raw file, as there are dots where there should be none, and the detail of the wall's brickwork (the orange part on the right) is all washed out.

Monochrome Cameras

A monochrome camera has the same light filter in front of all pixels, that is you get a black-and-white image and no demosaicing is required. Some of these cameras have no infrared filter and are thus sensitive to infrared light, which can be used for creative black and white photography.

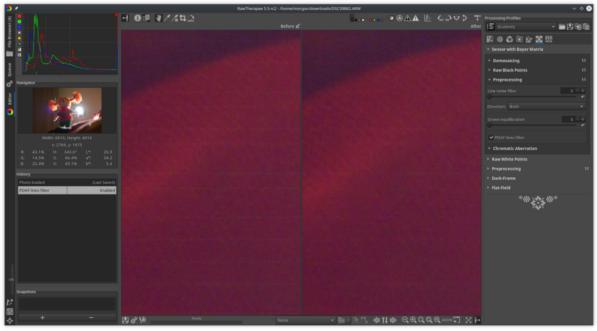
RawTherapee supports monochrome cameras, but the user interface is not adapted for it so when you load a monochrome file all color tools will still be enabled (they won't do anything meaningful of course). You will have to live with that, monochrome cameras are rare so we won't put any major effort into making the user interface morph into reduced monochrome-only version.

There are a few additional factors to consider when working with monochrome files: some monochrome cameras report that they have only a single monochrome channel and a neutral color matrix, like Leica M Monochrom, while others report RGB channels in a bayer configuration (like Phase One IQ260 Achromatic, or IR-modified cameras). If the camera reports only one channel, RawTherapee recognizes this and won't perform any demosaicing (the demosaicer selection is still enabled but does not do anything), and everything works normally. However, if the camera reports as an RGB bayer camera, demosaicing will be performed and a color matrix will be applied. To disable this, you should select the "Mono" demosaicing option, and select "No profile" as input profile in the Color Management panel.

Preprocessing

There are several preprocessing settings and they are split into two parts in the user interface, both in the "Raw" tab: one part appears as the main "Preprocessing" tool, and the other part appears within the "Sensor with Bayer Matrix" tool. The preprocessing settings in the main "Preprocessing" group affect raw files of all types, while the ones in the "Sensor with Bayer Matrix" tool affect only raw files from cameras which use a bayer filter.

PDAF Lines Filter



Line artifacts caused by PDAF pixels visible on the left, and fixed on the right.

Cameras with Phase Detection Auto Focus (PDAF) are susceptible to producing line striping artifacts when photographing backlit scenes with visible flare. Some mirrorless interchangeable lens cameras from Sony are known to suffer from this.

Enable the PDAF lines filter to fix these PDAF line striping artifacts.

RawTherapee 5.5 can fix PDAF line artifacts for the following cameras:

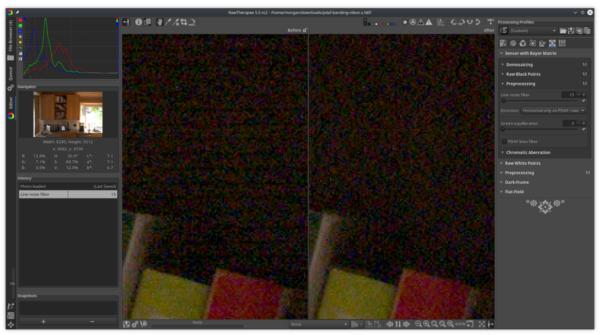
- Sony DSC-RX1RM2
- Sony ILCE-6000
- Sony ILCE-6300
- Sony ILCE-6500
- Sony ILCE-7M3
- Sony ILCE-7RM2
- Sony ILCE-7RM3
- Sony ILCE-9

Line Noise Filter

Line noise appears as horizontal or vertical banding most visible in noisy images. It is caused by noise in the sensor electronics which read the value of each photosite row by row or column by column.

You can see examples of what line noise looks like at this MagicLantern forum post: http://www.magiclantern.fm/forum/index.php?topic=10111.msg105001#msg105001 (http://www.magiclantern.fm/forum/index.php?topic=10111.msg105001#msg105001)

Nikon PDAF Banding



Banding artifacts caused by too-aggressive in-cameras PDAF filtering visible on the left, and fixed on the right.

Several mirrorless interchangeable lens cameras from Nikon use Phase Detection Auto Focus (PDAF), and they perform in-camera filtering to mitigate the PDAF striping artifacts. However, the filtering performed seems excessive and leads to "PDAF banding" - darker lines in shadow areas.

Use the line noise filter to fix PDAF banding artifacts with the direction set to "Horizontal only on PDAF rows".

RawTherapee 5.5 can fix Nikon PDAF banding artifacts for the following cameras:

- Nikon Z 6
- Nikon Z 7

The "PDAF lines filter" has no effect on Nikon PDAF banding.

Green Equilibration

Some cameras (for example Olympus, Panasonic, Canon 7D, and some medium format cameras) use slightly different green filters in the two green channels of the color filter array (https://en.wikipedia.org/wiki/C olor_filter_array) on the camera sensor. This is generally not a designed feature of the sensor, but rather a result of limitations in the manufacturing process when the color filters are applied to the sensor surface. One green filter may get a small pollution from the red filter and the other from the blue for example. Green equilibration suppresses interpolation artifacts that can result from using demosaic algorithms

which



Cross-hatch pattern due to imbalance of the two green channels.

assume identical response of the two green channels. The threshold sets the percentage difference below which neighboring green values are equilibrated.

Set the value high enough for the mazing to disappear but no higher. The DCB demosaicing algorithm is very sensitive to green split so it is good to use while trying to find the best value.

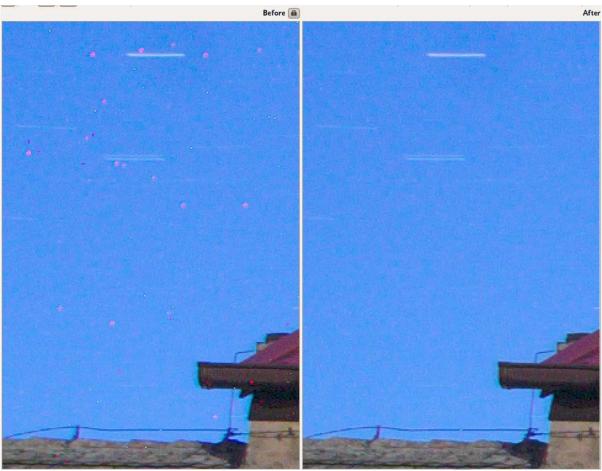
Green equiliberation can also be used to equalize green split caused by crosstalk. If you, for example via an adapter, use an analog ultra-wide angle lens with your digital camera the incoming light may arrive at such a low angle that some light passes through one color filter and gets registered in a neighboring pixel belonging to a different color channel - this is crosstalk. As one green channel has blue neighbors and the other red, the first will get crosstalk from blue and the other from red, hence they will separate which can cause mazing. Also red and blue channels will in such a situation suffer from crosstalk, but as they only receive from green there's no separation in those channels. Mild crosstalk will not have any visible effect if green is equalized, while heavy crosstalk will show as desaturated dull colors (as the color channels have been mixed). Note that crosstalk generally doesn't occur before heavy color cast, so in this case you will be using flatfield correction too.

Hot/Dead Pixel Filter

This tool suppresses hot and dead pixels (https://en.wikipedia.org/wiki/Defective_pixel) by replacing them by a neighborhood average.

"Hot pixels" appear as bright and saturated tiny dots. Each one is the result of a photosite on the sensor outputting a higher current than it should. Whether a single photosite on the sensor corresponds to a single pixel in the processed photo depends on the chosen demosaicing method (and other tools); most methods, such as the default AMaZE, do not have a direct photosite:pixel relationship, and so hot pixels may appear not only as single-pixel dots but also as tiny 3x3 pixel crosses or slightly larger blobs. Hot pixels are a completely normal phenomenon present in all cameras, however in typical daylight photography you will not encounter them. The longer the exposure, the higher the chance of hot pixels appearing and the larger their number, with exposures longer than two seconds generally being regarded as prone to this problem. Heat is also a factor, hot sensors being more prone.

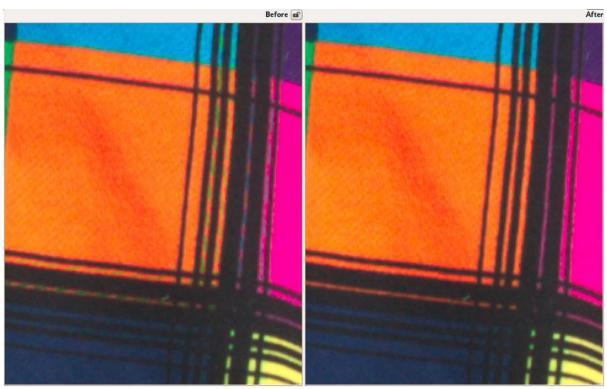
"Dead pixels" on the other hand appear as black dots (or crosses or blobs). They are the result of dead



Long-exposure shots are prone to suffer from hot pixels, but RawTherapee can easily fix that.

photosites on the sensor, and as such, exposure time does not have any influence on whether they appear or not - if a photosite is dead, every photo will have the dead pixel in the same spot. Because their position is static and always present when using the same camera body, you can fix them not only using the automatic "Dead pixel filter" but also by adding their coordinates to a *.badpixels* file; see Bad Pixels.

It is impossible to detect hot and dead pixels with absolute certainty by analyzing only one photo (as opposed to analyzing a whole series of photos), and as such one must find the balance between adequate removal and false positives (http://en.wikipedia.org/wiki/False_positives_and_false_negatives). The threshold slider allows you to set the sensitivity of the automatic detection of hot and dead pixels. Lower values make hot/dead pixel detection more aggressive, but false positives may lead to artifacts. If you notice any artifacts appearing when enabling the Hot/Dead Pixel Filters, gradually increase the threshold value until they disappear.



Artifacts introduced by a too-aggressive hot/dead pixel automatic detection threshold value. Increasing the threshold value from 40 to 120 eliminated the artifacts.

Chromatic Aberration

This "Chromatic Aberration" tool works on the image **before** demosaicing, that's why it's located in the *Raw* tab. The Chromatic Aberration Correction tool in the *Transform* tab works on the image **after** demosaicing.

Chromatic aberration correction on the raw level is currently only supported for raw files from cameras with a Bayer filter (https://en.wikipedia.org/wiki/Bayer_filter). If you need to remove chromatic aberration from raw photos from X-Trans sensor cameras (Fuji), then use the Chromatic Aberration Correction tool in the *Transform* tab.

Chromatic aberration can be corrected by using the "Red" and "Blue" sliders. Normally you won't see any chromatic aberration in the fit-to-screen preview, therefore it is highly recommended to open a detail window [47] or to zoom the main preview in to 100% (4) or more when you attempt this kind of correction.

This tools corrects bluish-green and magenta fringes due to lens lateral chromatic aberration that show mainly in the borders of the image. This correction is performed before demosaicing and can sometimes improve the quality of demosaicing.

Auto-Correction

If "Auto-correction" is checked, the "Red"/"Blue" sliders are disabled and an automated detection and correction of chromatic aberration is performed. Where manual correction applies a constant value across the image, auto-correction divides the image into many blocks and tailors the values required to eliminate chromatic aberration to each block. For this reason auto-correction usually performs better than manual



400% zoom showing correction of CA.

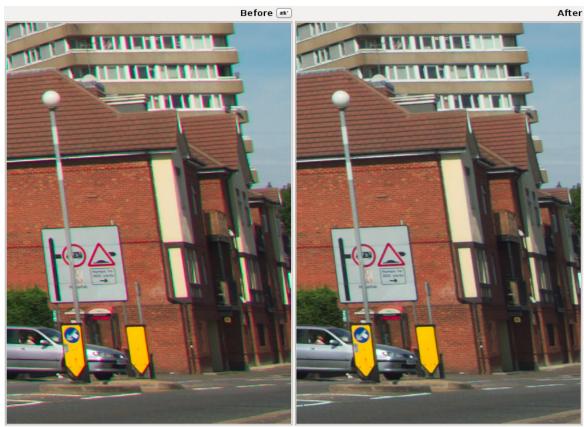
correction, and the auto-correction values cannot be displayed in the sliders.

Iterations

This setting is available if "Auto-Correct" is checked. Auto-Correction is conservative, means it often does not correct all Chromatic Abberation. To correct the remaining Chromatic Aberration, from RT 5.5 on you can use up to 5 iterations of automatic Chromatic Aberration correction. Each iteration will reduce the remaining Chromatic Aberration from the last iteration at the cost of additional processing time.

Red/Blue

If the "Red"/"Blue" sliders are non-zero the given values are used to correct chromatic aberration. They cannot be used at the same time as "Auto-correction".



CA easily visible even at 100% corrected.

Raw Black Points

It is unlikely you will ever need to use the Raw Black Points tool other than for diagnostic purposes.

These sliders allow you to add offsets to the raw channel values in case the values used in the program are incorrect. You would know that they are incorrect if the colors of your image were severely wrong, for example you would see dark green instead of black, or the whole image would have a very strong magenta color cast. Cameras with a Bayer filter (https://en.wikipedia.org/wiki/Bayer_filter) have two green channels, that is why there are two green sliders which are linked by default. Sometimes the two green channels are of different sensitivity, so you can unlink the two green sliders and adjust them individually.

Raw White Points

It is unlikely you will ever need to use the Raw White Points tool other than for diagnostic purposes.

The white-point correction adjuster scales the current white level for all channels. It scales the raw samples linearly after black-level subtraction and can be used to simulate a new white level. This can be necessary if RawTherapee's raw decoder has an incorrect white level, which can happen for new or rare cameras. Flat Field correction may also, in some cases, as a side-effect, scale down raw samples such that blown highlights become discolored (usually magenta/pink); in this case this setting can be used to scale up the raw data so it hits the white level again.

Dark Frame

Dark-frame subtraction (https://en.wikipedia.org/wiki/Dark-frame_subtraction) is a method of dealing with thermal, dark-current and fixed-pattern noise. It is not effective against high ISO noise because of its different random nature. In long exposure shots (more than 1 sec) the non-homogeneous thermal noise becomes evident, mainly due to unevenness of the sensor and surrounding electronics. A method to mitigate this effect is to subtract one (or more) shots taken in the same conditions, but with the lens cap on. Only raw images for the same camera model can be used as dark-frames, preferably taken around the same time as the photo they are being subtracted from. As taking a dark-frame shot is simply a matter of putting the lens cap on and pressing the shutter without changing any settings, this is not a problem.

In the "Dark-Frame" panel, you can specify a single shot to subtract from the image, or check "Auto-Selection" and let RT choose the best match from the directory specified in "Preferences > Image Processing > Dark-Frame". Under the widget, RT shows how many shots are found and how many groups of shots are found and averaged into a template. From now on, put your dark-frame shots there if not already done. You could also move a shot from the "File Browser" tab into the dark-frames directory by right-clicking on it and selecting "Dark-Frame > Move to dark-frames directory". RT chooses the best match looking for the same camera model with minimal difference in ISO, exposure duration and date. If more than one shot with exactly the same properties is found, then an average of them is used: this produces by far less noise, so it's better to have 4-6 frames taken in the same conditions of the actual photo.

When selecting a dark-frame (or with "Auto-Selection"), RawTherapee extracts from it all the positions of hot pixels and then always corrects them in the final image. This correction is better than applying only the "Hot/dead pixel filter", but works only for hot (=white) pixels not for dead (black) ones.

Auto-Matching Logic

Key for dark frames (dfInfo::key):

- camera manufacturer
- camera model
- ISO
- shutter speed

The search for the best match is two-fold:

- if perfect matches by key are found, then the list is scanned for lesser distance in time,
- otherwise if no match is found, the whole list is searched for lesser difference in ISO and shutter.

Bad Pixels

RawTherapee can correct a list of bad pixels (https://en.wikipedia.org/wiki/Defective_pixel) (pixels that are always black or white or stuck at one color) for your particular camera model. To do this, you need to write a text file with the absolute raw coordinates of these pixels: each line specifies a pixel with x<space>y<return> positions.

Important: RawTherapee cuts some pixels from the top and left edges of the raw image's border (because they can't be interpolated correctly). If you look at the pixel coordinates in RawTherapee, beware of the offset introduced by this cutting. To each coordinate you must add 4 for Bayer Sensors and 7 for X-

Trans Sensors! Alternatively you can specify the offset (4 for Bayer, 7 for X-Trans) in the first line of the .badpixels file.

The file has to be located in your dark-frames directory. Set it by going to "Preferences > Image Processing > Dark-Frame". The file has to be named exactly as your camera's make and model: "make model.badpixels". Get the make and model as RawTherapee expects them by opening a raw image you want corrected in the Image Editor tab and looking at the displayed name and model in the "Quick info overlay, shortcut "i", e.g.: "Pentax K200D.badpixels"

Remember that these .badpixels files must be saved to the dark-frames directory!

If you've done the steps correctly and it still doesn't work, you can verify that your badpixels file is being read by closing RawTherapee, editing the "options" file in a text editor and changing "Verbose=false" to "Verbose=true", then starting RawTherapee from a console, opening the photo you want fixed, and looking at the text output in the console. If you see a message like "Pentax K10D.badpixels not found" then you know what to rename your badpixels file to. Once you get it working, remember to set "Verbose" back to "false".

Pixels in the bad-pixels list will always be corrected in processed photos as long as the camera make and model matches the badpixels filename.

Bad pixel detection software

Programs exist to aid in the detection of bad pixels:

Dead Pixel Test

http://www.starzen.com/imaging/deadpixeltest.htm (dead)

Mirror:

https://web.archive.org/web/20140725130003/http://www.starzen.com/imaging/deadpixeltest.htm (https://web.archive.org/web/20140725130003/http://www.starzen.com/imaging/deadpixeltest.htm)

Mirror: http://rawtherapee.com/mirror/deadpixeltest.zip

This file is free from viruses (https://www.virustotal.com/en/file/11e7aodb897fd3ad9f3e24c97c73b178cfe9f9d246e3dadfe57113318e2defo6/analysis/1421736881/).

Pixel Fixer

http://www.pixelfixer.org/

Remember to fix the 4px or 7px cutting offset if your version of RawTherapee needs that.

Flat Field

Flat-field correction (https://en.wikipedia.org/wiki/Flat-field_correction) is used to compensate for the non-uniformity characteristics of the camera and lens combination. A well known example of such non-uniformity is vignetting - a peripheral darkening of the image, more pronounced in the corner areas. Another example, more familiar to users of digital medium format cameras, is the lens cast effect - both color and luminance non-uniformity of the image field. Both of these examples of non-uniform image capture can be further complicated



Example of flat-field correction. From the left: before correction, after correction, and the flat-field image used. Notice the correction of both luminosity and color non-uniformities.

by a possible misalignment of the lens mount or by the usage of a tilt-shift lens. Another set of examples of capture non-uniformity is due to light leakage in the camera, thermal non-uniformity of the sensor or defects/irregularities in the sensor readout electronics. An alternate name for this type of function used by some other raw converters is "LCC correction", ie "lens color cast correction" or just "lens cast correction", but flat-field correction is the better name as it can as decribed correct for many sorts of non-uniformity not just lens cast.

Manual correction of these effects in post-production is quite difficult, especially when needed to be reproduced on a series of images captured under various conditions, and would rarely be perfect.

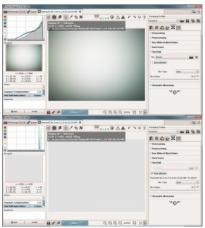
The "Flat-Field" correction tool in RawTherapee allows both automated and user-guided modes. Flat-field correction is performed only on linear raw data in the beginning of the imaging pipeline and does not introduce gamma-induced shifts. Thus in RawTherapee flat-field correction can be applied to raw files only.

Due to performance considerations, thumbnail images do not reflect flat-field corrections. At present only the main preview in the Editor tab and the output image can be flat-field corrected.

Accuracy of the flat-field correction is largely based on use of the appropriate flat-field image which should fulfill the requirements of the sort of correction you desire - either lens/camera non-uniformity, or the same including the removal of dust spots. The creation of flat-field images for both purposes is described in the next section.

To illustrate flat-field correction, a flat-field image has been applied unto itself. There is a noticeable light falloff asymmetry shown in the "before" image at the top, as well as a green color-cast. Flat-field correction removes both color and luminance non-uniformities and results in a perfectly uniform image. The histogram of the bottom "after" image indicates that the corrected image does not have any variation in the tones - exactly what one would expect from a uniform (flat) field. The same level of correction is applied to the "real" image when it is flat-field corrected.

All parameters in the "Flat-Field" panel are saved to the processing profile. These settings can be copied and pasted to other images just as any other setting. This includes the flat-field "Auto-selection" option. Pasting it to different images will result in independent auto-selection of the appropriate flat-field for each of them.



Demonstration of the effects of flat-field removal from the flat-field photo itself.

Creating and Using Flat-Field Images for the Correction of Camera/Lens Aberrations

It is sufficient to shoot a single flat field image using the same camera body and lens setup as the images you want to apply it to - you do not need several flat-field images with identical settings, unlike dark-frame shots where in order to get higher quality you do need to take several shots with identical settings and merge them. You do not need to merge flat-field shots since merging serves to increase the signal:noise ratio, and this is unnecessary as the flat-field image gets blurred and should have a low ISO. Shoot it at ISO-100 if possible. The white balance setting is irrelevant. The focus setting of the lens while taking the flat-field shot for camera/lens aberration correction does not need to match that of the image you want to apply it to (this is not true if you take flat-field shots for dust removal, as explained in that section). Keep the lens de-focused while taking the flat-field shot. Your flat-field image should be a de-focused photograph of a uniformly lit, uniformly colored and toned subject.

The most optimal method of shooting such a flat-field image is to use a uniform piece of milky, semi-opaque Poly(methyl methacrylate) (https://en.wikipedia.org/wiki/Poly(methyl_methacrylate)) (PMMA, also sold under the names Acrylite, Lucite, Perspex and Plexiglas). It should be plain, colorless, smooth, with no symbols or writing etched or embossed (https://en.wikipedia.org/wiki/Paper_embossing) into it. You can find sheets at hardware and art stores, or you may have luck finding an empty container made of such material - Nordsjö Vävlim (http://www.nordsjoidedesign.se/se/PRODUKTER/Ovrigt/Tapettillbehor-1261.html) containers work well. Cut out a shape the size of your lens barrel. If you use a lens with an extremely wide angle of view, such as Samyang 8mm fisheye lens, then the size will need to be several times larger than the lens barrel - at least 20x20cm. The thickness of the PMMA sheet should be at least 1mm, preferably 3mm or more, to sufficiently blur and distribute the light shining through it. If it is glossy, use fine sandpaper to make the surface matte. This is now your flat-field filter.

Hold your flat-field filter in front of the lens (at full contact with the lens/filter barrel), click the auto-exposure button, and photograph it against a uniform illumination (e.g. point the camera towards a clear sky). The idea is to provide a uniform entrance illumination into the lens. As a result, all non-uniformities of the lens/camera combination can be recognized in the captured raw image as deviations from the ideal spatially uniform (flat-field) response. Take one shot for every major aperture value (f/1.4, f/2, f/2.8, f/4, f/5.6, f/8, f/11, etc.). Remember to press the auto-exposure button each time to keep the histogram centered without clipping.

If you are using a zoom lens instead of a prime, you may need to check whether you need to take a series of shots at major zoom intervals. For example if you have an 18-55mm lens, you may need to take one series of shots at 18mm, one at 36mm and one at 55mm. Test this for yourself with your own gear, as maybe the 36mm shot will be almost identical to the 55mm one and so you could skip the 55mm series of shots.

If you use an electronic lens then the filename is irrelevant as RawTherapee can automatically find the correct flat-field image, however if you use an old manual lens which does not record the aperture or focal length, then name your files as follows:



This will make finding them manually easy.

Apply the flat-field images to your photos, matching the aperture values of the flat-field images with those of your photos, and keep the "Blur Radius" slider at 32 or above.

If you haven't yet found a PMMA sheet and you need a flat-field image ASAP, you can take a completely de-focused photo of a plain blank wall or a large sheet of matte paper. The difficulty here will be getting the light uniform - remember that the inverse square law (https://en.wikipedia.org/wiki/Inverse-square_law) is playing against you. The light source should be far away and diffused, as you will not get a good flat-field image using a room light or a diffused strobe. Do not replace the PMMA sheet for a piece of paper, as the paper is far too irregular, and all the variations in thickness and density will show up as lighter/darker regions which will get subtracted from the photo you apply it to. If you want to use paper then get a large blank sheet, place it in an evenly illuminated place such as on a wall opposite a large window or outdoors on the ground under the sky. Zoom into it, de-focused, and take the shots the same way as described above.

Creating and Using Flat-Field Images for Dust Spot Removal

Carefully follow the instructions from above, with the changes and additions described here.

If you intend to use the Flat Field tool for dust spot removal, it is advisable that you take a flat field shot not too long before or after you take the actual photos you want to apply it to, because as time goes by new dust spots on your sensor and lens can appear and old ones can shift or disappear, and if that happens perfect dust spot removal becomes impossible.

The main difference between using a flat-field image for the correction of camera/lens aberrations and for dust removal is that the image will not get blurred inside RawTherapee, and this requires a few changes. Using the flat-field filter described above will make taking these shots significantly easier, faster and more accurate. In addition to shooting one flat-field image at every aperture value, you will also need to take into account the focus setting of the lens. Trying to apply a flat-field image for dust removal taken with the focus setting at 0.5m to a photo taken with the focus at infinity will most likely fail because changing focus zooms the image slightly, even on prime lenses. Luckily, the zoom displacement is very small, so if for every aperture you take one shot at the closes focus (e.g. 0.5m), one at about 2m and one at infinity, this will surely suffice.

The same requirements for zoom lenses apply as described above, as does the file-naming procedure. Additionally, you will want to include the focus setting in your filename:

ff_<lens>_<YYYYMMDD>_<focallength>_<aperture>_<focusdistance>.<raw>

for example:

ff_20141009_pentax18-55mm_36mm_f11_2m.dng

Now you see why having a small PMMA flat-field filter is so useful - instead of shooting whole series of shots at home, you can just take the needed shot without changing any camera settings right after you take your actual photo, standing on the mountain top at sunrise or crouching shooting fungus in the woods.

When taking certain kinds of shots, for example macro photos, I tend to stick to f/11, as for my lens this is the optimal trade-off between depth of field and sharpness. If I make the aperture any smaller, diffraction (https://en.wikipedia.org/wiki/F-number#Effects_on_image_sharpness) will have a negative effect on the perceived

sharpness. Sticking to f/11 for my macros also means that I only need to take one series of flat-field shots. The series would consist of three shots: one at f/11 with the focus at minimum, one with the focus midway, and one at infinity. Having a handy flat-field filter makes that very easy.

When to Take Flat-Field Shots

As described above, make a library containing all the required series of shots for your lens(es). You will want to update this library when you get a new lens or a new camera body. Additionally, if you use flat-field shots for dust spot removal, you will want to take new shots after you clean your camera sensor, and when the dust and spot layout changes, i.e. new dust spots appear or old ones disappear.

Algorithm Specifics & Concise Summary

The user or auto-selected raw flat-field image need not have the same white balance as the image to which it is applied. The flat-field is blurred according to any of several user-selectable blur types and choice of blur radius. The blurred flat-field serves as a template for vignetting, lens cast, etc. and be used to correct the corresponding issues in the image file. "Area" blur uses a box blur (https://en.wikipedia.org/wiki/Box_blur) of the flat-field file and is the normal use of the correction. Use a large blur radius to smooth the flat-field image in order to mitigate any imperfections such as flat-field image noise, dust spots, etc. Use of a small radius will leave these effects in the



The Flat Field tool can be used to completely and automatically eliminate dust spots from hundreds of photos with a single click.

flat-field image, and will leave their imprint on the corrected image file. This can be used to advantage; for instance if the flat-field has the same dust spots as the image file, use of a blur radius of o-1 pixels will subtract the darkening caused by the dust, therefore eliminating dust spots from the image. Use the "vertical", "horizontal" or "vertical+horizontal" blur options if the camera generates repeatable line pattern noise, such as repeatable vertical lines.

Organising Flat-Fields

Non-uniformity of the captured field depends on the following parameters:

- Camera (camera & sensor combination in case when a digital back (https://en.wikipedia.org/wiki/Digital_back) is used),
- Lens,
- Focal distance,
- Aperture,
- Lens tilt/shift



Thumbnails of sample flat-fields displayed in RawTherapee's File Browser.

It is recommended to assemble a library of flat-fields for camera/lens combinations, taken at various aperture settings (ones which you use in your real photography). It is advisable to name flat-fields descriptively, so that files can be recognized easier by the user, preferably incorporating all parameters listed above. During the flat-field correction process these parameters are read from the Exif data only and the filename is irrelevant from this perspective. The flat-fields should be stored in a dedicated directory. RawTherapee allows to point to it via "Preferences > Image Processing > Flat Field > Flat-fields directory". Setting that directory initiates an analysis of its content, and a count of qualifying files and templates is reported in the UI once the scan is complete (this could take a short while depending on the number of flat-field images you have).

Flat-Field File Browser Context Menu Options

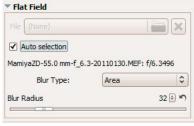
You can apply and manage flat-field images from within the File Browser tab, by right-clicking on a thumbnail and selecting the "Flat-field" option. You will be presented with three sub-options:

- "Select flat-field" brings up the file selection dialog to select a flat-field file to be applied to the selected images.
- "Auto flat-field" allows to run the "Auto-select" option on the currently selected images.
- "Move to flat-fields directory" moves the selected image into the directory specified in Preferences.



Flat-field actions in the File Browser tab.

Auta Calaction



Flat-field auto-selection capability can be engaged simply by checking the "Auto-selection" checkbox. RawTherapee will search through the files in the "Flat-fields directory" specified in preferences and select the exact or, if not available, the closest match to the image being corrected based on camera make, model, lens, focal length, aperture and date of the flat-field file. If a match is found, the filename of the selected flat-field image will be displayed along with its aperture setting. If a match is not found, flat-field

correction will not be applied and a message will be displayed to the user. If more than one exactly matching flat-fields are found, their data will be averaged and then used for flat-field correction.

Auto-selection does not account for the tilt-shift settings used on the lens, therefore such flat-fields should not be stored in the main flat-fields directory, but rather in a descriptively named sub-directory. Such unusual flat-field files should be applied manually.

Auto matching logic

Key for flat-fields (ffInfo::key):

- camera manufacturer
- camera model
- lens
- focal length
- aperture

The search for the best match is twofold:

- if perfect matches by key are found, then the list is scanned for lesser distance in time,
- otherwise if no match is found, the whole list is searched for lesser distance in lens and aperture.

Blur Type

Area

The default and generally most useful setting to apply blur action equally in all directions. Works well for correcting vignetting and lens cast.

Vertical

Blurs the flat-field vertically to compensate for vertical non-uniformities. This is useful if the vertical sensor readout has variation between columns.

Horizontal

Blurs the flat-field horizontally to compensate for horizontal non-uniformities. This is useful if the horizontal sensor readout has variation between rows.

Vertical + Horizontal

Blurs the flat-field sequentially horizontally and then vertically to compensate for both vertical and horizontal non-uniformities.

Note that the concept of vertical and horizontal is related to how the sensor is oriented in the raw file, which is always the same regardless if the camera was held horizontally or vertically when shooting. It varies between camera models if the sensor data is stored in landscape or portrait format so if you want to use Vertical or Horizontal modes you need to test which direction that is right for your model.

Blur Radius

The "Blur Radius" slider controls the degree of blurring of the flat-field data. The default value of 32 is usually sufficient to get rid of localized variations of raw data due to noise. Setting the blur radius to 0 skips the blurring process and allows to correct for dust and other debris on the sensor (as long as their position in the flat-field image is the same as in the photo the flat-field is applied to) at the expense of carrying noise from the flat-field file into the corrected image. If such correction is desired, it is advisable to create flat-field files with minimum amount of noise at the lowest ISO setting and optimal light exposure.

Clip Control

Applying a flat-field image can cause nearly-overexposed areas in the image to become overexposed due to the correction. Activating the Clip Control option will keep the flat-field image from clipping the actual image. Areas in the actual image which were already clipped before the application of the flat-field image may acquire a color cast, therefore as a rule of thumb, if your photo contains overexposed areas it is better to not use Clip Control.

To understand exactly how Clip Control works we need to get a bit technical. The Flat-Field tool works by adjusting the exposure of areas of the actual image whose corresponding areas in the flat-field image are different to the measured exposure of the center of the flat-field image. The factor by which the exposure of an area of the actual image is increased is proportional to how much darker the corresponding area in the flat-field image is relative to the measured exposure of the center of the flat-field image. When Clip Control is disabled or at o, pixels of the actual image can have their exposure increased beyond the raw white level. Clip Control works by calculating the factor against the raw white level, so that no corrected pixel will exceed the value of the white level.

For instance, if the required exposure increase of a pixel in the actual image would be 1.25 that of the raw white level (1), then the limiting factor would be 1/1.25 = 0.8. This factor would then be used in the final calculation to prevent any pixel's value from exceeding the white level. The formula for the slider is clipControlGui = (1 - limitFactor) * 100

Therefore with a limiting factor of 0.8, the value of the slider would be 20.

Metadata Tab

Exif Tab

The "Metadata" tab lets you control which Exif (https://en.wikipedia.org/wiki/Exif) metadata will be contained in the saved (developed) image file. The Exif metadata is usually created by the camera itself and implemented into the raw image file. Basic Exif information is directly visible. Extended Exif information and so-called makernote data (https://en.wikipedia.org/wiki/Exchangeable_image_file_format#MakerNote_data) is organized into a tree. Click on the arrow at the very left of the desired sub-tree and you'll see its contents.

You can "Remove", "Keep", or "Add/Edit" Exif metadata. Manipulating metadata does not change the source file in any way! If you want to restore a value you have changed or removed by accident, simply press "Reset". "Reset All" works similarly but is used for trees and works recursively, which means that all values changed/removed in this sub-tree are restored.

You can "Add/Edit" the following Exif information:

- Artist
- Copyright
- ImageDescription
- Exif.UserComment

Only the English names of the Exif fields are displayed for easy reference. They are not translated when you choose a different GUI language.

Technical Details

Reading Exif. UserComment:

- ASCII values are read without problems.
- For Unicode values, RT detects if the string has an odd length.
 - If so, it's assumed to be a UTF-8 string.
 - If not, it tries to detect if a UCS-2 byte order mark (BOM) is set.
 - If no BOM is provided, and the string is a valid UTF-8 string, it tries to auto-detect the byte order based on the "zeros" count.
 - If it is not a UTF-8 string, it is assumed to be a UCS-2 string with the platform's byte order.
- For undefined values, RT tries to convert the string from the local charset to UTF-8.

Saving Exif. UserComment:

- If the string is pure ASCII, it is saved as ASCII.
- Otherwise the string is converted to UTF-16 and saved without a byte order mark, using the metadata's byte order.

IPTC Tab

IPTC (https://en.wikipedia.org/wiki/IPTC_Information_Interchange_Model) values are metadata, as they are stored in the image files but do not affect the actual image. Basically the metadata summarized as IPTC contains additional information about your image. As this information is saved within the image file it cannot get lost. This eases the workflow a lot as you don't have to care about other files when backing up or sorting your images.

IPTC is usually used to describe the image in detail. There are a lot image database software that use the (IPTC) information saved in images to e.g. fill their descriptive fields. For example you can also use IPTC fields when you try to sell your images. Most online image-selling companies support IPTC tags and read them when you upload your images to their databases, thus you have less work. Adding keywords to your images on your home computer once while developing the photo is much more comfortable and efficient than doing so through the web browser every time you share an image. Multiple "Keywords" and "Suppl. Categories" (Supplemental Categories) can be added/removed using the plus + and minus – signs next to them.

There are three buttons:

- "S Reset" resets the IPTC values to those saved in your current processing profile.
- "\[\]" copies your current IPTC setting to the clipboard. This is especially useful when you you want to apply the same IPTC values to multiple images.
- "a" pastes the formerly copied IPTC settings from the clipboard to your current image.

Technical Details

IPTC strings are handled by libiptcdata. When saving strings, RT sets IPTC_TAG_CHARACTER_SET to UTF-8.

Retrieved from "http://rawpedia.rawtherapee.com/index.php?title=RawPedia_Book&oldid=3595"

- This page was last edited on 3 February 2018, at 00:41.
 Content is available under Creative Commons Attribution Share Alike unless otherwise noted.