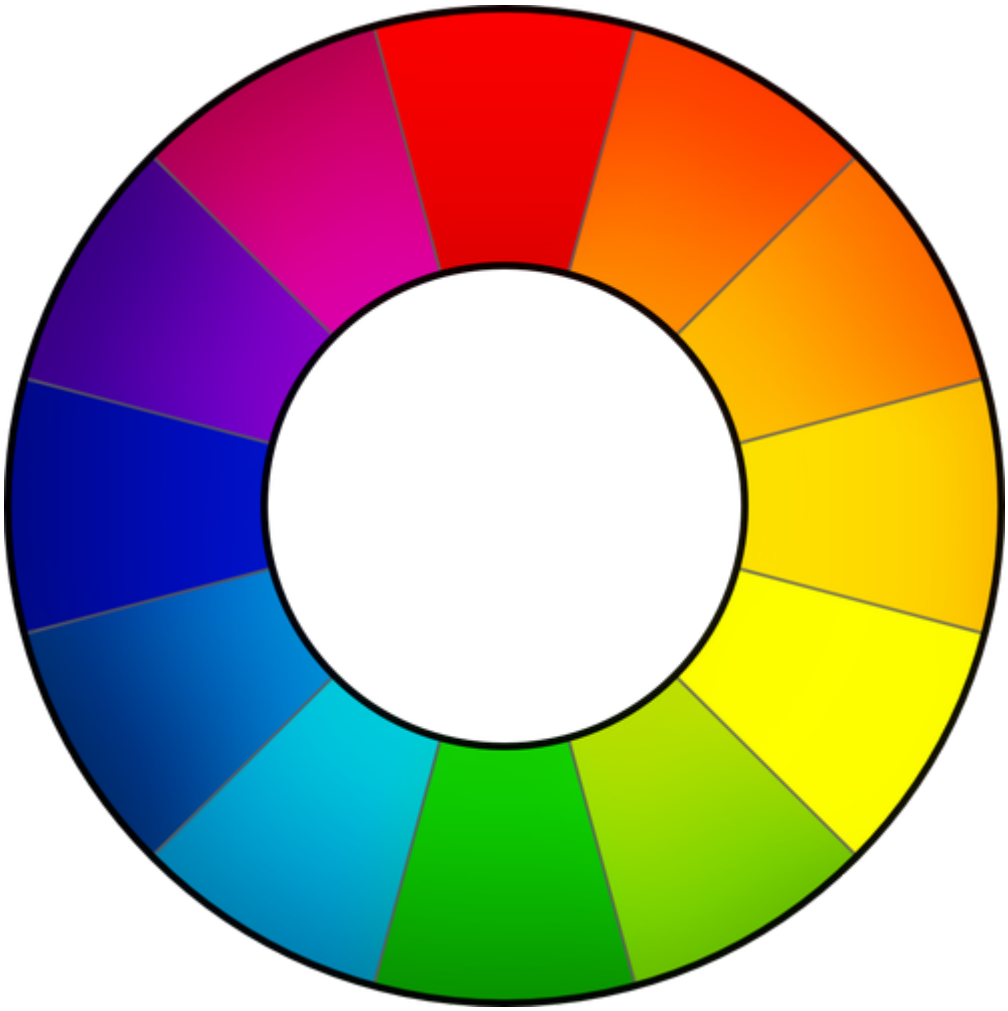

RawTherapee 4.0.9 User Manual

This is a live manual with frequent updates.
Downloaded from Google Docs on 2012-06-03
https://docs.google.com/document/d/1DHLb_6xNQsEInxiuU8pz1-sWNinnj09bpBUA4_Vl8w8/edit

If you would like to help translate this manual, read this forum introductory thread:
http://www.rawtherapee.com/forum/viewtopic.php?f=4&t=2943#p21647

Table of Contents

Introduction

Welcome!

Welcome to RawTherapee 4.0, the powerful 64-bit open source raw converter for Windows, MacOS and Linux! The RawTherapee project has been started in 2004 by the Hungarian programmer Gábor Horváth. In January 2010 Gábor decided to open his source code under the GNU General Public License, which resulted in talented developers from around the world joining the project. Based on their hard work, we are proud to present RawTherapee 4.0! We hope you'll enjoy it.

New Features in **RawTherapee 4**

- Floating point engine,
- Better color management - A new color management system and more precise handling, monitor profile auto-detection (Windows),
- Much enhanced file browser - Colored tagging, searching, metadata filtering, etc.
- New image tools - Enhanced auto-levels, black level compensation, output gamma, texture and border sharpening, post demosaic artifact & noise reduction, automatic distortion correction, vibrance, tone mapping using edge-preserving decomposition, RGB curves, etc.
- Better highlight recovery - A new “blend” method, a redesigned “color propagation” method, and more headroom with floating point operations,
- Flexible UI - You can hide panes and disable and move sections to adapt them to your needs and screen size, new theme-based icons,
- Enhanced before/after mode, new panning acceleration, better zooming and scrolling,
- Export Panel with Fast Export Options,
- Partial profile handling, with Partial Copy/Paste/Load/Save, usable in the command line too,
- Improved edit window - Grayscaled and precise clipping indicators, more color probings, raw histogram with linear & logarithmic mode, curves auto-open, RGB-L preview modes, Focus Mask, etc.
- New white balance presets for specific spectra,
- Support for Adobe DCP colorpofiles, including RT-specific camera profiles
- Support for Adobe LCP lens correction profiles
- Enhanced color matrices, new automatic input ICC profile detection, new Blend CMS with matrix colors,
- Additional crop ratios,
- Additional output ICC profiles,
- Additional keyboard shortcuts,
- Additional command line options - Allow you to automate RT using scripts or call it from other programs,
- Support for more cameras,
- Sound support - Informs you when e.g. the batch queue is done,
- XMP copy-unchanged - Keep all the metadata in your pre-tagged raws (e.g. DNGs)
- Stability improvements - Bugfixes and structural enhancements that did not make it into 3.0, esp. for Windows.

The Floating Point Engine

RawTherapee 4.0 is probably the only real time raw converter doing image processing in high precision 32bit floating point (in contrast to 16bit integer as used in and prior to RawTherapee 3.0).

Classical converters work with 16 bit integer numbers. A pixel channel has values ranging from 0-65535 in 16bit (to increase precision converters usually multiply the 12-14 bit camera values to fill the 16bit 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 also 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.

The downside is the RAM space floating point numbers require, which is exactly twice that of 16 bit integer. To cope with these increased requirements, a 64bit operating system is highly advised for stability. If you have problems running RT on a 32bit 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 the file browser tab while you are processing photos.
- Use 4-Gigabyte Tuning in Windows. See the page in the Microsoft Library for an explanation of what 4-Gigabyte Tuning is:
  http://msdn.microsoft.com/en-us/library/bb613473%28VS.85%29.aspx
  and find out how to do it in Windows XP, Vista and 7 by reading this guide:
  http://avatechsupport.blogspot.co.uk/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.
- The most memory-intensive tools are Tonemapping, High Quality Shadows & Highlights and Contrast by Detail Levels, so you might need to avoid using them if your machine and operating system are up to standard.

Using RawTherapee

First Run

The first time you start RawTherapee 4.0, you will most likely see an empty file browser tab. This is because you need to tell RawTherapee 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 it. Now RawTherapee will generate thumbnails for your photos in the central panel. The first time you open a directory full of raw photo files, RawTherapee will read each file and extract the embedded JPEG thumbnail from it (every raw photo has an embedded JPEG thumbnail, sometimes even a few of various sizes) to quickly present you with a view of all your photos. It will then proceed to re-read all those raw files, this time generating high quality thumbnails based on the adjustments you've made to each (if any), and it will replace the inaccurate JPEG thumbnails with these high quality ones. This can take some time, but it only happens the first time you read a directory. All subsequent times RT will just grab the previously generated matching high quality thumbnail from a local cache, which is almost instant.

Use the zoom icons on top of the file browser to make the thumbnails smaller or bigger. Each thumbnail uses RAM, so it is advisable to keep the thumbnail size low (Preferences > File Browser > Maximal Thumbnail Height).

You can filter the thumbnails you want to be visible using the toolbar buttons, as well as quickly rotating a whole selection of photos, and other adjustments.

Processing your first raw

Now let's get to work! Hover your mouse over a thumbnail and right-click on it. You'll see a number of options; forget about them for the moment and choose Open. By default, a new tab appears, showing your photo. In fact, by default, every time you open a raw it will be displayed in a new tab.

Single Tab Mode

Not everybody likes having each photo open in a separate tab. RawTherapee can also be used in single tab mode, where the thumbnails are shown in a film strip at the top of the window and where every raw opens in the same Image Editor tab. To activate this mode, click on the 'Preferences' icon in the top-right corner of the RT window, and choose 'General > Layout' and set 'Editor Layout' to 'Single tab mode'. Use this 'Preferences' window as well to select various languages for the user interface, to choose other color themes, change the font size, etc.

The Image Editor Tab

Take a moment to explore the 'Image Editor' tab. The central panel holds a preview of your photo. It will show you the effect of all the adjustments you make. Note that some adjustments, such as sharpening and noise reduction, are visible only when you are zoomed in at 100% (1:1) or more. On

either side are panels which show information and contain tools with which you can squeeze all the juice your photos have to offer.

The Left Panel

To the left is a panel which shows the RGB, HSV and Lab values of the pixel your cursor is currently hovering over. This is called the “Navigator”. Under 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. Under the ‘History’ panel is a panel called ‘Snapshots’. It’s 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, all history items and all snapshots will be saved to the sidecar file. For now, the history and snapshots are lost when you load a new photo in the Image Editor or close RT.

The Right Panel

To the right, at the top, is the histogram. It can show the histograms of the red (R), green (G), blue (B) and CIELab Luminance (L) channels of the photo as it would look if you saved it, or it can show you the histogram of the raw file before any transformations such as demosaicing are applied to it. Additionally, you can show or hide the RGB Indicator Bar, which is situated under the histogram and shows the exact place on the histogram of the enabled R, G, B or L values of the pixel your cursor is currently hovering over.

Just under the histogram is a drop-down list called ‘Postprocessing Profiles’. This is an important concept in RawTherapee. Every time you open a raw file, a profile is applied. A profile is a set of instructions regarding which demosaic algorithm to use to interpret the raw data, whether a sharpening filter is to be applied, if auto-exposure is enabled, etc. These profiles are simple text files and it's easy to make your own. For clarification: these postprocessing profiles have nothing to do with camera profiles or color profiles. By default, RawTherapee opens raw files with the “Default” profile, but you can change that if needed in ‘Preferences > Image Processing > Default image processing parameters’.

Preview modes

Enabling the preview modes allow to alter the modes in which preview image is displayed in the Editor. Preview modes are controlled via buttons in the Editor toolbar or via shortcuts.

Preview mode	Shortcut	Toolbar Buttons
Red channel	**r**	
Blue channel	**g**	
Green channel	**b**	
Luminosity	**v**	
Focus Mask	**Shift-F**	

Regular mode	*see note	

*note: Preview is returned to the regular mode with the repeated use of any preview mode shortcut.
Only one preview mode can be engaged at a time.

RGBL preview modes

The following preview modes are currently supported:
Red, Green and Blue channels,
Luminosity, which is previewed as 0.299*R + 0.587*G + 0.114*B.

When clipping indicators are engaged in the RGBL preview modes, shadow clipped areas are indicated in blue color and highlight clipping is indicated in red. As during normal preview the lightness of the clipping highlight is indicative of the degree of clipping.

Preview of individual channels may be helpful when editing RGB curves, planing BW conversion using channel mixer, evaluating image noise. Luminosity preview is helpful to instantly view image in black and white without altering development parameters (e.g useful in aesthetic image evaluation).

Focus Mask

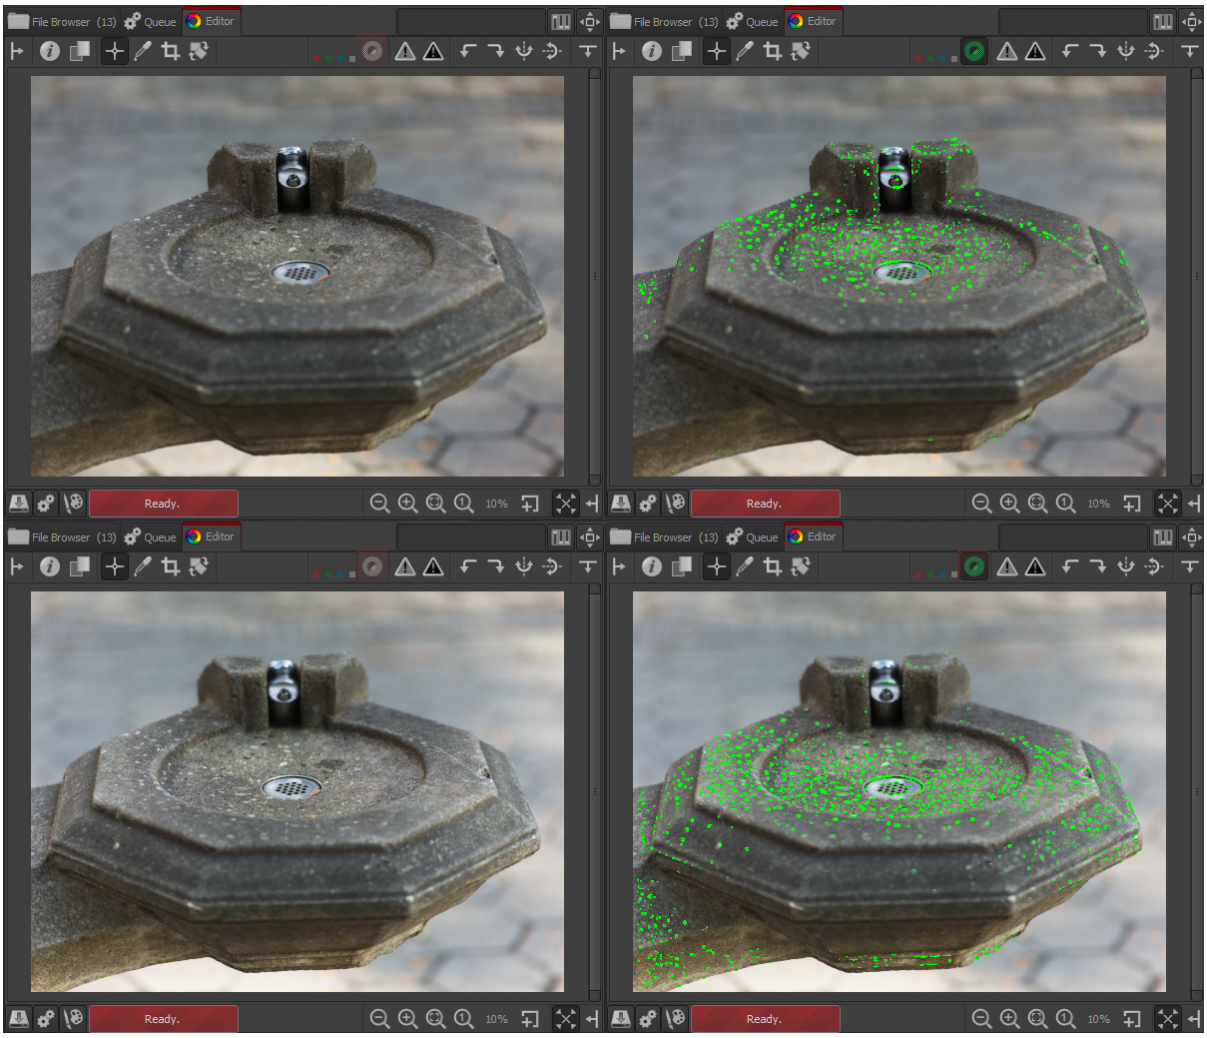

Focus mask indicating the focusing plane.

Focus Mask is designed to highlight sharper areas of the image. Naturally, focused areas are sharper, so the focused areas are being highlighted. Focus Mask is more accurate on images with shallow depth of field, low noise and at higher zoom levels.To improve detection accuracy for noisy images evaluate at smaller zoom 10-30%. Please note that preview is rendered slower with Focus Mask turned on.

Current beta implementation analyses 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 also may contain micro texture. At the same time, rescaling of the original image to the preview size compresses larger scale details into smaller size at which they may appear sharper; rescaling may introduce aliasing artifacts that could get sharp rendering and lead to false indication of the sharper feature in the full size image.This is one of the areas of future improvement in the Focus Mask algorithm.

Using current implementation of Focus Mask you may increase your confidence by viewing the mask at a couple of zoom levels. It is not always fault proof, but can be helpful in many cases; it proved to be quite accurate on a variety of test images.

Warning: please be sure to double-check your images if you decide to delete them

based on the focus mask.

Saving

The first thing one needs to know is how to process (or convert) a raw file and save it somewhere. Play around with the ‘Exposure’, ‘Highlight recovery amount’ and ‘Saturation’ sliders in the ‘Exposure’ tab. When satisfied, you have several ways of saving the file.

Click on the little hard disk icon at the bottom-left of the preview image (a 'Save current image Ctrl+S' tooltip appears when you hover your mouse over it). This works as a 'Save as' dialog, meaning you can select the name and location for the output file (RT will automatically add the extension based on the chosen format), choose between JPG, TIFF or PNG (8-bit or 16-bit for the last two), set compression ratio, etc. You can also choose to save the processing parameters (your adjustments to the three sliders) in a sidecar file next to the output image. The last option lets you choose whether you want to 'Save immediately' or ‘Put to the head/tail of the processing queue’.

Save Immediately

If you chose ‘Save immediately’, RT 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.

Try it. Click on the ‘Save current image’ icon in the ‘Image Editor’ tab (CTRL+S). Choose JPG, quality 95, click 'Save immediately' and click OK. After a couple of seconds you will find the resulting file with the JPG extension in the directory you specified. Use an external image browser to check the result. Although you can use RawTherapee for that, a specialized image browser will likely be faster and handier. Don't forget that RawTherapee is designed as a raw converter, not as an image browser!

Put to the Head / Tail of the Processing Queue

If you chose 'Put to head of the processing queue' or ‘Put to the tail of the processing queue’, your image will be kept in a queue of files to be processed once activated, so you can keep working on photos and adding them to this queue without having to suffer through a busy and not very responsive RT as it processes those photos. Then, once you are done tweaking a whole bunch of photos and adding each one to the queue, you can tell RT to start processing all of the ones in the queue while you go and enjoy some tea as RT grinds through them.

Whether you save the image immediately or put it to the queue using the ‘Save current image’ button, you can use the proposed output name or set a new one. RawTherapee adds the extension automatically based on the chosen output format.

Note that there is a shortcut for adding an image to the tail of the processing queue: at the bottom-left of the preview image, next to the ‘Save current image’ button, there is a ‘Put current image to processing queue’ button. It does the same things as

clicking on ‘Save current image’ and then selecting ‘Put to the end of the processing queue’ and clicking ‘OK’, except that it will use the file format settings as specified in the ‘Queue’ tab (read about them in the next section).

The Batch Queue

Open a photo for editing, tweak it, click on ‘Save current image’, add it to the tail of the processing queue and click OK. Go to the ‘Queue’ tab. You will see your photo there, waiting to be processed.

The ‘File format’ panel resides in the top-right side of the ‘Queue’ tab. You can save to JPG (8 bits per channel), TIFF (8 or 16 bits per channel) and to PNG (also 8 or 16 bits per channel). You can also select ‘Save processing parameters with image’ - this options writes a sidecar file with all your adjustments made to that photo in a plain text file. This file will have the same filename as your photo, but it will have a “.xml” extension.

You can set where you want the resulting JPG, PNG or TIFF image saved to by entering an appropriate template in the ‘Use Template’ field in the ‘Output Directory’ panel. To find out how to create a template, 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, %d1, %d2, ..., %p1, %p2, …, %r

These formatting strings refer to the different parts of the photo's pathname or some attributes of the photo. 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 unranked, *%r* will be replaced by '*0*'. If the photo is in the trash bin, *%r* will be replaced by '*x*'

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*" located in the directory of the opened image, write:

%p1/converted/%f

If you want to save the output image in a directory named "/home/tom/photos/converted/2010-10-31",

write:

%p2/converted/%d1/%f

Alternatively, you can save directly to a specific directory, but in the long run it is much easier to use a template.

On the left you see a ‘Start/Stop processing’ button, and an ‘Auto start’ checkbox.
If ‘Auto start’ is enabled, every time a raw is sent to the queue, processing will start immediately. Usually you will not want this, as this will use up your CPU on developing the photos in the queue, and as a result all adjustments you do while the queue is running will take much longer to get applied so that you can see their effect - RT will become sluggish.
If Auto start is unchecked, you will have to activate the queue manually by clicking the ‘Start processing’ button.
You can pause the queue by pressing the ‘Stop processing’ button, but RT will first finish processing the current photo.

You can delete the content of the processing queue by right-clicking on a thumbnail and choosing ‘Select all’, then ‘Cancel job’.

You can exit the program and restart it later; the batch queue will still be there. The queue can even survive a crash of RawTherapee, as the batch queue info is written to disk each time you add a photo to it, each time a photo is done processing and each time you delete a photo from it.

8-bit and 16-bit

“8-bit” when referring to image formats typically means that the program assigns 8 bits (8 digits, or a total of 255 possible decimal values) to each pixel’s color channel, and each pixel in the files RawTherapee saves has three color channels - red, green and blue.

Most, if not all, modern raw-capable 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. 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 postprocessed using a state of the art raw postprocessing program, such as your’s truly - RawTherapee.

Once you have postprocessed a photo in RawTherapee, you’re 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 and PNG, not JPEG). If you plan to post-process your photos after RawTherapee in a 16-bit-capable image editing program, it is better to save it in a non-destructive 16-bit format. Uncompressed TIFF is suggested as an intermediate format, as it is quick to save and stores all the metadata (EXIF, IPTC, XMP) 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 don’t actually need to read this to use RT, it’s 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 files 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 image formats capable of storing only black-

and-white images, such as a variant of TIFF designed for fax machines. 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 get’s 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 also 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 bits floating point per color channel, such as the EXR format, or 32 bits, such as the RGBE format.
To summarize, an “8 bit per channel” image can also be called a “24 bit per pixel” image. Use the former.
A “16 bit per channel” image can also be called a “48 bit per pixel” image. Use the former.

Naming

If your original raw file was called DSC_1000.RAW, the processed file name will be DSC_1000.jpg (or .tif or .png). There's 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 DSC_1000.jpg, DSC_1000-1.jpg, DSC_1000-2.jpg, etc. The same applies when you send different versions of the same raw to the batch queue. By the way, your original raw will never be altered by RawTherapee.

Sidecar files - postprocessing profiles

If you enabled ‘Save processing parameters with image’, you will notice that there is a second file next to your output file. It has the same file name as the output file, but it has the .xml file extension. This is a little text file that describes exactly what edits you made to your photo. Open it in a text editor and take a look. Whenever you re-open that raw in RawTherapee, this file will be read and all the settings from the last session will be restored (the profile displayed under Postprocessing Profiles will change to 'Last saved'), so you'll never have to wonder what the sharpening settings were for a certain photo that you edited two weeks ago, because everything is recorded in this sidecar file (to be more precise: nearly everything, as ranking information and the history panel contents are not stored in this file yet). By default this information is also saved to a cache on your hard disk (here the ranking info is available). In ‘Preferences > Image processing’, you can disable the generating of these sidecar files. But why would you? They represent all the time-consuming work you did on a photo, so keep your work! You might even want to backup these files together with the processed photos. In case you want to start a new edit session from scratch, simply open the raw photo and select the postprocessing profile you want to start from; ‘neutral’ will show you the most basic version of your photo without any modifications.

Creating postprocessing profiles for general use

In RawTherapee, we call the sidecar files ‘postprocessing profiles’. We supply a bunch of postprocessing 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 BW-2 - it will start off by turning your photo into a somewhat high contrast black and white version, which you can then fine-tune.

You can see the whole list of postprocessing profiles in the ‘Image Editor’ tab, under the histogram, if you expand the ‘Postprocessing 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 ‘Profile Operations > Apply Profile’.

You can create your own postprocessing profiles for general use. Open a photo you want to create a good starting point profile for. Let’s imagine this is a photo of a baby portrait. I suggest you start off with the ‘neutral’ profile, and work your way up from there. When you are done, just click on the ‘Save Current Profile’ icon in the ‘Postprocessing Profiles’ panel. Enter any name; you don’t need to specify the extension - RT will add it for you. But sometimes, you’ll 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 on the Control key while clicking on the ‘Save’ button. When you’ll have selected the output file name, a window will let you choose which parameters to save. 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 profiles 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 studio, as the little bugger 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’ with an appropriate ‘Clip’ value (more on this tool later on). 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 postprocessing profile is meant to make baby faces profile photos look soft and cuddly by a clever mixture of highlight recovery, auto exposure, RGB tone curve, Lab tone curves, Lab brightness and Lab contrast, then don’t enable noise reduction, don’t set custom white balance, 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.

Closing RawTherapee

To close RawTherapee, click on the little close icon on top of the screen. It depends on your window manager if this icon is located on the right or the left of your screen and if it's colored or just shows a 'x'. There's no 'official' exit button.

Send to editor

You can use the third icon at the bottom-left of the preview panel (the paint brush and paint board) to process your raw in RawTherapee and send it straight to an image editor such as Gimp or Photoshop for further processing. Note that RawTherapee produces 16-bit temporary intermediate output files when using this ‘Edit current image in external exitor’ button, but Gimp cannot yet handle 16-bit images, so it will down-sample them to 8-bit ones.

You can specify your external editor of choice in ‘Preferences > General > External editor’.

The Toolbox

What are all those sliders in the toolbox on the right side for and what do they do exactly, you may ask yourself? Well, 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's 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.

General comment about some widgets

Panels

A panel is a foldable element and is also frequently called a 'Tool', or 'Section'. Right-clicking over a panel's title opens it and closes all other panels of the same tab. You'll learn to love this shortcut when you'll consider the time you'll spend on folding sections manually...

Sliders

Each slider has three values in memory:

1. the current value, when you move the slider to any position,
2. the 'default' value, the one that the programmer has set as default. It can be called back by clicking on the 'Reset' button,
3. 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.

Warning: when using the mouse scroll wheel over a tab to scroll its content, you'll rapidly see a very annoying behavior: when a slider (or combo box button) comes under the cursor while scrolling the tab content, the slider get the focus automatically and the scroll wheel changes the slider's value!

Curve editors

Each 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).

The preview area

The preview is designed to show you the more realistic result possible but it still has to deal with the result/performance balance. That's why some tools will only be computed and previewed at the 1:1 scale only. That's the case for e.g. the Sharpen tool: it doesn't really make sense to show such fine details when you have a zoomed preview image.

Exposure tab

Exposure panel

Auto Levels

Open a raw photo and be sure the postprocessing profile under the histogram shows 'default'. If not, click on that dropdown box and choose default. As you can see, the Auto Levels button seems to be clicked on already and Exp. Comp. (exposure compensation) will most likely have a value other than zero. It is our default profile that instructed RawTherapee to activate these auto-levels. If you don't want this, click on the two-arrow icon on the right of the exposure compensation slider (it says 'Reset to default' when hovering the mouse over) and the value will most likely get lower but probably not zero.
This means that even without auto exposure, RawTherapee still applies some correction when opening a raw. If you want all the values in the Exposure section of the Exposure tab to be zero, then you must choose the neutral profile in the dropdown menu. More experienced users will use the neutral profile more often as a starting point, for starters or occasional users the default profile is recommended.
Auto Levels works well as a good starting point for your edits, but may fail on special images containing much noise or are very dark or light.

Clip

Auto levels uses the clip value to adjust the exposure. This number defines the amount of highlight pixels that may be clipped at the white point and the amount of shadow pixels that may be clipped at the black point. Default value is 0,0010, maximum value is 0,9999. Higher values increase the contrast.

Exposure Compensation

The values of the exposure compensation slider are ISO values. This means a value of +1.00 equals one stop of overexposure (+1 EV, exposure value; also known as +1 LV, light value)). If you make two photos, one without correction (so EV = 0) and one underexposed by one stop (EV = -1), you can get exact the same output in RawTherapee by setting exposure compensation for the second photo to +1.00. Also 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). This effect is best seen by watching the Luminance channel only, so click on the R, G and B icons next to the histogram to hide these.

Highlight recovery amount

The Highlight recovery amount slider (HR amount) can be used to compress the highlights in a photo, useful for 'dimming' (or burning) slightly overexposed areas. To see if your photo contains overexposed areas, click on the white triangle icon on the top right of the image window. Overexposed areas will show as black spots. (In case you exposed right, just add 1 EV to the Exposure compensation slider in order to see how Highlight recovery works). By dragging the HR amount slider to the right, the intensity of the highlights will decrease and the number of black spots indicating overexposure, will decrease accordingly. A value of 100 is the maximum

compression here. Note that the HR amount slider only works when Exposure compensation is greater than zero.
When increasing highlight recovery amount please also enable one of the highlight recovery methods (“Blend”, “Color Propagation” etc.).

Highlight recovery threshold

The Highlight recovery threshold slider sets the point where the HR Amount slider starts implementing compression. A value of 0 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 recovered when this slider is set to 0.

Note: as we'll see later, RawTherapee offers more ways to deal with blown highlights. 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. So while occasionally very useful, one should not expect miracles here. To create the best possible output, feed RawTherapee with the best possible input to start with - so expose well on forehand!

Black

Next is the Black slider, use this to set the black point. See the left side of the histogram move when you touch the slider. Values greater than 0 will make the image darker, negative values will lighten up the shadow parts of the photo.

Shadow recovery

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

Brightness

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

Contrast

This slider increases or reduces the contrast of the photo. Technical explanation: 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 less or more saturated. In more technical terms: it adjusts the saturation of the image by applying a multiplier to the Saturation level of pixels in HSV color space. Set this slider to -100 to create a black and white image.

Note: the above three curves are also available under LAB Adjustments, which can be found at the bottom of the Exposure tab. Although they do about the same, they give different results. More on these differences in the LAB section below.

Tone Curve

Here you can construct your own tone curve. It works on all three R, G and B channels at the same time (so you can't work on the R channel only). Three curve types are available, click on the 'v' icon on the right to choose which one you want to use, and the editor will open itself.

Linear curve

Represents the unaltered (or linear) image, so without any tone curve applied.

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. Press and hold the Control key to slow down the movement. Hold the Shift key to snap the point to key elements (displayed in red). Dragging the point down makes the image darker, while pushing it up makes it brighter. The dotted line you see now as well, marks the linear or unaltered state of the photo. Delete a point on the curve by dragging it out of the editor area.

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

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

Change the orientation of the line from left under and right up to left up and right under to produce a negative image.

S-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 you like the effect, you can save the curve for later use. Click on the hard disk icon next to the graph and give it a name. Use the Open icon to apply this curve later to another opened raw 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 use as many control point as you like.

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 highlight 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. Don't think on forehand this is a bad idea though! 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). You cannot save a specific parametric curve.

Control cage

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 1 cm (0.5 inch) to the left or the right. Now the curve passes nearby the black point, but doesn't touch it. Another difference is that 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). 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. Delete a point on the curve by dragging it out of the editor area. You can use all three curve types if you like, but only the one that is selected in the dropdown menu will be applied to the photo.

Highlight Reconstruction

Use this tool to try to restore blown-out highlights. It attempts to restore clipped (blown-out) channels in the raw data using nearby data from unclipped channels, if present.
Please note that you need to set “Highlight recovery amount” in the exposure section to see the effect of the reconstruction.
Three different methods are available:

Luminance Recovery

When selecting 'Luminance Recovery', recovered details - if any - will be gray.

Color Propagation

Color Propagation tries to restore the color information as well. Sometimes the last option 'guesses' the wrong colors. It has been much improved since 3.0 so it may

render differently than before

CIELab
CIELab reduces the luminance channel and tries to restore colors afterwards. Note that 'Color Propagation' needs a lot of calculations and is therefore slower.

Blend
This restoration mode attempts to guess clipped color channels by filling in their values from the closest match from unclipped highlight regions nearby.

Shadows/Highlights
Use this tool to independently influence the highlights and the shadows of the image (yes, another one!). Take care with the High quality option, as this one is quite cpu-intensive.

Highlights
The Highlights slider makes the brightest parts of the image less bright without touching the darker tones. To make the effect stronger, use higher values. A slider value of 100 will turn the whites in light gray.

Tonal Width for Highlights
This slider controls the strength of the Highlights slider. Higher values give a stronger effect. A value of 100, combined with Highlights 100, turns the whites into middle gray (you probably won't want that...).

Shadows
This slider lifts the shadows and applies an effect that is sometimes called 'fill-light' (or 'fill-flash') in other software. Higher values lighten the shadow areas more.

Tonal width for Shadows
Controls the strength of the Shadows slider. A maximum value of 100 gives strongest 'lift shadows' effect.

Local Contrast
Local Contrast is adaptive contrast adjustment depending on contrast within a specified area. It increases contrast in small areas while keeping the global contrast (which could be set with the contrast sliders in Exposure or Lab). The resulting image will look more 'three-dimensional'. This feature is very useful when you have a foggy image or took your picture through a window. The effect can be very subtle. In fact it is the same as an unsharp mask with high radius and small value. For optimal results the slider should be between 5 and 20.

Radius
The higher the Radius value, the stronger the effect of Local Contrast. Radius influences the highlights and shadows as well: a higher value makes the whites whiter and undoes the effect of the Highlights slider to a certain extent. Same story when you only try to adapt the shadows.

If you are bored, set the first four sliders to 100 and play with Local contrast to

transform your favorite raw processor into a cheap effect machine!

Tone Mapping

The tone mapping option is used to adjust 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 Tone and Detail Manipulation with some modifications. Note: tone mapping takes a lot of memory.

Strength

This controls the strength of the overall effect.

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.

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.

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, 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).

Lab Adjustments

Here we see again three sliders for brightness, contrast and saturation. The difference with the earlier mentioned sliders with the same name in the Exposure section is that here the curves are applied in Lab color space, while above they were applied to the RGB color model. Simple explanation of the differences: LAB is made to approximate human vision, while RGB is used to describe the colors that were captured by the photo sensor in your camera or that are displayed on your monitor, etc. (see Wikipedia to learn more about this complex topic).

Brightness

When using the Brightness slider in the Lab section, a tone curve is applied to the L-channel (the L stands for lightness) 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 luminance level. Tonalities above the average are lifted (lowered), while tonalities below the average are lowered (lifted).

Saturation

The Lab saturation slider increases or decreases the saturation of the image, by applying a contrast curve to the a- and b-channels of Lab space. Set this slider to -100 to switch to black and white.

Avoid color clipping

Prevents clipping of very saturated colors. Effect is best seen when saturation slider is set to +100. Has no visible effect with moderate saturation values.

Enable saturation limiter

Limits the amount of saturation. When Saturation limit is set to 0, the saturation does not increase with values greater than 0. A raw file processed with saturation +100 and the saturation limiter set to 0, produces the same result as the same file processed with saturation 0.

Curves

Here you find again the same curve types as described in the Exposure section above. Use them to manually adapt the lightness channel L and the two color channels a and b. When you are working on a black and white image (Lab saturation = -100), you can use the color channels to apply a 'toner'. Choose for example the Parametric curve in the b channel and set the Lights slider to 5, 10 or 20 to make the image 'warmer'. Use the middle control point under the histogram to fine tune the effect. If you want to apply the same tone to a couple of b&w photos, you need to copy the current profile to the clipboard via the appropriate icon at the top of the Exposure tab. Then go to the file browser, select the photos, right-click and say 'Paste profile'.

RGB versus Lab

Many people are wondering what the differences are between adjusting brightness, contrast and saturation in RGB and in Lab.

RGB calculates in the three channels red, green and blue. L*a*b is a conversion of the same information in lightness, a color and b color. The lightness is roughly derived from the eyes sensitivity, which is very sensitive to green but less to blue. So if you brighten in Lab space it 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 both brightness and contrast sliders is harder to tell, because it is subtle. A brightness setting of +30 in the Exposure section (so in RGB) produces an image that is overall a bit brighter than when using the same setting in Lab. The colors in Lab are somewhat more saturated. The contrary is the case when using an Exposure contrast of +45: now the colors are clearly warmer than 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. This applies as well when you set the saturation slider to -100 to produce a black and white image. LabSat gives a more neutral image to start with,

while ExpSat seems to apply a kind of red filter. Just open a portrait and look how the lips turn white with ExpSat.

Detail tab

Here you find the sharpening and noise reduction tools. Please note that the effects of some of the filters here are only visible at 100%. Use the detail window (click on the + sign right under) to inspect a part of the raw.

Sharpening

A big warning to start with! The Sharpening tool is computed BEFORE the Resize tool. So if you set a resize value and were planning to adjust a post-resize sharpness, it will not be possible at the moment. But if you resize your image by a factor of 0.5 for example, you might want to double your sharpening radius value. Unfortunately, sharpening cannot be previewed at scales lower than 1:1.
Sharpening in RawTherapee can be done with two methods: classical unsharp mask (USM) or with the RL Deconvolution algorithm.

Unsharp Mask

Radius

The Radius determines the width of the sharpening halo.

Amount

The Amount parameter controls the strength of the sharpening.

Threshold

The Threshold parameter is used to prevent the sharpening of noise: if a pixel 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 0 but then everything will be sharpened (even the noise). In general the quality of sharpening is best if the sharpening radius is as small as possible. For 'normal' images (not too blurry) it should be set between 0.8 and 2.

Sharpen Only Edges

If you activate Sharpen Only Edges the noise sharpening can be avoided completely. Two new sliders appear as well.

Radius

The Radius is used for the noise detection. If the noise is low, a lower radius could 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's 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 aggressive. When activated, a new slider appears.

Amount

When Amount is zero, Halo Control is off. At 100 it works at maximum, reducing the visual impact of the USM filter. Default value when activated is 85.

RL Deconvolution

The second method is RL Deconvolution, named after the makers of this algorithm, Richardson and Lucy. Here it is assumed that there is Gaussian blur (like applying a Gaussian filter) that might be produced by e.g. the lens or motion. In reality the blur might come close to Gaussian blur, but not exactly. Therefore some artifacts like halos might occur when you try to remove the Gaussian blur.

Radius and Amount

You can define the Radius of the gaussian blur you want to remove. When you set the Amount to 100 the gaussian blur will be removed completely, but as this gives a harsh result a lower setting is recommended.

Damping and Iterations

The Damping is used to avoid sharpening of noise on smooth areas. As deconvolution cannot be done perfectly at the first time several Iterations are necessary. How much is changed between each iteration is defined by the Richardson-Lucy algorithm. The more iterations are used the more perfectly the gaussian blur is removed. But with each iteration the speed decreases and the danger of halo artifacts rises. Normally you don't want to remove the gaussian blur perfectly due to personal visual taste and speed. The default settings should be fine most of the time.

Impulse Noise Reduction

Suppresses 'salt and pepper' noise. The slider adjusts the threshold that must be exceeded for the suppression to be applied.

Noise Reduction

Luminance

Amount of noise reduction applied to the L channel of Lab color space.

Chrominance

Amount of noise reduction applied to the a,b channels of Lab color space.

Gamma

Noise reduction is carried out after the image is transformed to the indicated gamma. Lower values weight the noise reduction to be stronger in shadows.

Defringe

Radius

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

Threshold

Sets a threshold for the application of defringing.

Contrast by Detail Levels

This filter decreases or increases the local contrast and thus works in two directions: either it reduces noise by blurring the image somewhat, or it enhances the sharpness impression. It uses four sliders. Slider 0 (Finest) has a pixel radius of 1, sliders 1 to 3 have a pixel radius of approximately 2, 4 and 8 pixels.

Contrast +/- and Neutral

Use the Contrast- button to move all four sliders by preset amounts to the left (noise reduction). Use the Contrast+ button to move all four sliders by preset amounts to the right (sharpening). Use the Neutral button to reset all sliders to 0. 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 this. Click twice on the Contrast- button and use USM sharpening with an amount of 80.

Threshold

The Threshold slider helps to suppress noise amplification when raised above zero.

Color tab

White Balance

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.

Auto

Automatically corrects the white balance.

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)
  These temperatures correspond to the “daylight” setting temperature used in these camera manufacturers’ cameras (they will be referred to as “camera-specific” temperatures in this manual), not to the actual daylight temperatures as calculated from the ICC or DCP profiles (these will be referred to as “ICC/DCP-specific”). When available, use the more precise ICC/DCP-specific daylight temperature and tint values.
  - Canon: 5200K
  - Leica: 5400K
  - Minolta: 5100K
  - Nikon: 5200K
  - Olympus: 5300K
  - Panasonic: 5500K
- Cloudy
  Camera-specific.
  - Canon: 6000K
  - Leica: 6400K
  - Minolta: 6500K
  - Nikon: 6000K
  - Olympus: 6000K
  - Panasonic: 6200K
- Shade
  Camera-specific.
  - Canon: 7000K
  - Leica: 7500K
  - Minolta: 7500K
  - Nikon: 8000K
  - Olympus: 7500K
  - Panasonic: 7500K

- Tungsten 2856K
- Fluorescent
    - F1 Daylight: 6430K
    - F2 Cool White: 4230K
    - F3 White: 3450K
    - F4 Warm White: 2940K
    - F5 Daylight: 6350K
    - F6 Lite White: 4150K
    - F7 D65 Daylight simulator: 6500K
    - F8 D50 simulator Sylvania F40 Design: 5020K
    - F9 Cool White Deluxe: 4150K
    - F10 Philips TL85: 5000K
    - F11 Philips TL84: 4150K
    - F12 Philips TL83: 3000K
- Lamp
    - HMI lamp studio Osram (for film, spectacles, studio, etc.): 4800K
    - GTI lamp Graphiclite & ColorMatch for Photography: 5000K
    - JudgeIII lamp D50: 5100K
    - Solux lamp: 3500K
    - Solux lamp: 4100K
    - Solux lamp near daylight (e.g. in “Musee d’Orsay”): 4700K
    - Solux lamp near daylight (e.g. in the “National Gallery”): 4400K
- LED
    - LSI Lumelex 2040 (e.g. in the “National Gallery”): 3000K
    - CRS SP12 WWMR16 (e.g. in the “National Gallery”): 3050K
- Flash
    - Daylight (Leica): 5500K
    - Daylight (Canon, Olympus, Pentax, standard): 6000K
    - Daylight (Minolta, Nikon, Panasonic, Sony): 6500K

ICC/DCP Specific Daylight White Balance Values

The daylight white balance temperature values used by your camera and commonly referred to in literature and on the internet are averages. The precise temperature and tint combination you should use for daylight photos is specific to the ICC or DCP input profile auto-matched to your camera by RawTherapee. This will give you a technically color-accurate image where whites are white, though aesthetically it may not be what you want.

MAKE	MODEL	TEMP	TINT
Canon	EOS-1D MKIII	4871	1.113
Canon	EOS 20D	4733	0.969
Canon	EOS 40D	5156	1.049
Canon	EOS 400D	4862	1.030
Canon	EOS 450D	4950	1.050
Canon	EOS 5D	4993	0.998
Canon	EOS 550D	4915	0.916
Canon	EOS 7D	5770	0.971
Canon	EOS D60	4723	1.237
Canon	G12	5821	0.994
Nikon	D200	4936	1.064
Nikon	D300	5277	1.070

Nikon	D3000	5302	1.109
Nikon	D3100	5087	0.955
Nikon	D3S	5100	0.970
Nikon	D50	5321	1.180
Nikon	D5100	5621	0.989
Nikon	D700	5000	1.100
Nikon	D7000	5398	0.986
Nikon	G10	4885	1.078
Olympus	E-1	5118	1.154
Olympus	E-P2	5000	1.060
Panasonic	DMC-G1	5550	1.000
Panasonic	DMC-G3	5800	1.050
Panasonic	DMC-GH1	6280	1.036
Panasonic	DMC-GH2	5680	1.043
Panasonic	DMC-FZ150	5950	1.020
Panasonic	DMC-FZ35/38	4840	1.100
Pentax	K10D	5420	1.094
Pentax	K200D	4465	1.129
Sony	A700	5280	1.076
Sony	A900	5258	1.042
Sony	NEX-5N	5832	0.890
Sony	A55 SLT-A55V	5254	1.081

Spot WB

When you click on the Spot WB button (shortcut: W), the cursor changes into a pipette (when it is on the photo). Click on a gray or white area to determine the correct white balance. You may do this several times on different places in the photo. 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.

Temperature and Tint

Moving the Temperature slider to the left makes the image cooler (bluish), moving it to the right makes it warmer (yellowish). Moving the Tint slider to the left makes the image purplish, moving it to the right greenish.

Vibrance

Vibrance is an intelligent saturation adjustment tuned to correlate with color sensitivity of human vision. Vibrance effect is applied with higher accuracy on a correctly white balanced image when RGB and Lab saturation sliders are set to 0.

Enabled

Vibrance effect can be enabled and disabled without altering other settings of the control

Pastel Tones

Saturation control over the pastel tones in the image

Saturated Tones

Saturation control over the saturated tones in the image

Pastel/Saturated tones threshold

This threshold setting value is used differentiate between pastel and saturation tones to control the area of influence of the saturation sliders for pastel and saturation tones.

Protect skin tones

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

Avoid color shift

When enabled exercises extra control to avoid hue shifting

Link pastel and saturated tones

When enabled Vibrance level is adjusted with a single slider equally controlling saturation for both pastel and saturated colors.

Channel Mixer

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

The effects are hard to describe in words, so try for yourself. Among the uses of the channel mixer:

a. IR and black & white photography, using the mixer to achieve desired tonal values for different parts of the scene.
b. Extending the range of White Balance beyond the 1200-12000 range of the WB tool.

In the channel mixer, for the values that are 100, enter values other than 100 and you are in effect changing the WB (ie. 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 temp/tint sliders of the usual WB tool to fine tune. In principle this allows an arbitrary range for WB.

HSV Equalizer

Use the HSV Equalizer (Hue, Saturation and Value) to change selective colors, e.g. emphasize a blue sky by making it somewhat darker or more saturated without touching other colors. 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 Lab Saturation slider on the Exposure tab to -100. The HSV Equalizer does not work when the raw file is desaturated by

setting Exposure Saturation on the Exposure tab to -100.

H(ue) button

The HSV Equalizer is of a so-called flat curve type, which needs some explanation. So open a raw file, click on the v-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. If you look (very) well, you'll notice that those lines are colored; from left to right: red, yellow, green, aqua, blue and magenta. Now click on the very left dot (cursor changes into a little hand) and move it slightly upward and downward. Result: red colors are quickly changing to green, blue and magenta when the cursor is moved up, and to pink, blue and green when moved down.

Notice the new horizontal line that appears when you start dragging a color point and see how its color changes. In the example on the right we change the input color green (vertical line) to the output color blue (horizontal line).

When you click and drag a vertical line, the very first movement will determine the kind of move: vertical or horizontal (so take care of 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. It's easy to see if a point is on its neutral value (i.e. on the middle line) because it's displayed in green color, like you can see on the screenshot on the right. Using the 'snap to' function by holding down the Shift key while moving the point will help you in that task.

Other things to know about this flat field curve. The very right vertical line equals the very left line (think circle). You can see this by dragging the red line on the left a bit to the left. Now the left point of the graph is at the same position as the very right point. 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.

Dragging all the color points to their maximal values (completely up) gives an image that is about the same as when no points have been changed (ie. Linear). Same when all sliders are set to their minimum positions (down).

Reset the Hue curve to linear (no changes) by clicking on the blue icon next to the V(alue) 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.

S(aturation) button

Use the same flat curve editor to change the saturation of the different colors or to accentuate one color. Dragging all sliders down desaturates the image to black and

white. Moving the slider of a single color up, re-introduces this color, and this color only.

V(alue) button

Here you can change the lightness value of a color, or its 'presence' (think 'volume'). Dragging the slider down makes the color somewhat darker, dragging it up lighter.

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, but my_hue, my_sat and my_val if you want to save all three channel curves.

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.

Color Management

Profile types

RawTherapee supports the following types of color profiles:

Type	Color fidelity, speed
Adobe DCP (below version 4), RawTherapee auto matched profile	Best color fidelity, not restrictions. Fast.
Adobe DCP (version 4)	Not supported, too bright (and not needed, RT does not have the problems solved with V4)
ICC - LUT based	Highlights clipped (default) or inaccurate (using Blend with Matri option). Blend with Matrix only usable if profile rendering is similar to RawTherapee internal default matrix (so e.g. not with indoor profiles). Slow.
ICC - CaptureOne	Supported like ICC LUT based, however colors are not guaranteed to be exactly like CaptureOne. Slow.
ICC - Nikon NX2 (hacked from temporary directory)	Highlights clipped and not recoverable. Colors are not guaranteed to be like NX2, but very near. Slow. See chapter below on how to obtain Nikon profiles.
ICC - Matrix based, RawTherapee Camera Default	Color fidelity reduced in comparison to the other types, but no highlight restrictions. Fast.

Input Profile

No Profile

Bypasses camera input color profile. Colors are closer to camera raw data. Useful in exceptional cases where camera records colors far outside conventional color gamuts.

Camera standard
Uses a simple color matrix by DCRAW, enhanced RawTherapee version (whichever is available based on camera model) or embedded in DNG.

Auto-matched camera-specific profile
Uses RawTherapee's camera-specific DCP input profile that is more precise than a simpler matrix (and fall back to lesser powerful ICC profiles if no DCP is available). Available for some cameras, these profiles are stored in /dcpprofiles (or fallback /iccprofiles/input) directory and are automatically retrieved based on file name matching the exact model name of the camera.
Please note that when ICC profile is used Highlight Reconstruction is not as efficient. As an mitigation, select enable option “Blend highlights with matrix”. DCP profiles are faster and more precise (escp. in the highlights), so they don’t have this restriction.

Custom
Specify a custom DCP or ICC camera input profile stored on your computer.
RT also parses ICC profiles by CaptureOne and Nikon (hacked from Nikon NX2 temporary directory). These need some pre- and postprocessing which is applied automatically based on the profile.

Preferred DCP profile
While RawTherapees own profiles are all single illuminant (standard daylight), some other DCP profiles contain calibrations for two light sources (e.g. Adobes ones, which come with tungsten and daylight). With preferred profile you can select which one you’d want to be taken if available. If the profile just contains one illuminant, this option has no effect. If the preferred profile type could not be found, the most daylight like profile is taken.

Blend highlights with matrix
Enable to recover blown highlights when using LUT based ICC profiles. Not needed for DCP profiles.
This function merges the imprecise, but recoverable highlights from the matrix profile into the precise ICC picture. So the color precision is reduced there, but the highlights are available.
Note that this only works correctly with RTs default profile, since the (RT DCRAW embedded) matrix must match the ICC profile.

Save reference image for profiling
Clicking this button saves the linear TIFF image before the input profile is applied. This file can be used for profiling i.e. creating a new camera profile.
However since this file does not contain metadata needed for correct profiling (e.g. white level), you must transfer this info manually. It’s generally more advisable to profile the RAW file directly.

Working Profile
Specify the working color space; this color space will be used for internal calculations, for instance for calculating saturation, RGB brightness/contrast and tone curve adjustments.

Output Profile

Specify the output RGB color space; this color space will be used for image output and embedded in metadata.
Use “No ICM” (Integrated Color Management) if you aim for maximum compatibility with many viewers, especially when publishing images to the web. It is also the fastest option.

How to obtain Nikon profiles from NX2

The Capture NX 2 and the related View NX software is proprietary software for Microsoft Windows. The Nikon input ICC profiles are not usually distributed. Fortunately, the software extracts and generates input profiles on-the-fly in a temporary folder during runtime.

1) First you have to install Capture NX 2 (version 2.3.8 was tested), but View NX 2 also should work. You do not need to purchase Capture NX 2, the trial version is sufficient. View NX 2 is free.

2) Open the NEF file in Capture NX 2. While the program is running and your picture is open by the program, look at the following hidden folder:
<Windows directory>:\Users\<your username>\AppData\Local\Temp\Nkn<random strings>.tmp
In this temporary folder will the ICC profiles created, similar to this one:
Nkx_D90_962_1_03_0_00_10_00_00_00_05_00_0200_0_0_476.icm

3) Select/change the „picture control” type in Capture NX 2, and see how new ICM files are generated by NX! The reason: Nikon applies generic profiles to your pictures. These are named „picture control” settings. For example, Capture NX 2 has „neutral”, „standard” and „monochrome” profiles. For use with RT you must simply know that Nikon generates different ICM profiles with every different setting. To ease the choices, the „neutral” ? „standard” ? „vivid” ? „landscape” settings seem to be similar, apart from the increased saturation and contrast. You should be able to tune the tone curves in RT, so in general, a single „neutral” profile should be good starting point.

WARNINGS: other settings may also generate new ICM files! You must be particularly aware that pictures shot under artificial lighting will have different profiles, and using such profiles with daylight shots will result in color casts. D-lighting dramatically changes contrast.
If in doubt, open your particular image you want the color profile for and use that profile in RT. In most cases, a generic color profile will be sufficient.

Transform tab

Crop panel

Quite self-explanatory. Use the Select crop button to drag a crop with your mouse. Use the Shift key to move the crop over the image. Resize a crop by placing the mouse on one of the sides, cursor changes into a double arrow. Click inside of the crop to show (temporarily) the whole image. Go to the Preferences window, tab General, and click on Cut overlay brush to change the color and/or transparency of the area around the crop. Use Fix ratio to set the

crop to a fixed size like 3:2, 16:9, etc. 1:1 is for square crops. DIN means paper sizes like A4 and A5, those have fixed width/height proportions. You must tell RawTherapee if you want to have a horizontal (landscape) or vertical (portrait) crop. Use Guide types if you like. The DPI value informs about the print size of the current crop (or the whole photo if there is no crop). For prints up to 13x18 cm (5x7 inch), 300 dpi is a good choice, bigger prints can have less pixels per inch.

Resize

Can be applied to a crop or the full image. Choose your favorite method. Resize by Scale: use the slider to set scale. Or specify width, height or bounding box (developers lingo for width and height).

Warning: the Resize tool will not be previewed. This is a limitation of the current RawTherapee's engine, but the output image will of course be resized.

Lens / Geometry panel

Auto Fill

When checked, the photo will be enlarged to an extent that no black borders are visible anymore. Tip: to better understand what happens with the following tools, uncheck this option for the moment.

Rotate

Rotate the image between -45 and +45 degrees. Use the Select Straight Line button to tell RawTherapee what line has to be straight. Use the mouse to draw this line.

Perspective

Horizontal

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', eg. when photographing architecture.
Higher values for both sliders produce heavily distortion, so use with care. Or don't care at all and have fun!

Lens correction profile

Select an Adobe LCP file here to automatically correct geometric distortion, vignetting and lateral chromatic aberrations. It uses its own, more precise correction model and is not linked to the other controls. However you can still use the other controls in addition to the LCP profile e.g. for artistic reasons or if the LCP fails to correct a parameter (which happens on some extreme distortion occasions, like with some heavily distorting compact cameras).
The following restrictions apply:

- LCPs are only supported for RAW files (JPGs will only correct distortion and CA, but not vignetting).
- While distortion is visible in the full image preview, CA and distortion are not reflected in the detail crop windows, only in the fully processed result image. Auto filling is also not supported.
- CA correction is only supported if the EXIF information contains the focus

distance (e.g. in DNGs from Nikon files).

Distortion

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.

Chromatic aberration

The chromatic aberration can corrected by using the two sliders Red and Blue. Normally you won't see any C/A in the preview therefore it is highly recommended to use the Image Details (I even suggest to magnify to 200%) when you apply 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.

Vignetting Correction

The difference between a cheap telelens and an expensive one is, among others, that the first one likely will produce vignetting, while the expensive lens will not or less. Vignetting means light fall-off in the corners of the image, making them darker than the center. This can easily be corrected with this tool.

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, to the level of an effect filter. Set Amount to -100, Radius to 50 and move Strength from 1 to 100 to see how this works.

Center X

Moves the 'circle of correction' to the left or the right of the image. Use this to apply vignetting correcting to a cropped photo or to darken an area on the left or the right that is a bit overexposed (think flash).

Center Y

Moves the 'circle of correction' up or down over the image. Useful to 'burn' skies or lighten the foreground.

Raw tab

Demosaicing panel

Method

To display a raw file - any raw file made by any camera brand - its data has to be 'demosaiced'. That's why opening a raw file always takes a bit longer than opening a jpg or tif file, where the data are already 'display-ready'. RawTherapee offers several demosaic algorithms, each with its own characteristics. The differences between them are often very subtle; one needs to zoom in to 300%, 400% or even 500% to really see what's going on. The choice of a certain algorithm influences the quality of very fine details in the image and decides how well colored edges are rendered, among others.

RawTherapee always uses the Fast algorithm to open the image for editing (so to display). After this the selected demosaic 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 is the default demosaic method of RawTherapee, as it yields the best results for most users. The Hphd method was the default method in earlier versions of RawTherapee and is still quite capable. It is faster than Amaze. Eahd also comes from earlier versions of RawTherapee. It is somewhat slower than other methods because it can only use one processor core. VNG-4 used to be (in version 2.4) the preferred algorithm for Olympus cameras, as it eliminates certain mazing artifacts that might be created by the other methods. In version 3 Olympus users might prefer the Amaze method though. DCB is similar to Amaze, it also removes maze artifacts. Completing the list are the Ahd algorithms. Just try what works best for you.

False color suppression steps

Sets the number of median filter passes applied to suppress demosaic 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.

Preprocessing panel

Line Noise Filter

This filters pattern noise. The slider sets the strength of the effect.

Green Equilibration

Some cameras (for example Olympus, Panasonic, Canon 7D) use slightly different green filters in the two green channels of the color filter array on the camera sensor. Green equilibration suppresses interpolation artifacts that can result from using demosaic algorithms which assume identical filter properties in the two green channels. The threshold sets the percentage difference below which neighboring

green values are equilibrated.

Apply Hot/Dead Pixel Filter

Suppresses hot and dead pixels by replacing them by a neighborhood average.

Apply post demosaic artifact/noise reduction

Refinement from Luis Sanz based on EECI demosaicing algorithm by L. Chang and Y.P. Tan. You can see the effects at 100% view with sharp edges, e.g. with faint tree lines.

Raw White point panel

Linear correction factor

Applies a linear exposure correction on the raw file before the demosaic algorithm is used. This can be useful to deal with difficult lightning situations that resulted in overexposed areas in the raw. The value of this slider represents the multiplier of the RGB channels. The correction factor is independent of color space.

Be careful when using it, since it is shifts the input before the color management system translates them, so you might see color shifts. On the other hand reducing the raw white point can recover highlights which would be lost when going through the color management system with ICC profiles.

HL preserving corr. (EV)

Use this slider to preserve the highlights. The value here represents an EV value, so the equivalent of a f-stop. When the correction factor (the first slider) is positive, a high value for the second slider makes highlights somewhat darker, thus preserver them. On the contrary, when the first slider is negative, a high value of the second slider makes the highlights lighter. This slider only works when the Linear correction factor slider has a value other than 1.0.

Black levels

These allow you to add offsets to the channels still in raw (that’s why there are two greens, because a bayer sensor has two green sensors). It helps to fine tune color tints in shadow areas.

Dark Frame panel

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 frame".

Attention, this method is not effective with high ISO noise because of its different random nature.

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 one directory specified in Preferences. To select the directory go in "Preferences...", tab "Image processing" and select "Dark Frame Directory"; you may wait a little after selection because RT scans all files

in that dir. 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 darkframe shots there if not already done. You could also move a shot that you find to be a dark frame with right click menu in thumbnails browser "Move to dark frame directory". RT chooses the best match looking for same camera model shots with minimal difference in ISO, exposure duration and date. If more then 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 autoselection), RT 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.

Bad pixels

RT can correct a list of bad pixels (pixels that are always black or white or stuck to 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. (If you look at the coordinates of pixels in RT, beware of the offset introduced by cut border: you must add +4 to each coordinate).

The file has to be named exactly as your camera:"**maker model.badpixels**" (maker must be short name) es: PENTAX K200D.badpixels; if you have two cameras with same model you can specify also the serial numbers (look at exif data): "**maker model serial.badpixels**". The file has to be located in the "dark frames" directory specified in preferences.

Pixels in the bad pixels list will **always** be corrected in processed photos.

Bad pixel detection software

There exist programs to aid in detecting such pixels:

- Dead Pixel Test
  http://www.starzen.com/imaging/deadpixeltest.htm
- Pixel Fixer
  http://www.pixelfixer.org/

Flat Field panel

Original image

AFTER Auto Flat Field correction

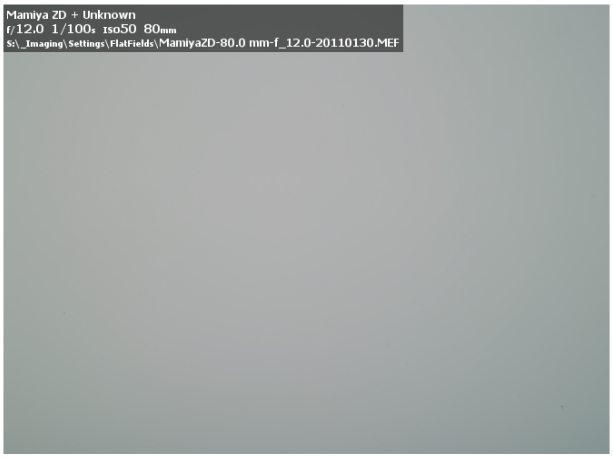

Flat field image at similar aperture

The same flat field image shown with exaggerated contrast and saturation to visibly demonstrate color and luminance non-uniformity due to the lens cast effect

Flat field corrections are 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 the 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 by a possible misalignment of the lens mount or by the usage of the tilt-shift lenses. 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.

Manual corrections for these effects in post production are quite difficult; especially when to be reproduced on series of images captured under various conditions, and would rarely lead to perfect corrections.

Flat field correction tools in RawTherapee allow 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 time only preview in the Editor 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 file. Flat Field raw image is usually a de-focused photograph of a uniformly lit uniformly colored and toned subject, such as a piece of white matte paper. Another method of capturing flat field reference image is to use a uniform piece of white Plexiglas® in front of the lens (at full contact with the lens/filter barrel) and photograph it against a preferably uniform illumination (e.g. point camera towards the 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.

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 used to correct the corresponding issues in the image file. Area blur uses a 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 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 0-1 pixels will subtract out the darkening caused by the dust. Use the vertical, horizontal, or vertical+horizontal blur options if the camera generates repeatable line pattern noise. For instance, consistently repeatable vertical line patterns can be subtracted off by using vertical blur on the flat field file.

Organising Flat Fields

Non-uniformity of the captured field depends on the following parameters:

- Camera (camera & sensor combination in case when digital back is used),
- Lens,
- Focal distance,
- Aperture,
- Lens tilt/shift

It is recommended to assemble a library of flat fields for camera/lens combination, taken at various (expected to be used for real image capture) aperture settings. It is advisable to name flat fields descriptively, so that files can be recognised easier by the user, preferably, incorporating all parameters listed above. During flat field correction process these parameters are being read from exif data only and contents of the file name 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:

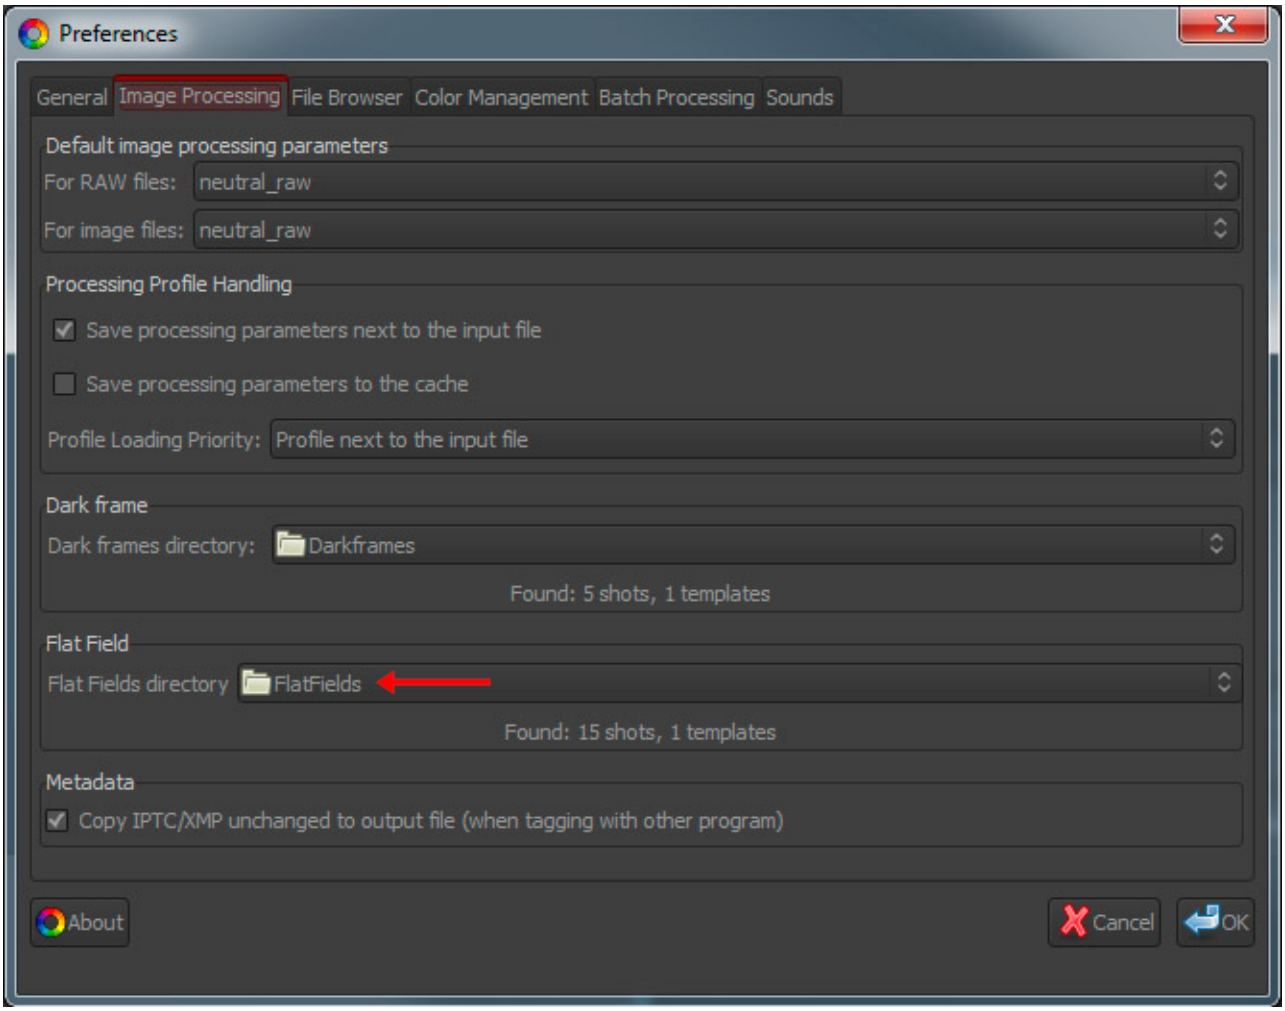

Setting flat fields directory in the preferences window initiates analysis of its content. Counts of qualifying files and templates are reported in the UI.

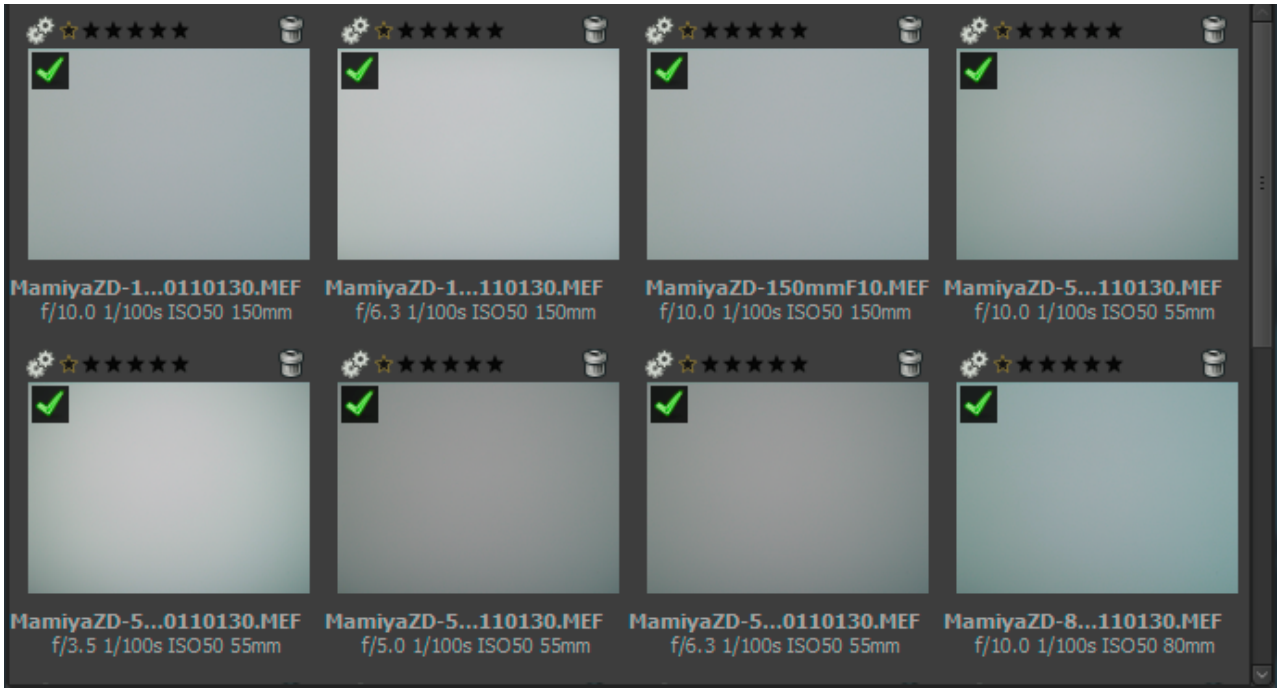

Thumbnails of sample Flat Fields displayed in RawTherapee file browser

Auto Selection

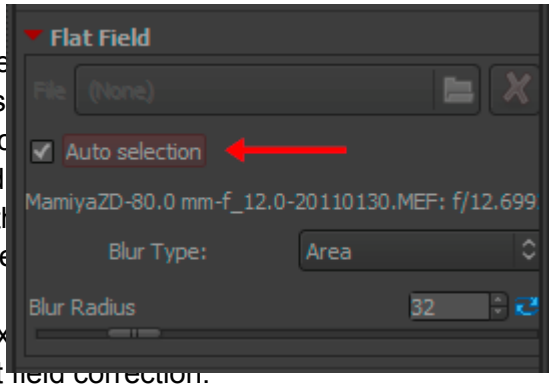

Flat field Auto Selection capability can be
Selection" checkbox. RawTherapee will s
directory specified in preferences and selec
match to the image being corrected based
Focal distance, Aperture and the date of t
name of the selected flat field image will be
If match is not found, flat field correction
displayed to the user. If more than one ex
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. It is recommended to apply such flat fields correction via manual selection of the required file by launching the file browse dialog in the Flat Field panel or a right-click menu option (see below).

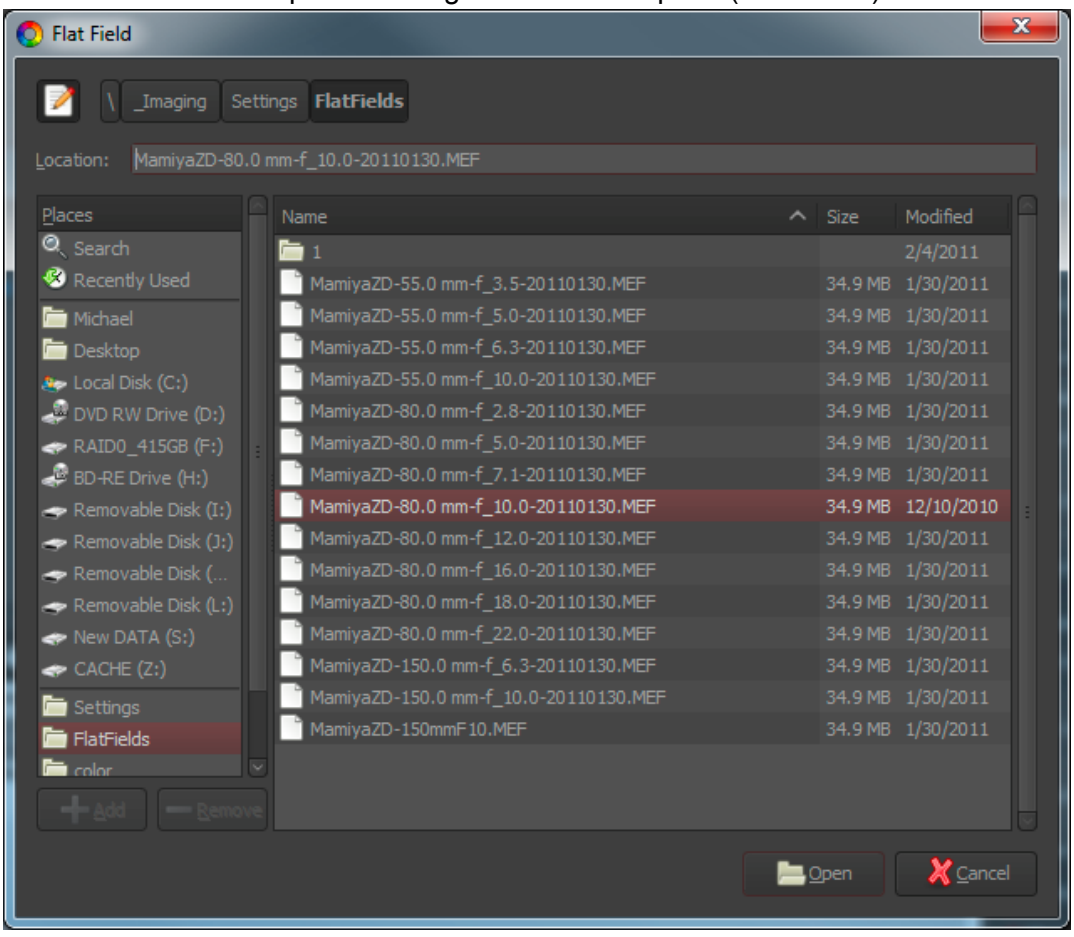

Illustration of the flat file selection dialog.

Menu Options

File browser allows to manage flat fields via the Flat Field menu options.
"Select flat field" brings up the file selection dialog to select a flat field file to be applied to the selected thubmnail(s).
"Auto flat field" menu allows to set "Auto Select" option to the selected thumbnail(s).
"Move to flat fields directory" moves selected image(s) into the directory specified in

the preferences.

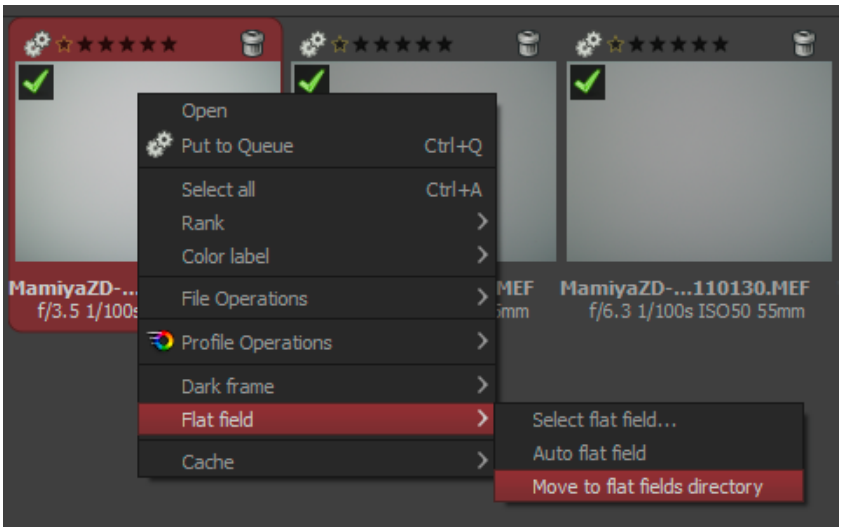

Blur Type

Blur Type	Purpose
Area	A default and generally most useful setting to apply blur action equally in all directions. Works well for correcting vignetting and lens cast.
Vertical	Blurs flat field in vertical direction to compensate for vertical non uniformities. This is useful if vertical sensor readout has variation between the colomns.
Horizontal	Blurs flat field in horizontal direction to compensate for horizontal non uniformities. This is useful if horizontal sensor readout has variation between the rows.
Vertical + Horizontal	Blurs flat field sequentially in horizontal and then vertical directions to compensate for both vertical and horizontal non uniformities.

Blur Radius

Blur radius controls the degree of blurring of the flat field data. Default value of 32 is usually sufficient to get rid of localised variations of raw data due to noise. Setting the Blur Radius to 0 skips blurring process and allows to correct for dust and other debris on the sensor (as long as their position has not changed) 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.

All parameters in the flat field panel are saved into the development profile of the image being edited. These settings can be copied and pasted to other images just as all other development settings. This includes 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.

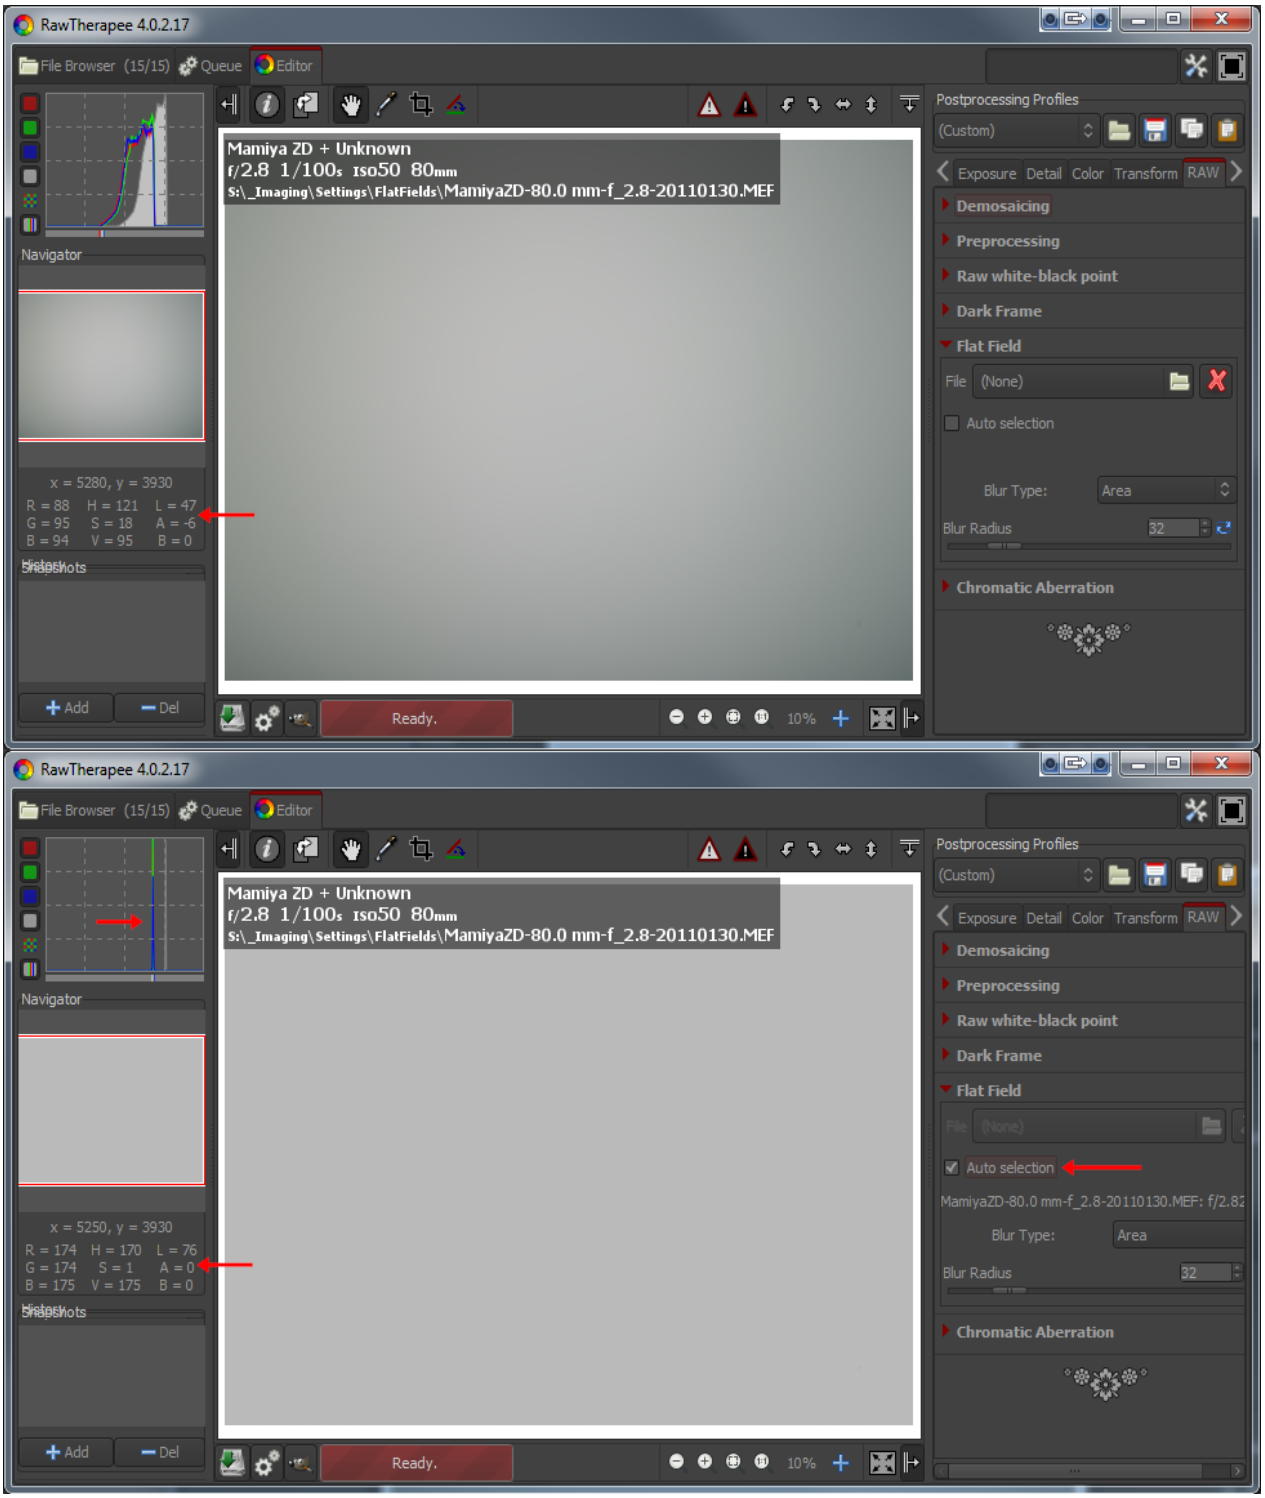

Demonstration of the effects of flat field removal from the flat field photo itself

The figure above illustrates the application of a flat field file onto itself. There is a noticeable asymmetry shown in “before” image on top in the light falloff as well as in the green color cast. Flat field correction removes both color and luminance non uniformities and results in a perfectly uniform image. Histogram on the bottom (after) image indicates that 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.

Chromatic Aberration panel

This tools corrects bluish-green and magenta fringes due to lens lateral chromatic aberration that shows mainly in the borders of the image. This correction is performed before actual demosaicing and can sometimes improve the quality of demosaicing step, but it is rather slow, so its results are shown only when zooming 100% or more.

Auto correction

If auto correction is checked the Red/Blue sliders are disabled and an automated detection and correction of chromatic aberration is attempted.

Red and Blue

If the Red/Blue sliders are non-zero the given values are used to correct chromatic aberration.

Metadata tab

Exif tab

Here you can control which 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 informations are directly visible. Extended EXIF information and so-called makernotes are organized into a tree. Click on the arrow at the very left of the desired subtree and you'll see its contents. Makernotes are informations varying between camera manufacturers and even between camera models.

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 similar but is used for trees and works recursive, which means that all values changed/removed in this subtree are restored.

You can Add/Edit the following EXIF information: Artist, Copyright, ImageDescription and 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.

Tip. If you want to add to every processed photo the Artist field for example (so that every jpg contains your name), just do this. Open a raw (it will be opened with the profile you use by default), go to the Exif tab, click Add/Edit, choose Author, fill in your name and save this profile with a new name, like default-exif. Now open the Preferences window, section Image processing, and say there that every raw has to be opened with this default-exif profile.

IPTC tab

IPTC values belong also to the group of metadata, as they are integrated into the image files but do not add any picture details (additional or different pixels) to the 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 the workflow a lot as you don't have to care about another file when e.g. backing up or sorting your images.

IPTC is usually used to describe the image in detail. There are a lot image database softwares 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 companies willing to sell your images support IPTC tags when you upload your images to their databases, thus you have less work. Adding e.g. keywords on your computer at home is much more comfortable than doing it through the web browser. Multiple Keywords and Suppl. Categories can be added/removed using the plus and minus signs next to it.

The button Reset (under) resets the IPTC values to those saved in your current profile. There are two more buttons, one 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. The other one pastes the formerly copied IPTC settings from the clipboard to your current image.

Preferences window

General tab

Workflow

Choose here between single editor tab and multiple editor tabs mode. If you own a second monitor, you can use that one as well. With ‘Histogram in left pane’ and ‘show profile selector’ you can control what and where the tools are shown in the edit panels. If helps you to safe GUI space.

A restart is required for these options to take effect.

Default language

Select the language for the GUI out of a list of thirty languages. English (US) is the default ('mother') language, translations are based on that one. On Win Vista/7 64bit you can have the language automatically read from the operation system.

Again, a restart is required to change the language of the GUI.

Default theme

Choose between several themes for the GUI, from light to dark. The effects are visible after a few seconds, so no need to restart here. Checking 'Use System Theme' might change the appearance of RawTherapee, although this depends on the platform and the window manager in use. Just see if it works for you.

Cut overlay brush: this is the color of the area outside of a crop. By clicking on the colored button, a new window appears where you can also set transparency. If set to 75, the cropped area is still somewhat visible. Useful to move the crop around and to find the best composition (hold the Shift key and move the crop with the mouse).

Choose the font of your liking here. With smaller fonts more tools can be displayed on the screen.

Clipping indication

When checked, the raw photo blinks at underexposed and overexposed areas. The both threshold sliders determine at which values blinking begins (0..255). In RawTherapee 4.0 the clipping indicators are calculated on the final image in the output space selected for the image. The clipping indicator takes the distance to the maximum/minimum and renders these as grayscale values. Mention the clipping is only calculated at the end of the processing chain. However, since RawTherapee 4.0 has a floating point engine, there are fewer positions within the processing chain where clipping can occur at all.

Date format

Determines the date format of the thumbnails in the file browser window.

External editor

When you send a raw from within RawTherapee to an external editor for further

processing, RawTherapee needs to know which program you want to use. If that's not the Gimp, you can specify the path to your favorite photo editor here.

Custom profile builder

Executable (or script) file called when a new initial profile should be generated an image. The executable receives command line parameters to allow a rule based .pp3 generation:
[Path raw/JPG] [Path default profile] [f-no] [exposure in secs] [focal length in mm] [ISO] [Lens] [Camera]

In the ‘samples’ folder you can find the RTProfileBuilderSample.cs / cs.config file which can serve as a basis for your own custom profile program. It is heavily documented in the files.

This feature is very powerful to automatically set e.g. lens correction parameters or noise reduction base on image properties. It is called just once on first edit of the picture.

About

Shows information about the original author of RawTherapee and the current version, details of the build, names of developers and other contributors and the licence under which RawTherapee is published: GPL 3.

Image Processing tab

Default image processing parameters

Specify which profile RawTherapee has to use when opening a raw and when opening a non-raw photo. When you have made your own default profile, you can tell RawTherapee to always use that one.

The profile for non-raw files like jpg or tif can probably best set to neutral. The neutral profile just loads the photo as it is, without applying for example Auto exposure or sharpening.

Processing Profile Handling

Save processing parameters next to the input file: When checked, RawTherapee writes a pp3 file with all the edits you made to your photo next to the input (raw) file. This represents your work (e.g. sharpening settings used) and can be reloaded later.

Save processing parameters to the cache: Instead of creating a pp3 file next to the raw, this option - when checked - writes the pp3 to the cache. When you check the last option only, chances are that you lose your work (the+ edits) when installing RawTherapee on a new pc for example.

It’s usually a good idea to only save the processing parameters next to the input file, since you can e.g. back them up along with the your raws.

Dark Frame

Specify the directory on your hard disk for searching for the dark frame shots for long

exposure noise subtraction. File with coordinates listing of the bad pixels must be placed in the same directory for auto correction.

Flat Field

Specify the directory on your hard disk for searching for the flat field reference images.

Metadata

The option “Copy IPTC/XMP unchanged” changes RTs metadata behavior. Usually RT will remove all IPTC/XMP data from the input image and write only its own tags in IPTC section. This may become a problem when you tag your input files using other programs, raw usually contain XMP data. This would be lost. By checking this option RT will not touch IPTC and XMP data at all, just passing them through. On the downside all your tagging within RT will not be saved though.

File Browser tab

Image directory at startup

At the top you can define the Image directory at startup. It could be the RawTherapee Installation directory, the Last visited directory, the Home directory, or you define another directory.

Next you can set some File Browser / Thumbnail Options. This option defines whether you can see date and time or EXIF information below the images in the file browser.

The Menu Options determine the grouping of the right-click context menu in the file browsers.

You can also define a list of parsed extensions, i.e. which files are recognized as images. Some extensions are defined by default. Those can be deactivated by removing the checkmark in front of it. If a desired extension is missing you can easily add it by using the plus button.

At the bottom of that page there are several Cache options that influence the speed of loading/generating the thumbnails. These options are quite self-explaining. The JPG option is usually the fastest and uses the least disk space, however the thumbnail quality is the least.

Color Management tab

Here you can define the directory where ICC profiles could be found. You should also define the ICC profile of your monitor when you've done a calibration. If you don't do it, the image will be displayed with wrong colors.

The option “Automatically use operation systems main monitor profile” is currently only supported on Windows. Please mention that RT only knows about one monitor. So if you have multiple monitors connected, it will always take the main monitors profile (the one with the task bar).

You can also set the colorimetric intent which defines how the ICC profiles are used for translation between gamuts or color spaces

Perceptual
If the color gamut of your image is higher than that one your device (monitor or printer) is able of then it is compressed a bit to fit the gamut of your device as far as possible. This might result in an image with reduced saturation, the hue is still kept. It might look a bit dull. But this is not really that much visible as the color relations stay the same. This method is activated by default (recommended)

Relative Colorimetric
The colors existing in the color gamuts of both your image and your device are kept and displayed 100% perfect. If the color does not exist within the color gamut of your device the nearest possible value is taken. This might lead to some banding effects, especially visible in blue sky. The white point will be corrected.

Saturation
Very similar to Perceptual, but here it is tried to keep the saturation and change the hue instead. This is very useful for e.g. screenshots or similar. It could also be used when you do not care of some possible color shift as long the image does not look dull.

Absolute Colorimetric
Similar to Relative Colorimetric. It tries to reproduce the exact colors recorded in the original scene. The white point will not be corrected. It is normally used, when the gamuts of your image and your device are nearly the same. For example for reproduction of specific colors like reproductions of fabric or logo colors.

Batch Processing Tab

The Batch processing is the capability of editing several images at the same time in the 'File Browser' tab. That's why there's a tool editor on the right of it.

This is done by selecting more than one file by using the Shift or Control key in the File Browser, then you can edit those images with the so-called 'Batch tool panel editor'. The way the sliders' value is used to modify the image depends on the options set in this 'Batch Processing' tab.

When you select a single image, the sliders get the values of the processing parameters of that specific image. These can be the values of the default profile or the values from your last edit session of this photo. If your image is currently being edited in an Editor tab, the editors' values will be reflected in realtime in the Batch tool panel editor, and vice versa, so take care of what you're doing.

But if you select several images, the behavior of each sliders can be used in two different ways.

The 'Add' mode
The slider set to work in this mode will change its range and default value right

after the image selection, in order to add the value of the slider to the value of the according parameter of each image.

For example, if you select two images by holding the Ctrl modifier key, one image that has an 'Exposure compensation' of -0,5 EV, and another one which has an Exposure compensation of +1,0 EV, moving the 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! This mode may also be understood as 'Relative mode'.

Using the 'Reset' button will move the slider to its default (zero) position and will then bring back the initial value of each images for this parameter.

The 'Set' mode

In this mode, the value of the slider will set 'as is' the value of the corresponding parameter.

If we use the same example as before, moving the slider up to 0,3 will result in setting a value of 0,3 EV for both images (one value for all images). This mode may also be understood as 'Absolute mode'.

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.

You can select the desired behavior for each slider in this tab. Please note that all sliders are not listed, because they are supposed to work in SET mode only.

The option “Overwrite existing output file” sets RT to overwrite an existing e.g. output jpg if it exists. If this is unchecked and you e.g. send an image a second time for conversion to the batch queue, RawTherapee will add a numeric index to the file name (e.g. myoutput_1.jpg).

Sound Tab

Here you can get notified with a sound when a lengthy operation has ended. It is currently only supported on Windows.

The *Queue processing done* sound is played after the last queue entry finished 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 enabling the mute option or by setting fields with sound file references to blank values.

File paths to sound files can refer to WAV files. Alternatively, the sounds configured in Windows system settings may be specified by using one of the following values: *SystemAsterisk, SystemDefault, SystemExclamation, SystemExit, SystemHand, SystemQuestion, SystemStart, SystemWelcome.*

Command Line Options

RawTherapee may be called from the command line in the following different ways:

`rawtherapee [<selected dir>]` : start RT GUI browser inside dir.
`rawtherapee <file>` : start GUI editor with file.
`rawtherapee -c <inputDir>|<file list>` : convert files in batch with default parameters.

Other options used with `-c` (that must be last option)
`rawtherapee [-o <output> | -O <output>] [-s|-S] [-p <file>] [-d] [-j[1-100]|-t|-n] -Y -c <input>`

`-o <outputFile>|<outputDir>` : select output directory.
`-O <outputFile>|<outputDir>` : select output dir and copy pp3 file into it
`-s` : include the pp3 file next to the input file (with same name) to build the image's parameters, eg: for IMG001.NEF there should be IMG001.NEF.pp3 in the same dir if absent use default
`-S` : like -s but skip if pp3 file not found.
`-p <file.pp3>` : specify pp3 file to be used for all conversions. You can specify as many -p options as you like (see * below).
`-d` : use the default Raw or Image pp3 file to build the image's parameters
`-j[compression]` : specify output to be jpeg.(default)
`-t` : specify output to be uncompressed tiff.
`-t1`: specify output to be compressed tiff.
`-n` : specify output to be png.
`-Y` : overwrite output if present.

Note:

- Your pp3 files can be incomplete, RawTherapee will set the values like that:
- the pp3 file is built with internal default values;
- then overridden by those found in the default Raw or Image pp3 file (if -d has been set);
- then overridden by those found in the pp3 files provided by `-p`, each one overriding the previous values;
- * then overridden by the sidecar file if -s is set and if the file exists; the time where the sidecar file is used depend of the position of the -s switch in the command line regarding to the `-p` parameters (e.g. "`-p first.pp3 -p second.pp3 -s -p fourth.pp3`")

Examples

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 dir as “foo.tif”, and overwrite the file foo.tif if it exists:

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

In the next example, we'll assume that you want to quickly process all your raw photos from the /tmp/Jane01 directory to a web subdirectory 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 -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 like that:

1. A default profile is created with internal values
2. The Raw default profile (-d) will override the default profile with its values
3. iptc.pp3 will override the profile with its values
4. if it exist, the sidecar profile will override the profile with its values, so you can force some IPTC tags even if already set by iptc.pp3
5. exif.pp3 will override the profile with its values, so you can force the profile to erase some tags
6. web.pp3 will override the profile with its values, 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 regarding to the other -p parameters. That's not the case for the -d switch.

In the third example, we will see how long it takes to process every raw file in a directory, assuming each raw photo has a corresponding postprocessing profile created previously in RT:

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

Shortcuts

Following is a list with the keyboard shortcuts and some special mouse actions.

Mode	Shortcut	Action
Any	**Ctrl-F2**	Switch to File Browser
Any	**Ctrl-F3**	Switch to Queue
Any	**Ctrl-F4**	Switch to Editor (when in single tab mode)
Any	**F11**	Toggle fullscreen.
Editor	**1**	Zoom to 100%
Editor	**- or _**	Zoom out.
Editor	**+ or =**	Zoom in.
Editor	**<**	Show/hide highlight clipping.
Editor	**>**	Show/hide shadow clipping.
Editor	**Shift-B**	Before/After view.
Editor	**C**	Crop tool. **Shift-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.
Editor	**Ctrl+Shift+Z**	Redo.
Editor	**Ctrl+Z**	Undo.
Editor	**Double-left-click on canvas outside photo**	Cycle canvas color (black / theme background color / white).
Editor	**F**	Fit to screen.
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**	Show/hide EXIF info overlay.

Editor	**Right-click in the preview area**	Get back to default cursor (after having used the pipette for white balance or the straightening tool). If RGB indicator bar is enabled, this freezes/un-freezes its indicators
Editor	**S**	Select straight line (Rotate tool).
Editor	**Shift+drag crop**	Allow to move crop border.
Editor	**W**	White balance.
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	**Shift-F**	Toggle Preview mode - Focus mask
File Browser	**~**	Show only images without a rating.
File Browser	**1 through 5**	Show only images with specific rating
File Browser	**6**	Show only not edited images.
File Browser	**7**	Show only edited images.
File Browser	**Alt + ~**	Show only images without color label.
File Browser	**Alt+ 1 through 5**	Show only images with specific color label.
File Browser	**Alt+6**	Show only images that were not recently saved.
File Browser	**Alt+7**	Show only images that were recently saved.
File Browser	**Ctrl+O**	Sets focus to Browse path, selects contents. Browse Path shortcuts: ~ - user's home directory ! - user's picture directory Press "Enter" to refresh. Press Ctrl+Enter to reload
File Browser	**Ctrl+F**	Sets focus to the Find box, selects contents. Press "Enter" to execute search.
File Browser	**D**	Clear all filters: ratings, color labels, edited(not), saved(not), trash can
File Browser	**F2**	Call up the dialog for file rename. If multiple thumbnails

		selected, dialog will be called for each.
File Browser	**T**	Show trash contents.
File Browser, Editor	**Alt+E**	Switch to exposure tab.
File Browser, Editor	**Alt+D**	Switch to details tab.
File Browser, Editor	**Alt+C**	Switch to color tab.
File Browser, Editor	**Alt+T**	Switch to transform tab.
File Browser, Editor	**Alt+R**	Switch to raw tab.
Editor	**Alt+M**	Switch to metadata tab.
File Browser, Editor	**Ctrl+L or L**	Toggle left panel
File Browser, Editor	**Alt+L**	Toggle right panel
Editor	**Shit+L**	Toggle Top panel
Editor	**Ctrl-Shift-L**	Toggle Left and Top panels
File Browser, Editor	**Alt-Shift-L**	Toggle Right and Top panels
File Browser, Editor	**Ctrl-Alt-L**	Toggle Left and Right panels
Editor	**Ctrl-Shift-Alt-L**	Toggle Left, Top and Right panels
File Browser, Editor	**M**	Toggle all peripheral panels (*m*aximize main panel)
File Browser, Editor	**Ctrl+Q**	Add current image to processing queue.
Tool Panel	**Ctrl+click on a slider's reset button**	The slider is set to the value of the initial profile, that has was used when image was loaded in the editor, or when the image was selected in the File Browser.
Tool Panel	**Ctrl+drag a curve point**	The point moves slowler than the mouse cursor. Allows finer adjustmets to the curve
Tool Panel	**Shift+drag a curve point**	The point is snapped to key positions highlighted in red.

Tool Panel	**Right-click over a tool panel title (foldable section)**	Unfolds clicked section and closes all others.

Windows-only shortcuts

Mode	Shortcut	Action
File Browser, Editor	**F5**	Open result image in default 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 directory and select current image.
File Browser, Editor	**Shift+F5**	Open Explorer in the queue output directory and select current image if already converted.
	Alt	When pressed while working in edit window, it temporarily enables both shadow and clipping indicators during preview image updates.

Contributors

The following people contributed towards this manual:

- Emil Martinec
- Fabio Suprani
- Jean-Christophe Frisch
- Maciek Dworak
- Michael Ezra
- Oliver Duis
- Paul Matthijsse

Licence