darktable 1.1

darktable 1.1

The owner of the darktable project is Johannes Hanika. Main developers are Johannes Hanika, Henrik Andersson, Tobias Ellinghaus, Pascal de Bruijn and Ulrich Pegelow.

Table of Contents

Preface to this manual

The current user manual is a draft issue. It is at an early stage of writing. User manual version and applicable darktable version are listed below:

	version	date
user manual	1.1.2 (draft)	December 2012
darktable	1.1	December 2012

For the course of this documentation the arrow keys on your keyboard are referenced as with arrow symbols: ↑←↓→.

Chapter 1. Overview

darktable is an open source photography workflow application and RAW developer. A virtual lighttable and darkroom for photographers.

It manages your digital negatives in a database, lets you view them through a zoomable lighttable and enables you to develop raw images and enhance them.

General Features

- darktable runs on GNU/Linux / GNOME, Mac OS X / macports and Solaris 11 / GNOME.

- Fully non-destructive editing.

- All darktable core functions operate on 4x32-bit floating point pixel buffers for high accuracy processing; preventing banding and color breaks.

- Heavy use of SSE instructions of your CPU for speedups.

- GPU acceleration via OpenCL (runtime detection and enabling).

- Most of the actions on the pixelpipe are done in CIELab color space, which is much larger than the gamut of modern displays, printers or even human vision.

- Full color managed display with softproofing and gamut-check. Built-in ICC profile support for export: sRGB, Adobe RGB, XYZ and linear RGB.

- A collect module allows you to execute flexible database queries, search your images by tags, image rating (stars), color labels and many more. Filtering and sorting your collections within the base query or simple tagging by related tags are useful tools in your every-day photo workflow.

- Import a variety of standard, raw and high dynamic range image formats (e.g. jpg, cr2, hdr, pfm, ..).

- darktable has a zero-latency fullscreen, zoomable user interface through multi-level software caches.

- Tethered shooting.

- The powerful export system supports Picasa webalbum, flickr upload, disk storage, 1:1 copy, email attachments and can generate a simple html-based web gallery. darktable allows you to export to low dynamic range (JPEG, PNG, TIFF), 16-bit (PPM, TIFF), or linear high dynamic range (PFM, EXR) images.

- darktable uses both XMP sidecar files as well as its fast database for saving metadata and processing settings. All Exif data is read and written using libexiv2.

- darktable comes with 47 image operation modules which cover everything from basic operations, tonal value changes, color manipulation, correction of common image defects to artistic effects.

- Many of darktables modules can be combined with blending operators for even more development options.

- A powerful blend mask feature gives you fine control on module's effect to different parts of an image.

1. User interface

This section describes the layout of the user inter-
face. You can change darktable's contrast by using *F7*
and *F8* and darktable's lightness by using *F9* and *F10*.
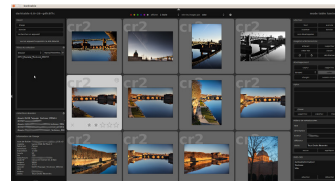
1.1. Views

darktable consists of several views also mentioned as modes. For the time being, there
are four available views as described in this section. The current view is identified by the
label at the top of the right panel. Some key accelerators, allowing to access the different
views, are described in the table below:

l	switches to lighttable view
d	switches to darkroom view
t	switches to camera tethering view
m	switches to map view

1.1.1. Lighttable

The lighttable view is where you manage your images and filmrolls; it's in that view that
you rate, add tags, add colorlabels, export images among other actions.

You can perform some actions on the image under the mouse pointer using key accelera-
tors; see "Lighttable" chapter.

1.1.2. Darkroom

In the darkroom view you develop a specific image using the available modules.

1.1.3. Tethering

This view is where you do camera tethering, ie. connecting your camera to the computer
and remotely capturing images that will be downloaded and shown on your computer
screen.

1.1.4. Map

This view shows images that are geo-tagged on a map and allows you to manually geo-
tag new images.

1.2. Filmstrip

The filmstrip shows the same images as lighttable with respect to filters and sort order; it
is turned on/off with key accelerator *ctrl-f*. You can navigate along the filmstrip by wheel
scrolling. The filmstrip allows you to interact with images when you are not in lighttable
mode. For example you can, while being in darkroom mode and developing an image,
switch to another image to develop by double clicking the thumb in the filmstrip. You can
also rate the images as you do in lighttable, copy/paste history stack, etc.
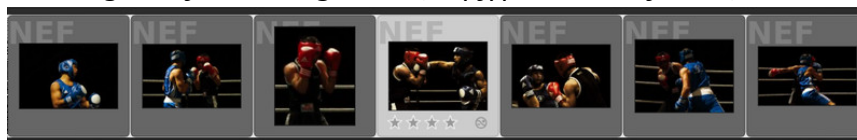
1.3. Preferences

A button ⚙ located at the upper panel allows you to define various parameters.

The options are quite self-explaining. In case you need more information, just let your mouse hover a text label and/or entry box to display more information. Furthermore all configuration parameters are explained in a later chapter of this manual.

2. darktable basic workflow

2.1. Importing Images

To begin with darktable, you first need to import images. You can either import from your filesystem or from your camera if supported.

2.1.1. Importing images from filesystem

In the left panel of lighttable view you will find the import module, here you have the choice to import either a single image or a folder, darktable will analyse its content, detect images that are already imported and only import new images.

2.1.2. Importing from camera

Connect your camera to your system and if your distribution tries to automount it, select the option to abort the mount operation. Otherwise the camera will be locked and not accessible from within darktable.

In the left panel of lighttable view you have the import pane. If you don't see your camera there, you need to hit the scan for devices button. Your camera will then appear in same pane with additional actions: import and tethering.

2.2. Basic development steps

2.2.1. Introduction

This section will guide you through the basics about developing your image in darkroom view.

To begin you need to open an image in darkroom mode which you will do by double clicking an image thumbnail on the lighttable, it's in the darkroom mode you do the actual adjustments for an image and you have an arsenal of modules at your hand to help you reach your goal.

Each change you make on a module when developing your image is turned into a history stack item, the history is stored in a database and in an XMP sidecar file for the specific image. Changes are stored automatically without the need to press a "save" button. Therefore you can safely leave darkroom mode and darktable and come back later to continue your work.

On the left panel in darkroom mode you have the history stack showing changes starting from bottom and building up with each change you have made to the image. You can select a point in this history to show how the image looked in that point for comparison of changes. The stack can be compressed and will discard changes done during the development but will make it optimized. When you think you are done and have a good feeling of what you have done, just compress the history stack.

darktable ships with a bunch of modules and each of them is arranged into groups named module groups, these module groups are accessed via toggle buttons in the right panel just under the histogram. There are also two special module groups named active and favorites, where active will only show modules enabled in the history for the current image and favorites will only show modules selected as a favorite by you. This is done in the "more modules" at the bottom of the right panel, just click your favorite module until it gets a star displayed in front of the icon and you have chosen your favorite.

4

2.2.2. White balance

Module *white balance* controls, well, the white balance of the image. It's always enabled and reads its default values from metadata embedded in the image which will be set by your camera. The most common change here will be fine-tuning the white balance and this is done using the temperature in slider. Taking it down will make the color balance cooler and steering it up will make it warmer.

2.2.3. Exposure correction

Module *exposure* might be the simplest module of them all, fine-tune your exposure and use the histogram to guide you if you don't know what you are doing. You can also boost the black level by increasing the black but be careful, use small amount like 0.005 steps. If you feel lucky use the auto correction feature, this might fail for several technical reasons …

You can also click-and-drag in the histogram to change exposure / blackpoint.

2.2.4. Noise reduction

You have several options in darktable to reduce noise. You'll find them in the correct group. There is *raw denoise*, *denoising based on bilateral filter*, *non-local means*, and *equalizer* based on wavelets. Most convenient to use is probably denoising based on non-local means, as it allows you to treat color and luminance noise separately.

2.2.5. Fixing spots

Sometimes you will need to remove spots caused by sensor dirt. Module *spot removal* is at your hand and can also correct other disturbing elements like skin blemishes. If your camera has stuck pixels or tends to produce hotpixel at high iso values or longer exposure times, have look at module *hotpixels* for an automatic correction.

2.2.6. Geometrical corrections

Quite frequently you want to only show part of the captured scene in your image, e.g. to take away some disturbing features close to the edge. In other cases the horizon in your image maybe leveled or you observe some perspective distortions. All this can be corrected in module *crop and rotate*. If you need to correct typical flaws of a camera lens like cushion distortion, transversal chromatic aberrations or vignetting there is a separate module in darktable called *lens correction*.

2.2.7. Bringing back detail

Digital RAW images often contain more information than you can see at first sight. Especially in the shadows of an image there are lots of details hidden. Module *shadows and highlights* helps to bring these shades back into tonal values where they contribute to the visual impression. Structural details in fully blown-out highlights by nature of the digital sensor can not be recovered. However, you can correct unfavorable color casts in these areas with module *highlight reconstruction*.

2.2.8. Adjusting global contrast

Almost each workflow will cover as one basic element the adjustment of global contrast. You will want to control which of the available tonal ranges of your output medium are represented in your image. Darktable offers several alternative modules to take care of that. In one of them, module *tonecurve*, you adjust tonal values by constructing a gradation curve. Module *levels* offers a much simpler interface with just three sliders. In addition

there is module *zone system* which allows to control tonal values in certain zones, inspired by the work of Ansel Adams.

2.2.9. Enhancing local contrast

Local contrast enhancement will emphasize detail and clarity in your image; carefully used it can give your photograph the right pop. Darktable offers several modules for this task. Module *local contrast* is easy to handle with just a few parameters. A much more versatile but also more complex technique is offered by module *equalizer*. Have a look at equalizer's preset to get a feeling for how it works. Equalizer is darktable's "Swiss Army Knife" for many adjustments where spatial dimension plays a role.

2.2.10. Color adjustments

darktable offers many modules to adjust colors in your image. A very straightforward technique is implemented in module *color correction*. Use it to give your image an overall tint or to adjust overall color saturation. Module *color zones* offers a much finer control to adjust saturation or lightness and even hue in user defined zones. Also darktable's *tonecurve* module - in addition to the classical adjustment of tonal values - gives you fine control over the colors in your image. Finally, if you intend to convert your image into black&white, a good starting point with an easy to use and intuitive user interface is offered by module *monochrome*. Alternatively you might consider to use darktable's *channel mixer*.

2.2.11. Sharpening

If you start your workflow from an RAW image, you will need to have your final output sharpened. Module *sharpen* can do this with the classical USM (unsharp mask) approach, well known from more or less all image processing software around. Another very versatile way to enhance edges in your image is offered by module *highpass* in combination with darktable's rich set of blending operators.

2.3. Exporting images

Your changes to an image are not saved as in a regular image editor. darktable is a non destructive editor which means all your changes are stored in a database and the original image is untouched and therefore you need to export your images to get an "image" file to distribute with your special processing.

Images are exported from lighttable using the export module in the right panel. In general, export means: save my developed RAW image as a JPEG.

The export is modularized into storage and format. darktable is shipping several storage modules such as save on disk, picasa and flickr webalbum and more. Format modules are the actual image formats such as JPEG, PNG, TIFF, OpenEXR and more.

Select images on your lighttable that you want to export, then choose your target storage and format to use and set the maximum width and height image restraint. This means that none of the images will be bigger than any of the width/height restraint. Leave the width and height restraint as zero if you want the original resolution and hit the export button.

Chapter 2. Lighttable

The lighttable is where you manage all your images, ratings, export and much more …
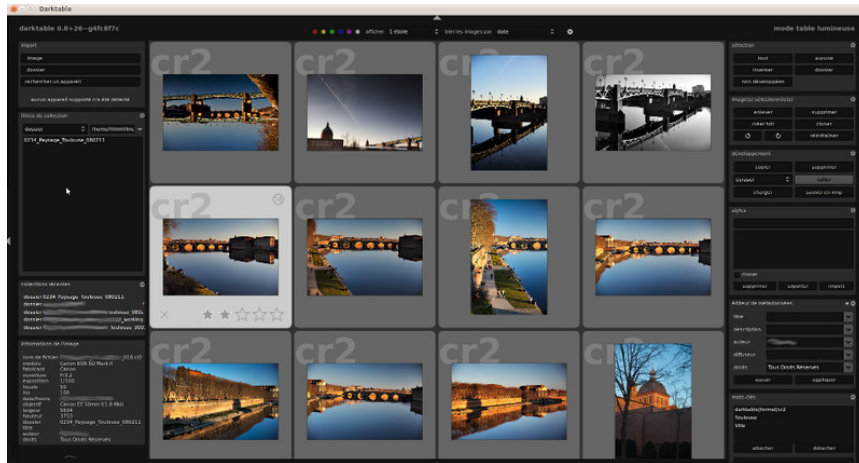
1. Overview

1.1. Lighttable

In the central view, your images are shown as thumbs; each thumb has a frame and while the mouse is over an image, the information of rating, color labels and if image has already been altered in darkroom 🖾 is shown. Also when the mouse hovers over an image frame the specific image metadata is shown in the metadata panel.

In order to locate where you are in your collection, you have read-only sliders at the extreme border of the window: left/right for your vertical position in the collection, top/bottom (when in zoomable lighttable mode, see below) for your zoom level.

While the mouse is over an image frame, you can access additional actions from your keyboard to apply on the image; here follows a table of key accelerators and assigned actions.

0 − 5	set the rating of the image; if an image has 1 star and you hit the *1* key, the image will be unrated. Pressing *r* rejects the image.
F1 − F5	set a color label for the image
ctrl-c	copy the history stack of the image
ctrl-v	paste the copied history stack on the image
d	open the image in darkroom view for developing
z	show a fullscreen view of the image

At the bottom you have an option to choose between zoomable lighttable view or filemanger view of the thumbnails.

Using the zoomable lighttable view, scroll with your mouse wheel to zoom in and out. Moving the mouse while *pressing the left mouse button* allows you to navigate through your collection.

While in filemanager you can use *ctrl-(mouse wheel)* to change the number of thumbs displayed in a row. That action is a key accelerator of the counter that can be found beside the filemanager option. Use your mouse wheel to navigate through your collection.

While in filemanager mode, by pressing ↑/↓ you can scroll (not select) up and down through your collection. In zoomable lighttable ←/→/↑/↓ allow you to move left/right/up/down through your collection. Pressing *g* goes to the top, *shift-g* to the bottom.

1.2. Filtering and sort order

The filtering and sort order of images on lighttable are accessed and changed in the top bar where you can select among a few predefined filters and ordering like date, color labels or rating.

1.3. Image grouping

Image grouping helps you to improve structure and clearness of your image collection as displayed in lighttable view.

You can combine images into a group by selecting them and then clicking button "group" in panel *selected image(s)* or by typing *ctrl-g*. Likewise you can remove images from a group by clicking button "ungroup" or typing *shift-ctrl-g*. Images that you generate by duplicating an existing one, will automatically constitute an image group. If you import images from file system or camera, images which have the same base name but different extensions will form a group (eg. IMG_1234.CR2 and IMG_1234.JPG). Images which are members of a group are labelled with a symbol "G" in their thumbnails.

With button G in the top panel of lightroom mode you toggle grouping mode on and off.

If grouping is *off* each image is displayed as an individual thumb.

If grouping is *on* images of a group are by default collapsed, which means they are represented by one thumbnail. The image you see is called the group head. If you press symbol "G" in the thumbnail of a group, only this group gets expanded; if another group was expanded at that time it gets collapsed. To collapse an expanded group again, just click on symbol "G" of its group head.

You can define which image constitutes the group head. When you are in expanded view of a group, just go to the thumbnail of the desired image and click on the symbol "G" within that thumb.

If you are in collapsed view and enter darkroom mode with an image group (eg. by double-clicking on the thumbnail) only the group head will be opened.

his is a convenient way to protect an existing history stack against unintentional changes. Suppose you have just finalized an image; all you need to do now is generate a duplicate and make sure grouping is switched on and the group is collapsed. Whenever you open the image group again in darkroom and do further changes, only the group head will be altered. The underlying duplicate remains unchanged. Please note that duplicating images only means that a second copy of your history stack and a second small XMP file is generated. There still is only one RAW file, so you don't waste a lot of disk space.

2. Panels

2.1. Import

In this panel you import your filmrolls. You can either import a complete folder by pressing 'import filmroll' or a single image 'import single image'. You can also import directly from a connected camera.

All your filmrolls are available in a list which can be filtered using the editbox to fast find the filmroll of interest. Double click on a filmroll in the list and it will open. You can also click the recent filmroll items to open the latest ones you have worked with.

2.1.1. Import from filesystem

A filmroll is the same as a folder in the filesystem. You can import a single image or a folder. If a folder is imported the images in that folder will show up in a filmroll whose name becomes the name of the imported folder. If single images are imported they will show up in filmroll 'single images'.
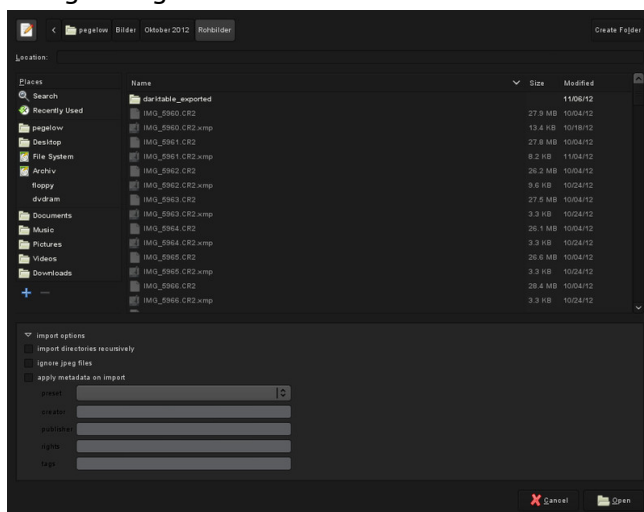
When you click on 'import single image' or 'import filmroll' a file selector dialog will pop up. Navigate through your filesystem and select the item you want to import. On the lower part of the dialog you can define some further import options.

As the name implies checking 'import directories recursively' will recurse from the currently selected directory through all subdirectories and import all images found. It is not recommended and wasting resources to this on an exhaustive list of images. darktable would generate thumbnails for all of them but in the end keep the recent ones in its cache. Better import images part by part as separate filmrolls.

Checking 'ignore jpeg files' is a good choice if there are jpeg images in the same folder that you do not want to process with darktable; eg. if your camera stores RAW+JPEG you often will only want to work further on your raws and let the jpegs alone. You can also apply some metadata already on import; see the section on 'Metadata editor' later in this manual for more details.

Importing a folder does not mean that darktable copies your images in another folder. It just means that the images are visible in lighttable and thus can be developed. It means that if you delete an image/a folder after having imported them, darktable will not be able to access them anymore. Importing an image/folder in darktable is not a backup of your filesystem! Moreover darktable does not currently watch changes in filesystem. Thus if on

10

your filesystem you add an image after having imported that folder in darktable, the new image will not be shown until you explicitly import it.

2.1.2. Import from cameras

When a camera is detected, it will show up in the device panel; if you hover your mouse over the camera tab label a tooltip will pop up with some information about the camera such as model, firmware version and a lot more...

Depending of the support of your camera, buttons with actions will be available such as import images and tethering.

Import images

This will bring up an import dialog, showing the images on camera that can be selected for import into a filmroll in darktable.

Tethering

Tethering is used to integrate darktable with your camera. While you take images with your camera, they will automatically be imported into darktable so you can review the result of the shoot. You can also setup remote capture jobs, controlling number of images and time between captures along with camera settings such as exposition time, f-length and more.

If supported by your camera, tethering will take you into capture view for tethered shoot. Read more about tethering in chapter 4.

2.1.3. Supported file formats

darktable has its main focus on managing and developing camera RAW files. Therefore a huge number of different file formats from various camera manufacturers is supported. In addition darktable can read specific *low dynamic range* and *high dynamic range* images - mainly for data exchange between darktable and other software.

Camera RAW files

darktable reads RAW files using two open source libraries: RawSpeed (developed by Klaus Post) and - failing that - with LibRaw. The number of supported cameras and file formats is constantly increasing. It would be beyond the scope of this manual to give an exhaustive list. Most modern camera models are supported and new ones tend to get in very quickly. darktable does not decode images from cameras with non-Bayer sensors (e.g. Fuji X-Pro1 or Sigmas with Foveon X3 sensor).
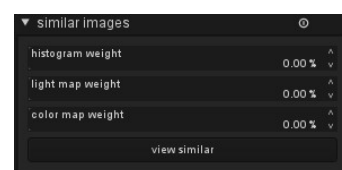
LDR image files

darktable reads "ordinary" images in JPEG, 8-bit/16-bit PNG and 8-bit/16-bit TIFF format.

HDR image files

darktable reads high dynamic range images in OpenEXR, RGBE and PFM format.

2.2. Similar images

Similar images will match and collect images with similar characteristics to the selected one and show them on lighttable.

2.3. Collect images

With this panel, you can manage your collection (current view) by filtering images with various criteria. The left combo box lets you choose from different filters, such a 'filmroll', 'historystack', 'date', 'camera' and much more. The table below lists all available data for the selected filter. You can then select the data of your choice by double-clicking.
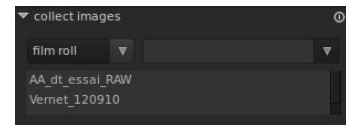
2.4. Keywords

Keywords modules shows a hierarchial tree of your tags. You can arrange tags within this tree using drag'n'drop a tag onto another.
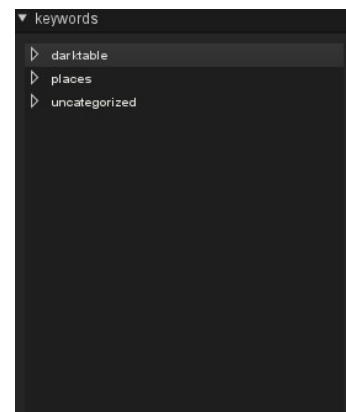
2.5. Recent used collections

This panel keeps tracks of your latest collections you have used, so you can jump to the latest used without remebering what rules was specified in the collection.
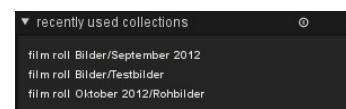
2.6. Metadata

This panel shows embedded metadata of an image, read from EXIF data stored in your pictures files. When hovering with your mouse over the thumbnails, darktable will update its view and display metadata of the image currently under your mouse cursor. This panel is also available in darkroom, tethering and map view.
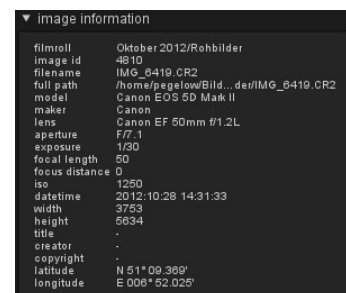
2.7. Select

This panel allows to quickly select images, according to various criteria.

select all

Select all images in the current view (collection), with respect to filters.

select none

Select no image.

invert selection

Select all images that are not currently selected.

select film roll

Select all images that are in the same filmroll as the currently selected images.

select untouched

Select all images that have not been developed yet.

2.8. Selected image(s)

This panel allows to perform some actions on select-
ed images.

remove

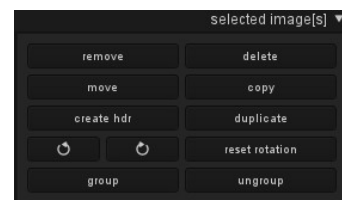Remove the selected images from the darktable database. Those images will not be shown in lighttable anymore, but remain on filesystem. As darktable stores XMP files with your development parameters on disk, you can later fully reconstruct your work by just re-im-porting the images. When backing up your RAWs make sure to also save the XMP files!

delete

Physically delete selected images from filesystem. See also preference option 'ask before erasing images from disk'; if this option is not active, darktable will delete the file(s) with-out further question! This is not reversible and will also erase your development work an these images.

move

Physically move selected images (parent file plus accompanying .xmp sidecar file) to an-other filesystem folder.

copy

Physically copy selected images (parent file plus accompanying .xmp sidecar file) to an-other filesystem folder.

create hdr

Create a high dynamic range image from the selected images and store it as a new source file in DNG format. Images need to be properly aligned.

duplicate

> Create a virtual copy of selected images within darktable but not physically on your filesystem. It allows to test and try different developments for the same image for example. Duplicate images will share the same parent input file but each have their own .xmp sidecar file.

rotation

> Perform a counter-clock wise/clockwise rotation on selected images. The third button allows to reset rotation to EXIF value.
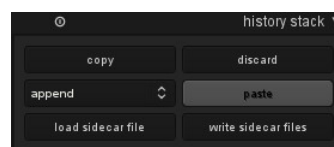
group

> Create a new group from selected images.

ungroup

> Remove selected images from the group.

2.9. History stack

> This panel allows to manipulate development. For each image, development is written in a sidecar file (.xmp) and is fully non-destructive.

copy

> Copy the history stack (development) of the first selected image.

discard

> Physically delete the history stack (development) of the selected images. Beware, this action can not be undone!

overwrite/append

> Describes how a new history stack must behave if pasted on an image having already an history stack. 'Overwrite' will delete previous history stack whereas 'append' will add the two history stacks.

paste

> Paste a previously copied history stack on another image. This button is greyed out until a history stack is copied from another image.

load sidecar file

> Prompt a dialog box to allow you to select a .xmp file, thus loading an history stack that you will paste (thanks to 'paste' button) on images.

write sidecar files

> Write .xmp sidecar files for all selected images. By default darktable will do this automatically for you whenever you work on an image and change history stack. You can deactivate automatic sidecar file generation in the preferences dialog; this is not recommended.

2.10. Styles

This panel lets you use a powerful functionality of darktable: to store and apply a history stack. Styles are created in darkroom mode thanks to a button placed below the history stack. Then they are managed within this panel, that allows you to apply and delete a style.
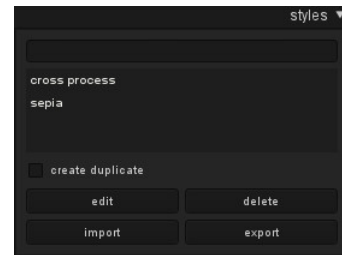
create duplicate

When applying a style on selected images, activating this box lets you create a duplicate before applying the styles. De-activate this option if you want to try various styles; this will avoid that you create multiple duplicates.

edit

Styles are a collection of history stack items. After pressing edit you are prompted with a dialog in which you can deselect specific items. Check option "duplicate" if you want to create a new style instead of overwriting the existing one; you need to give a new style name in this case.

delete

This will delete the selected style without further question.

import

You can import a style which has been previously saved by darktable. darktable stores its files in xml files with the extension .dtstyle.

export

This option saves a selected style to disk as a .dtstyle file. This allows to share styles with other users.

2.11. Geotagging

Using this panel you can import and apply GPX track data on selection of images, you also have the ability to add an time offset to GPX tracks for correction of time differences between your camera and gps receiver. You can also manually geotag your images within the 'Map' view. See chapter 5 for more details.
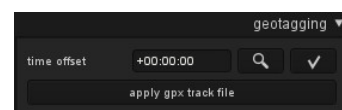
2.12. Metadata editor

Edit some metadata of your image. You can define your own presets if you want to apply specific settings frequently.
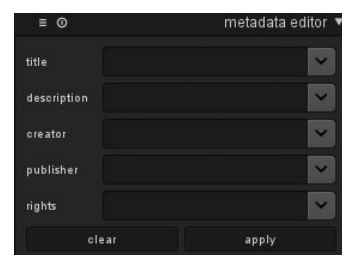
clear

> Delete existing metadata from the selected image(s).

apply

> Apply new settings as defined in the fields above to the selected image(s).

2.13. Tagging

This panel allows to manage tags for your images. Those tags are currently stored in sidecar files (.xmp) and within the darktable database for a faster access. The panel is divided into two parts: the upper part contains the tag(s) currently set for the image either under mouse (if mouse is over an image) or selected (if mouse is outside the lighttable). The lower part contains all tags available; those can be filtered thanks to the upper text box.
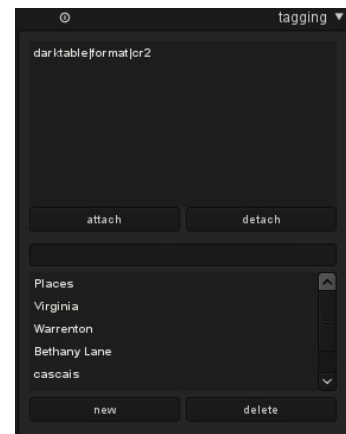
attach

> Attach the selected tag(s) from the list below to all selected images.

detach

> Detach selected tag(s) from the list above from all selected images.

new

> Create a new tag for the list.

delete

> Delete a tag from the list.

2.14. Export selected

Each workflow has to end by this panel: the exportation of your developed images. Currently, you can export either to a file on your disk, or to various on-line storages. *TIP*: you can use *ctrl-alt-e* from within darkroom mode to export.

target storage

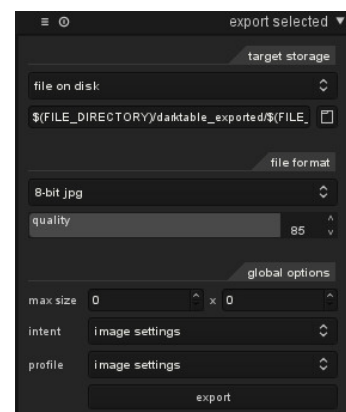Where to store your selected images. Different back-ends are implemented, including file on disk, a LaTeX book template and various web albums. Depending on the selected target you will be asked to give additional information like filenames, account name and password.

filename template

You can define the filenames that darktable generates for export. Several pre-defined variables can be used as placeholders:

$(ROLL_NAME)	roll of the input image
$(FILE_FOLDER)	folder containing the input image
$(FILE_NAME)	basename of the input image
$(FILE_EXTENSION)	extension of the input image
$(SEQUENCE)	a sequence number within export job
$(YEAR)	year at date of export
$(MONTH)	month at date of export
$(DAY)	day at date of export
$(HOUR)	hour at time of export
$(MINUTE)	minute at time of export
$(SECOND)	second at time of export
$(EXIF_YEAR)	exif year
$(EXIF_MONTH)	exif month
$(EXIF_DAY)	exif day
$(EXIF_HOUR)	exif hour
$(EXIF_MINUTE)	exif minute
$(EXIF_SECOND)	exif second
$(STARS)	star rating
$(LABELS)	colorlabels
$(PICTURES_FOLDER)	pictures folder
$(HOME)	home folder
$(DESKTOP)	desktop folder

output directory

Pressing button ▣ will open a dialog to select the parent directory for export.

file format

darktable knows various file formats to export to. For some of them you need to additionally define your desired bit depth. If you export to a jpeg file you can additionally define an output quality. Higher values will lead to larger file sizes.

max size

Set the maximum width and height of your output images in pixels. Set both to a value of "0" to export with full resolution. darktable currently can only do down-scaling; the maximum output resolution is defined by the parent image.

intent

This option lets you define the intent, i.e. the way darktable will deal with out-of-gamut colors. See module "output color profile" for a more detailed description of the available options. Select "image settings" if you want the settings in module "output color profile" of the individual images to take precedence.

profile

This defines the output color profile. Select "image settings" if you want the settings in module "output color profile" of the individual images to take precedence.

export

Pressing this button will start a background job to export all selected images. A bar at the bottom of the left side panel will display the progress. Whenever a file is successfully exported a notification message will pop up. You may press on the pop-up to make it disappear.

Chapter 3. Darkroom

The darkroom view is where you develop your image...
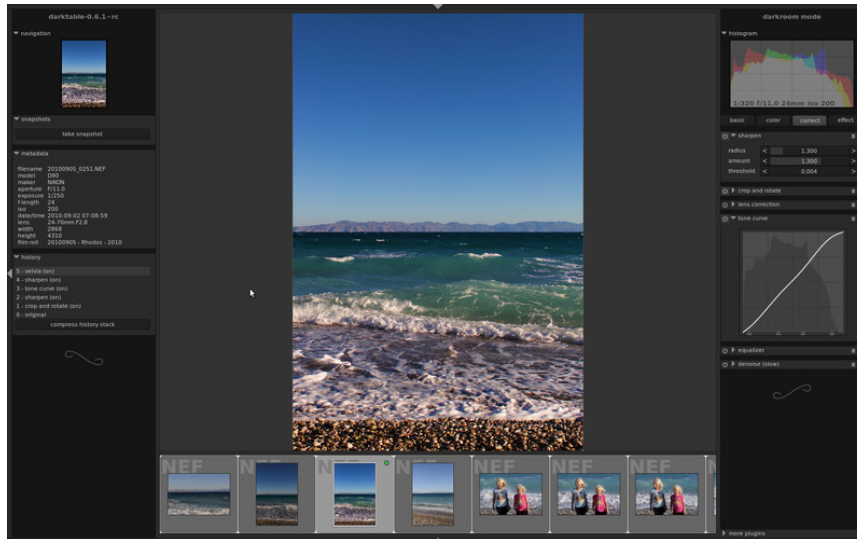
1. Overview

In darkroom you are in develop mode for the specific image; here you have a lot of tools (named modules) to develop you image.

On the left side you have the navigation, snapshots and history panels, each of them being described below. In the right panel you can see the histogram and then a list of available modules to use when working with your image. At the bottom of the right panel you can enable/disable view of individual modules.

Since darktable 0.6.1, modules are grouped into four groups: basic, color, correct, effect. However, you can view all modules by selecting no group.

For those interested in how it works, modules are applied on images following a bottom-to-top order. That means, that if you activated the view of all modules (by selecting no groups), the first module to be applied will be "import a RAW" at the bottom of the right panel, whereas the last one will be "splittoning" at the top of the right panel. That order can not be changed by user, so modules are always applied following the same order.

You can use *middle-click* to zoom 1:1. A double *middle-click* takes you to 2:1.

You can export directly from the darkroom by using *ctrl-e* shortcut. Export parameters are those currently selected in lighttable.

2. Filmstrip

The filmstrip is used to quickly switch between images to develop in darkroom view, the images viewed are the same as the ones in the lighttable view.

The filmstrip can be switched on and off using *ctrl-f* shortcut. You can scroll with your mouse to quickly navigate through the images. The height of the filmstrip panel can be changed by dragging the corresponding marker.
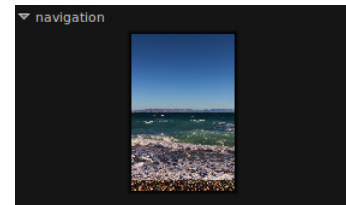
3. Panels

This section contains documentation for panels that are specific to the darkroom view.

3.1. Navigation

The navigation panels shows a full preview of your image and the current viewport of zoom. Drag the rectangle around to pan the zoomed-in view.

3.2. Snapshots

You can take a snapshot of what is displayed in the center view; this snapshot can then be selected and is overlayed in the current center view to help you with a side by side comparasion (left: snapshot, right: active) when you are tuning parameters of a module. In the active view, history remains usable in order to compare with the snapshot.
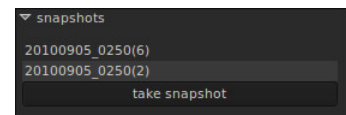
You can control the split view by moving the splitline back and forth and if you hover mouse over the split line, a small rotation icon will appear on the center of the line, click it to change to a horizontal split view.

3.3. History

The history stack lists every change of state (activate/de-activated) for all modules. Here you can select a point in stack to return to that point of development history. If you then activate a new module, all modules above the current point will be discarded.
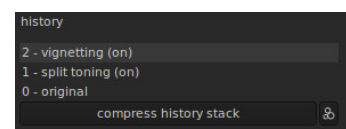
Be cautious: also activating any action by key accelerators will discard all modules above the current one, e.g. if you press "o" to see over/under exposed areas. The same happens if you press the according toggle button ▣. It is easy to lose all development work on an image this way.

Hitting "compress history stack" generates the shortest historystack that produces the current image, i.e. suppressing all de-activated modules. This will also discard all modules above the currently selected one.

The button to the right ⚇ lets you create a new style by popping up a dialog window. Use the first line to name your style, the second to add a searchable description. The "advanced presets" allows you chose which modules are to be taken into account for the styles.

3.4. Color picker

Within the global color picker panel you can take colour samples from your image, display their values in multiple ways and compare colors from different locations. Color picker is activated by pressing icon ⬈. Various parameters control how color picker works. All settings remain in effect until you leave darkroom mode.
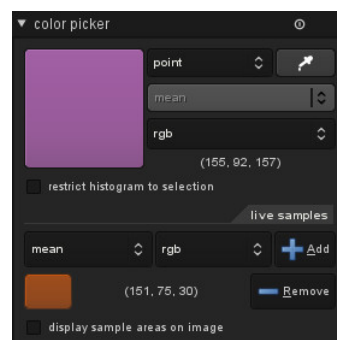
Besides the global color picker, which we describe here, there are local color pickers in various modules (eg. tonecurve). Global and local color pickers are different. The global color picker works in monitor color space; it takes samples after the complete pixelpipe has been processed and colors are transfered into RGB suited for your monitor. The local color pickers run in the color space of the individual module, which is most often Lab; they only see input and/or output colors of that specific module within pixelpipe.

Color picker can be run in point or area mode. When in point mode only a small spot under your cursor is taken as a sample. In area mode you can draw a rectangle with your mouse; darktable samples the area within that rectangle. The combobox to switch between point and area mode can also be used to toggle the mode of local color pickers.

If samples are taken in area mode, darktable will calculate mean, min and max color channel values. A combobox allows you to select which of those are displayed. For obvious reasons mean, min and max are identical in point mode.

A color swatch representing the sampled point or area is displayed. Numerical values are shown as well. As said before global color picker works in monitor RGB color space. You can also let darktable translate these numerical values into Lab color space. Beware that Lab values are approximated here; darktable does not take your specific monitor profile into account.

On activation of checkbox "restrict histogram to selection" only the values of your selected area or point are taken into account by the histogram on top of right panel. This is a useful tool if you want to know which tonal values are present in a specific area.

Picked colors in area or point mode can be "stored" as live samples by pressing the "Add" button. darktable shows a color swatch and numerical values for each stored sample. You can once again select which numerical value (mean, min, max) is to be displayed and if this is to be done in RGB or Lab color space.

Newly created live samples are not locked. If you change the parameters of your image these changes will be reflected in your stored samples. Use this if you want to get an overview how changing a parameters has effect to different parts of an image. If you click on the color swatch of a stored live sample it gets locked - a lock symbol is displayed. Further changes of parameters will not affect this live sample any longer. You can for example take two live samples from the same location and have one of them locked. This way you can keep the locked one as a reference sample for comparison.

Live sample locations are indicated in your image if you check option "display sample areas on image".

23

3.5. Histogram

Show a histogram of the developed image; you can change exposure by left mouse click in the right side of histogram and drag right for increase or left to decrease the exposure. The same goes for the black level but you left click in the left part of histogram. You can switch between linear and logarithmic histogram view by toggling a corresponding button. In its default state curves for all three color channels RGB are displayed; toggle one of the colored squares to enable/disable a specific color channel.
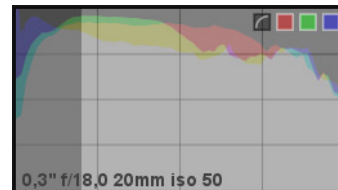
3.6. Modulegroups

The modulegroups button bar give you quick access to the set of modules available in darktable, each one categorized into groups, see the specification below.
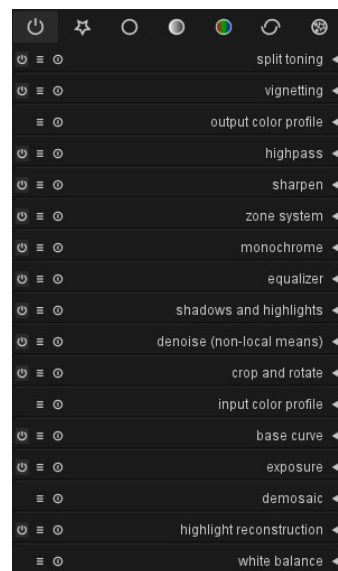
Here follows a description of the module groups available:

	Active	Shows the list of active modules for the currently processed image
	Favorites	Modules you have marked as favorites using "more modules" will show up here
	Basic	The basic group of modules that is frequently used when processing raw, such as exposure, temperature etc.
	Tone	These modules does only work with the image tone, like levels, tonemap etc.
	Color	Modules that just affects colors, such as color correction, vibrance etc.
	Correction	Modules that makes corrections to the image, denoise, ca correction etc.
	Effect	Modules with a more artistic output, vignetting, softening etc.

Clicking on one of the group symbols will show the modules in that group. If you once again click on the symbol grouping will be de-activated. Now you see all non-hidden modules in one long list. This list shows the sequence in which modules are applied from bottom to top. As a general principle darktable applies modules in a pre-defined sequence.

The long list also helps you to figure out in which color space a specific module acts. In fact there are only three modules which cause a switch from one colorspaces to another: *demosaic*, *input color profile* and *output color profile*. Here is an overview:

up to *demosaic*	Image is in raw data format with only latent colors. Each pixel carries lightness and color information for only one base color. Some of the modules in this part can also act on non-RAW input images in RGB format.
between *demosaic* and *input color profile*	Image is in RGB format within the color space of the specific camera or input file.
between *input color profile* and *output color profile*	Image is in Lab format. This is a very huge universal color space which covers all colors visible to the human eye (and even more). As darktable processes images in 4x32-bit floating point buffers, we can handle the Lab color space without risking banding or tonal breaks.
after *output color profile*	Image is in RGB format as defined by the selected display or output ICC profile.

3.7. More modules

More modules at the bottom of right panel is used to hide modules you don't use. As a default only standard modules are shown to the user, so if you are missing anything you should check here and show it.
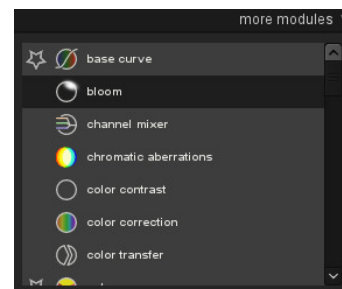
Each modules is shown as a small icon and the module name. Left-click with your mouse to toggle a module's status: visible, hidden or favorite. Favorite modules are indicated by a star in front of the icon and will additionally to their normal module group show up in module group "favorites". This is a good way to get fast access to modules that you use very frequently. Visible modules are indicated in the list by a light grey background. Hidden modules have a dark grey background; they will not display any controls whatsoever.

4. Modules

Each correction, enhancement or effect is implemented a module.

This section includes documentation of each module and its specific parameters. A module consists of an expander bar ⟨image: toolbar with sharpen⟩ where the ⏻ icon is used to turn the module on or off. Then there is the name of the module; press this to expand the module to see the parameters user interface. Icon ☰ is the preset icon to access the available presets stored for the module or to create new ones. The ⟲ icon is used to reset the module parameters to their default values.

4.1. Interacting with modules

The most frequently used control elements are sliders, comboboxes and curves.

For each slider, you can interact in four different ways, depending on the level of control you need.

1. Triangular marker

   You left-click with your mouse on the triangular marker and move it around.

2. Mouse wheel

   Hover a slider with your mouse, then use your mouse wheel to adjust by step the value.

3. Right-click

   Your mouse being over a slider, right click gives you a multi-functional pop-up below the slider for fine control with your mouse and for entering numerical values with your keyboard. A bent line that extends from the triangular marker moves as you move your mouse. The closer your mouse pointer is to the triangular marker the coarser the control; the further away from the triangular marker the finer is your control. Left-click with your mouse to accept the current value and go back to normal control. You can also directly type in the desired value with your keyboard; there is no cursor, you just start typing.

   darktable's innovative way of value input: coarse and fine value adjustments in one single control element.
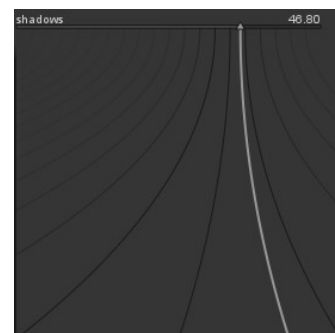
4. Double-click on parameter label

   In some modules, you can double-click on a parameter label to reset its value to default.

Clicking on a combobox will open a list of available options. Click on the item you want to select

Some modules are controlled by adjusting curves. More detail is given later in this chapter when the respective modules are explained.

4.2. Blending operators

Blending operators are a functionality that works on the incoming image information and the output of the current module. Not all modules support blending because it would not work and break the use of the module.

Modules that do support blending, inherit additional controls at the bottom of module user interface to control parameters of the blending.

4.2.1. Usage

blend mode

There are several modes implemented and more can be added. For now all the common ones are there and you will recognize a few of them from other imaging software. A good introduction on many common blend modes is given in The Gimp Manual (Chapter 8.2, "Layer Modes"). Therefore we only discuss a few blend modes here in more detail.

off

If set to "off" no blending will be done and blending related controls are hidden.

normal

This will probably be the most used blend mode, it does just mix input and output and will be used to control the amount of a module's effect.

inverse

This blend mode acts similar to blend mode "normal", only that the role of input and output of the current module are reversed. Can be used together with conditional blending (next chapter) to get the effect of an inversed blend mask.

unbounded

This blend mode acts similar to blend mode "normal", only that input and output data are not clamped to a certain min/max value range. This already tells you that all other blend modes clamp their input and output. In some cases (e.g. highly color saturated extreme highlights) it is important to let unbound values travel through the pixelpipe in order to properly deal with them at the right place (e.g. in module "output color profile").

lightness

This blend mode mixes lightness from input and output image. Color data (chroma and hue) are taken unaltered from the input image.

chroma

This blend mode mixes chroma (saturation) from input and output image. Lightness and hue are taken unaltered from the input image.

hue

This blend mode mixes hue (color tint) from input and output image. Lightness and chroma are taken unaltered from the input image. For modules which drastically modify hue (e.g. generating complementary colors) this blend mode can result in strong color noise.

color

This blend mode mixes color (chroma and hue) from input and output image. Lightness is taken unaltered from the input image. For modules which drastically modify hue (e.g. generating complementary colors) this blend mode can result in strong color noise.

coloradjustment

Some modules act predominantly on the tonal values of an image but involve some adjustment of color saturation as well, e.g. module *levels* and *tonecurve*. This blend mode takes the lightness only from output data and mixes colors from input and output. This way you can control the module's automatic color adjustment.

opacity

This slider controls the amount of blending. A value of 100% gives the module's full effect into the pixelpipe; a value of 0% means no effect of this module. With blend mode Normal and a proper selection of opacity you can therefore control the strength of a module's effect.

blend

Most modules offer additional control on blending by a user-defined blend mask. In order to activate these controls you need to switch from "blend uniformily" to "blend only, if..". More details on conditional blending are given in a separate chapter.

4.2.2. Examples

Texturing an image

The watermark module supports svg files with embedded images and thanks to it we can easily texture an image using the watermark module in conjunction with a blending operator.

Gritty details

When blending operators were introduced into dark-table, a new module named highpass was added. It produces a highpass filter of the image to be implicitly used with blending. It allows to produce a gritty detailed image and is a widely used workflow with other imaging softwares.
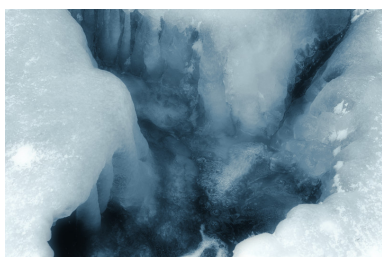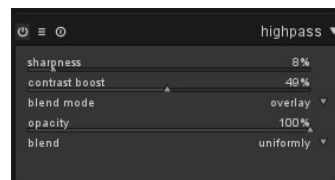
This is the original image, pretty heavy processed: first black and white, then some blue splittoning but as you see it lacks pop in details and is a bit out of focus...
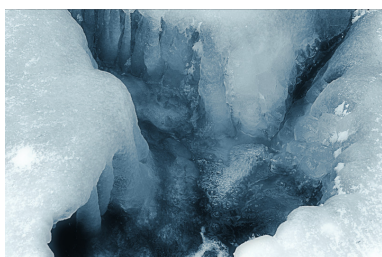
Here we applied the highpass filter with the values shown in the images above of the module userinterface; you can now see that the details are very boosted and we got a really gritty detailed image.

4.3. Conditional blending

Conditional blending offers a fine-grained control as to what extent blending is applied to different pixels of an image. It does so by generating an intermediate blend mask, constructed with user defined parameters. These parameters act in color coordinates, not in geometrical coordinates.

That said, conditional blending is a powerful tool with a certain level of complexity.

4.3.1. Working principle

For each data channel of a module (Lab, RGB) and additionally for several virtual data channels (e.g. hue, saturation) users can construct a per-channel opacity function. The function has the form of a trapezoid. Depending on the data channel's value this function determines a blending factor between 0 and 1.

Each pixel of an image thus has different blending factors for each of its data channels (real and virtual). All blending factors are finally pixel-wise multiplied together with the value of the global opacity slider (see chapter "Blend Modes") to form a blend mask for the image.

If for a given pixel the blend mask has a value of 0, the input of the module is left unchanged. If for a pixel the blend mask has its maximum value of 1 (or 100%), the module has full effect.

4.3.2. Usage

When conditional blending is activated with option "blend only, if.." an additional set of controls consisting of several notebook tabs is shown.

Notebook tabs

Each notebook tab represents a data channel - real or virtual. Modules acting in Lab color space have data channels for L, a, b, C (chroma of LCh) and h (hue of LCh). Modules acting in RGB color space have data channels for g (grey), R, G, B, H (hue of HSL), S (saturation of HSL), and L (lightness of HSL). Consult http://en.wikipedia.org/wiki/Color_space for a deeper look into different color spaces.

Within each notebook tab two color channel sliders are given; one for the input data that the module receives; and one for the output data that the module produces prior to blending.

Control buttons

Control buttons help you when designing a blend mask. With the eye button ⬤ you can temporarily deactivate the blend mask; blend mode and global opacity slider remain in effect. Switch this button on and off to see if the module is acting on the image as intended. With the mask button ⬛ you can directly see the blend mask as a yellow overlay over a black and white copy of the image; the stronger the yellow color, the higher the blend mask value. With the color picker button ⬈ you can select a probe from your image. The

corresponding values for the real and virtual data channels are displayed in each color channel slider. With the reset button ⊙ you can put all conditional blending settings back to their default state.

Color channel sliders

With the color channel slider you construct a trapezoidal opacity function. For this purpose there are four markers per slider. Two triangles above the slider; they mark the range of values where opacity is 1. Two triangles below the slider; they mark the range values where opacity is zero.

The filled triangles indicated the closed (mostly narrower) edge of the trapezoidal function. The open triangles indicate the open (mostly wider) edge of the trapezoidal function.

In its default settings the markers are at the extreme positions left and right of the slider. The sequence of the markers always remains unchanged: they can touch but they can not switch position.

A polarity button right to the slider defines what pair of triangles defines the closed edge of the trapezoid: positive for the upper triangles (corresponding to opacity 1); negative for the lower triangles (opacity 0). This allows for two principle types of trapezoidal functions which are represented graphically in the following images.
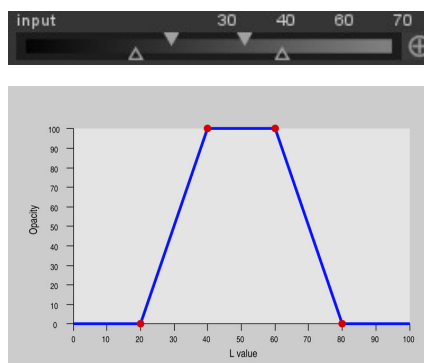
A trapezoidal function that selects a narrow range of values for blending.
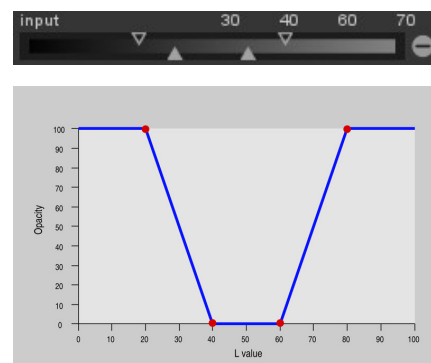
A trapezoidal function that excludes a narrow range of values from blending.

Mask blur

The mask blur slider applies a gaussian blur to the final blend mask. The higher the radius, the stronger the blur; set to 0 for an unblurred mask. This option is useful to get a softer transitions between blended and unblended parts of an image and avoids artifacts.

4.3.3. Examples

Colorkey effect

To make a colorkey effect for this red poppy blossom, we need to apply module monochrome to all parts of the image except of the saturated red petals.

These settings in hue channel construct a blend mask that excludes the red petals. The small white bar in the gradient is generated by color picker taking up a probe from within the petals.
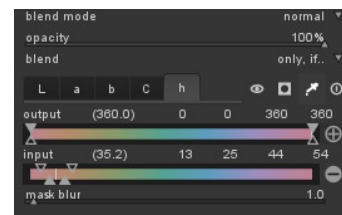
The resulting blend mask.
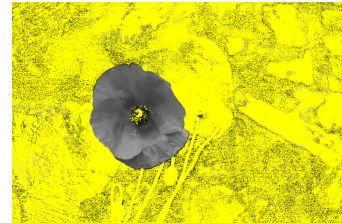
The final image after module monochrome is applied.
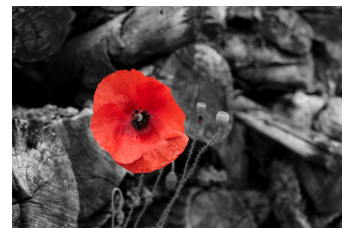
4.4. Presets

Presets are stored configurations for the module parameters. Some modules come with internal, pre-defined presets. You can also define your own ones. Both internal and user-defined presets are accessible through clicking the ▤ icon. A list of available presets is shown for your choice; the activated preset is in bold text.

The preset system does also support automatic preset selection based on image data such as focal length, iso, camera model and many more fields.

4.4.1. Creating a new preset

Configure the specific module of choice as you like then you click the ▤ icon and select store new preset menu choice. Following dialog will be shown to configure the preset:
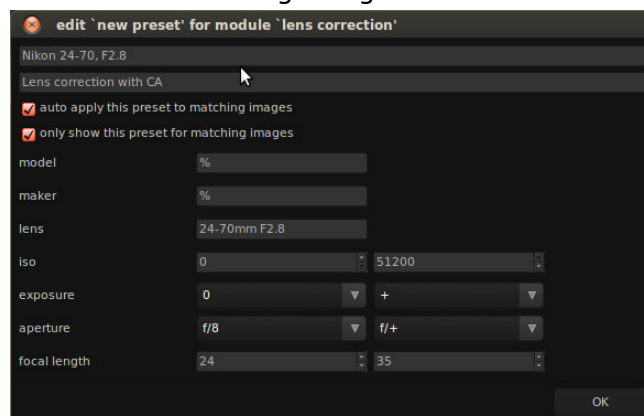
The first two fields are used for setting a name and a description of the preset.

In the example above we have also checked the auto apply option. This brings up additional selection fields where you can define a filter, that is used to decide if the preset should

be applied the image or not. The example dialog above does set up following rules, if lens name matches and aperture is greater or equal to 8 and focal length is in range of 24 to 35mm the preset will be automatically applied. Also the second checkbox is clicked, so that this preset will only show in the preset list, if the image does matches this rule.

darktable finds these data in the EXIF information of your images. If you want a preset to be applied to all images from a specific camera, then leave all fields at default values, only lookup model string in metadata panel and enter it into the model field.

4.5. Basic group

The basic group of modules contains the modules for basic develeopment. These are ones you probably will use most often, such as exposure, temperature etc.

4.5.1. Crop and rotate

Overview

This module is used to crop, rotate and correct perspective of your image. It also includes many helpful guidelines that assist you using the tools.
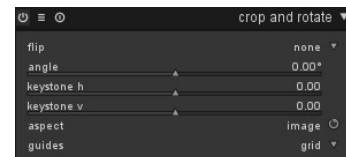
Usage

Whenever the user interface of this module is in focus, you will see the full uncropped image overlayed with handles and guiding lines.

First off, select what aspect ratio you want and size the crop bounds by dragging border and corner handles. Use the button right to the aspect box, to swap between portrait and landscape of the cropping rectangle. You can move around the crop rectangle by holding down left mouse button and move around. When you are done and want to execute the crop, just give focus to another module. You can at any time change your crop area by just giving back focus to this module.

angle

This is the rotation angle to correct, for example, a non-leveled horizon. You can either set a numerical value or use your mouse directly on the image. To use your mouse, right-click, hold it down and draw a line along the horizon; as soon as you release the mouse button the image is rotated so the line you drew matches the horizontal axis.

keystone h and v

This tool is used to correct any perspective distortions in your image. Useful when you shoot an high building from ground, aiming upwards with camera with short focal length for example.

aspect

Here you can change what aspect ratio you want to have on the result, thus constraining the ability to drag and crop rectangle out of the aspect ratio of your choice. Many common numerical ratios are pre-defined. You also select any other ratio after opening the combobox and typing it in the form of "x:y". A few special aspect ratios deserve explanation:

free   free forming the rectangle without any ratio restrictions

image   this option constrains the ratio to be equal to image ratio

| *golden cut* | this option constrains the ratio to be equal the golden number |
| *square* | this option constrains the ratio to be 1 |

guides

Many self-explaining guides are available to help you compose your image.

Examples
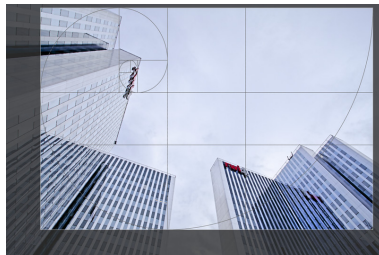
A cropped image in center view when module crop and rotate is in focus. The cropped area is visible as well as some guiding lines.

4.5.2. Overexposed

Overview

This module is a useful feature that displays pixel outside dynamic range.
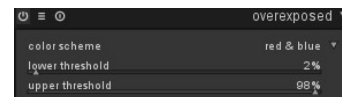
Usage

Overexposed is used to visualize pixels that are above and/or below a user-defined threshold. You can activate this module the usual way by pressing the activate button ⏻. You can also press icon ▣ in the lower panel as a shortcut or activate it with keyboard accelerator "o".

color scheme

In its default color scheme this module will display deep shadows in blue and highlights in red. These colors will be easy to identify in most cases. In some cases you might want to change that to black & white or purple & green, eg. if you experience overexposed highlights in red blossoms.

lower threshold

Set lower threshold, expressed as a percentage of the maximal brightness.

upper threshold

Set upper threshold, expressed as a percentage of the maximal brightness.

4.5.3. Shadows and Highlights

Overview

The shadows and highlights module allows to adjust the tonal range of darker parts of an image (shadows) and lighter parts (highlights); it can bring back details in shadows and highlights by enhancing local contrast.

Usage

shadows

This slider controls the effect on shadows; positive values will lighten up shadows while negative values will darken them.

highlights

This slider controls the effect on highlights; negative values will darken highlights while positive values will lighten them up.

soften with

This combobox choses the underlying blurring filter, gaussian or bilateral. Try bilateral filter if you experience halos with gaussian blur.

radius

This slider controls the radius of the involved blurring filter. Higher values give softer transitions between shadows and highlights but might introduce halos. Lower values will reduce the size of halos but may lead to an artifical look. As said, bilateral filter is much less prone to halo artifacts.

compress

This slider controls how strong the effect extends to midtones; high values reduce the effect to the extreme shadows and highlights; low values cause strong adjustments also to midtones. You normally only need to touch this parameter if you want to limit the effects to the extreme shadows and highlights; increase the value in this case. At 100% this module has no visible effect any longer as only absolute black and abolute white are affected.

shadows color adjustment

This slider controls the color saturation adjustment made to shadows; high values cause saturation enhancements on lightened shadows; low values cause desaturation on lightened shadows. It is normally safe to leave this at its default of 100%. This gives a quite natural saturation boost on shadows - similar to the one you would also expect in nature if shadows would receive more light.

highlights color adjustment

This slider controls the color saturation adjustment made to highlights; high values cause saturation enhancements on darkened highlights; low values cause desaturation on darkened highlights. Often highlights do not contain enough color information to give convincing colors when darkened. You might need to play a bit with this parameter in order to find the best fitting value depending on your specific image; but be aware that sometimes results still might not be fully satisfying.

Examples
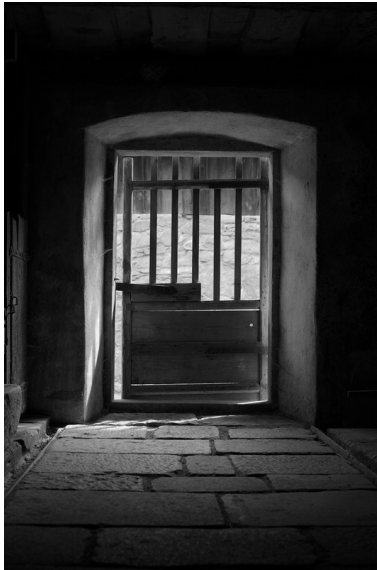
Original image exposed for the outer sunlit wall to avoid clipped highlights. As a consequence the interior of the barn has pitch black shadows.

Shadows get lightened; highlights are untouched; overall effect toned down a bit by blend mode "normal" and an opacity of 65%.

Resulting image.

4.5.4. Basecurve

Overview

darktable comes with general enhanced basecurve presets for several models. These are automatically applied to raw images for better colors and contrast.
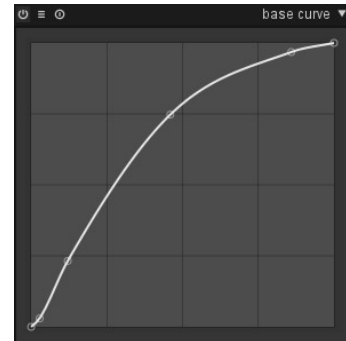
Usage

The basecurve works in camera linear RGB. It is defined by two or more nodes. You can click on any node and drag to adjust; you can also create additional nodes by clicking on a curve segment between two nodes. In order to remove a node drag it outside of the diagram area.

4.5.5. Exposure

Overview

This module is used to tweak the exposure. It is directly linked to the histogram panel. Indeed, if you correct exposure graphically, using the histogram, you automatically activate the exposure module. The histogram simply acts as a view for the exposure module.
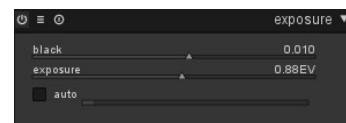
Usage

This module is responsible for one of the most basic steps in each raw image development. An exposure adjustment value allows you - within certain limits - to correct for under- or overexposure. A shift by 1EV is equivalent to a change of exposure time by a factor of 2.

Positive exposure corrections will make the image brighter. As a side effect noise level gets higher. Depending on the basic noise level of your camera and the iso value of your image, positive exposure compensations with up to 1EV or 2EV still give reasonable results.

Negative exposure corrections will make the image darker. Given the nature of digital images this can not correct for blown out highlights (see also module "highlight reconstruction").

A black level adjustment is a basic tool to increase contrast and pop of an image. The value defines at what threshold dark gray values will be cut off to pure black. Use with care as these dark gray values can not be recovered in other modules further down the pixelpipe.

black

Adjust the black level.

exposure

Adjust the exposure correction [EV].

auto

Calculate a correct exposition for the rectangular view appearing in the centre of the image. You can draw your own selection using your mouse. An adjustment slider right to the auto exposure checkbox lets you define what percentage of bright values to be clipped out in the calculation.

4.5.6. Demosaic

Overview

This module allows you to control how the demosaic is processed.
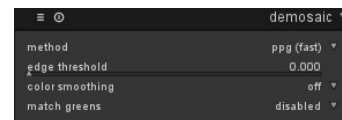
Usage

Demosaic is an essential step for any raw image development process.

A detailed description would be beyond the scope of this manual. In a nutshell, the sensor cells of a digital camera are only able to record different levels of lightness, not different color. In order to get a color image, each cell is covered by a color filter, either in red, green or blue. Due to the color sensitivity of the human vision, there are two times more green cells than red or blue. Filters are arranged in a certain mosaic, called Bayer pattern. Therefore each pixel of your image originally only has information about one color channel. Demosaic reconstructs the missing color channels by interpolation with data of the neighboring pixels.

As all this is prone to produce artifacts, various different demosaic algorithms have been developed in the past. Artifacts would typically be visible as Moiree-like patterns when you strongly zoom into your image. Currently darktable supports PPG and AMAZE. Both algorithms produce high quality output with a low tendency for artifacts. AMAZE is reported to sometimes give slightly better results. However, as AMAZE is significantly slower, darktable uses PPG as a default.

Some further parameters of this module can activate additional averaging and smoothing steps. They might help to reduce remaining artifacts in special cases.

Demosaic is always applied when exporting images. Demosaic is done on monitor display only when zoom is greater than 50% or when the according preference setting "demosaicing for zoomed out darkroom mode" is set accordingly. Else color channels are taken from neighboring pixels without an expensive interpolation.

method

Set the demosaic method. darktable currently supports PPG and AMAZE.

edge threshold

Set the threshold for an additional median pass. Defaults to 0 which disables median filtering.

color smoothing

Activates a number of additional color smoothing passes. Defaults to off.

match greens

> In some cameras the green filters have slightly varying properties. This parameter adds an aditional equilibration step, which suppresses artifacts that otherwise could result. Available options are "diabled", "local average", "full average" and "full and local average".

4.5.7. Highlight reconstruction

Overview

This module tries to reconstruct color information that is usually clipped because of the information being not complete in all channels. If you do nothing, your clipped areas are often toned to the not clipped channel. For example, if your green and blue channels are clipped, then your image will appear red in the clipped areas.
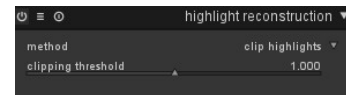
Usage

method

> You can choose between two methods: clipping highlight or reconstructing in LCh. Clipping highlight analyses each pixel having at least one channel clipped. Then it sets all channels to the minium value found among the channels. Reconstruct in LCh analyses each pixel having at least one channel clipped and transforms the information in LCh color space to linearly mix the channels.

clipping threshold

> Manually adjust the clipping threshold against magenta highlights. It's normally not needed to touch this.

4.5.8. White balance

Overview

This module is used to set the white balance. You have three ways to interact with it: (a) Set up tint, temperature in and temperature out, (b) define the value of each channel, or (c) choose from predefined white balances.
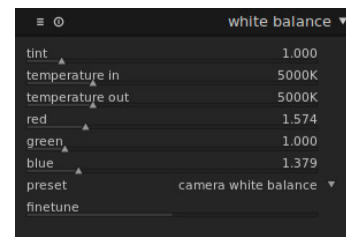
Usage

tint

> Alter the colour tint of the image, from magenta (value < 1) to green (value > 1). The channel sliders will be updated when you adjust this parameter.

temperature in/out

> Set the temperature in and temperature out (in Kelvin). The channel sliders will be updated when you adjust this parameter.

red, green and blue channels

> Set the channel values on a scale from 0 to 8.

preset

Select a preset white balance.

camera white balance (default)	White balance reported by the camera.
spot white balance	Select a square area in your image containing mostly grey pixels. The white balance is calculated based on that the selected area only contains grey pixels.
passthrough	Show without adjusting for white balance.
camera presets	Camera specific white balance presets. Examples: direct sunlight, flash, cloudy, shade and a number of indoor lighting options.

finetune

Some cameras offer additional finetuning parameters if one of the camera presets is selected. Depending on camera white balance can be adjusted step-wise within a certain range. The adjustments are usually towards yellow (value < 1) or blue (value > 1).

4.5.9. Invert

Overview

The main purpose of this module is to invert scanned negatives.
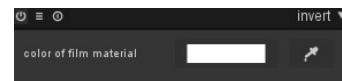
Usage

color of film material

The only control element of this module is a color selector which allows to adjust for different colors of your film material. Clicking on the colored field will open a color selector dialog which allows to define a color in HSL or RGB color space. You can also activate a color picker by pressing [icon] and take a color probe from your image - preferrably from the unexposed border of your negative.

4.6. Tone group

This group contains the modules to work with the tone in the image, this means only changin the lightness and leaving the color intact.

4.6.1. Fill light

Overview

This module allows a local modification of the exposure based on pixel lightness.
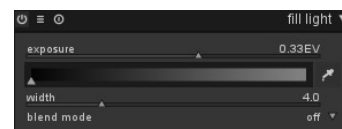
Usage

The module pushes exposure by increasing lightness in the form of a Gaussian curve, centered on a given lightness, with a given width in lightness.

exposure

Set the fill-light exposure in [EV].

center

Set the median lightness being impacted by the fill-light. A color picker is activated by pressing ⬀. It will show the picked lightness value in the gradient bar and thus helps to find the right center value.

width

Set the width of the Gaussian curve. It is expressed in zone, the whole dynamic range being 10 zones. As the Gaussian is symmetric, only even numbers can be entered.

4.6.2. Levels

Overview

A level tool that most of you will recognize how to use from other imaging softwares.
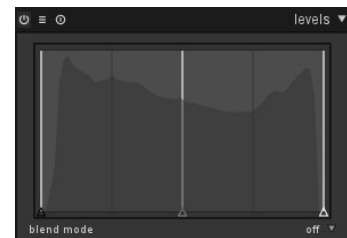
Usage

The levels tool shows the histogram of the image and displays three bars with handles for you to drag and correct the tone in the image. Those bars represent black, middle gray and white.

This module is especially useful if the histogram of your image does not span the whole horizontal range from pure black to pure white. You can adjust the left and right bar to match the left and right border of the histogram, which will make the output image span the full available tonal range. A previously flat looking image will get more contrast and pop.

Shifting the middle bar will modify the middle gray tones. Shifting it left will make the image look brighter, shifting it right will make it darker. This is often referred to as a change of image gamma.

4.6.3. Tonecurve

Overview

This module is a classical tool in digital photography. darktable's special feature is to have this module act in Lab color space. Thus, is offers three independent curves for L, a, and b channels.

40
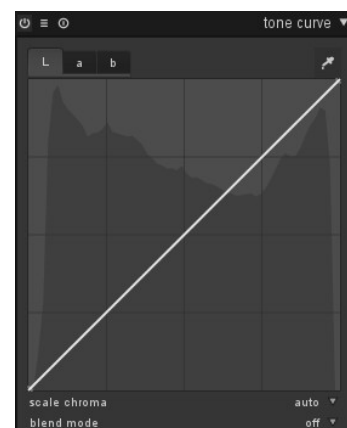
Usage

In its default state curves will be straight lines, defined by few anchor nodes. You can move the nodes with your mouse to modify the curve. You can also generate new nodes by clicking on the curve. Up to 20 nodes per curve can be defined. In order to remove a node, move it out of the diagram area. A color picker is activated by pressing ⬈ and will show the picked values in the diagram. Numerical Lab values of input and output (see below) at the selected spot are shown on top left of the diagram.

L-channel curve

The tonecurve in L-channel works on Lightness. For a better overview a lightness histogram is displayed in the diagram.

On the horizontal line is represented the lightness of the input image pixels. The vertical line represents the lightness of the output image pixels. A straight line does not change anything. A point above the default diagonal increases the lightness whereas a point under decreases it. Shifting the center of the curve upwards will lighten the image, shifting it downwards will darken the image. An S-like curve will enhance the contrast of the image.

a/b-channel curves

The curves in a-channel and b-channel work on color values. They are only displayed and active if combobox "scale chroma" is set to "manual". On the horizontal line is represented the color channel value of the input image pixels. The vertiacal line represents the color channel value of the output image pixels. Positive a-values correspond to more magenta colors; negative a-values correspond to more greenish colors. Positve b-values correspond to more yellowish colors; negative b-values correspond to more blueish colors.

A straight line does not change anything. Shifting the center of the curve will give the image a color tint: upwards shifting a-channel gives a magenta tint; upwards shifting b-channel gives a yellow tint; downwards shifting a-channel gives a green tint; downwards shifting b-channel gives a blue tint.

Increasing/decreasing the steepness of the curves without shifting its center will increase/decrease the color saturation of the respective channel. With properly defined curves you can excert fine control on color saturation depending on input pixel's colors.

scale chroma

darktable does an automatic adjustment of color saturation if this combobox is set to "auto". The level of this adjustment is dependent on the pixel's color values and its L-value modification by the L-channel tonecurve. It is designed to give an overall boost in color saturation if the L-curve gives a contrast boost. Look at blend mode "coloradjustment" to adjust the strength of this effect. Only if this combobox is set to "manual", you can do your own modification of color saturation by the curves in channels a and b.

Examples

Original image
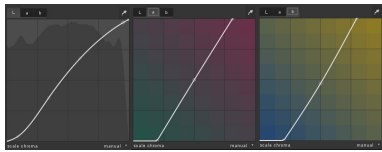
Tonecurve settings. Please note how the center node of our b-curve was shifted down to negative values. This gives the image its blue tint.

Resulting image

4.6.4. Zone system

Overview

This module is another way to change the lightness of your image. It is based on the Ansel Adams system. It allows to modify the lightness of a zone taking into account the effect on the adjacent zones. It divides the lightness in an user-defined number of zones.
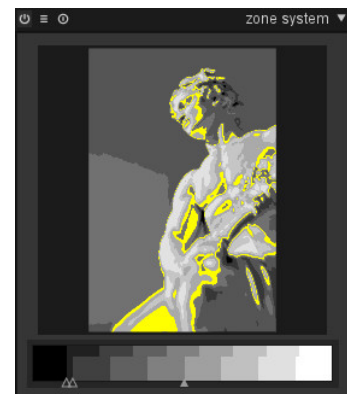
Usage

The lightness is processed on the L* channel from CIELAB.

The center view shows the image broken down in zones.

When hovering above a zone on the lightness scale, the current zone is highligthed on the preview. The number of zones can be changed by mouse-scrolling on the lightness scale.

Left click and drag a handle in the zonebar to modify the zonemapping, use right click to remove a controlpoint.

Examples

This is the original image use it as reference for the changes below...

Here i have compressed some darker zones and lighter zones to increase contrast, then expanded the upper parts of darker zones...

4.6.5. Local contrast

Overview

This module allows to enhance local contrast. It uses the unnormalized bilateral filter and works on the L channel from Lab.
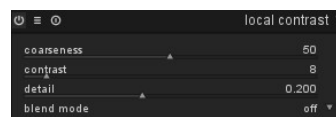
Usage

Local contrast boosts the details of your image much like the equalizer does, but in only one frequency band.

coarseness

Make the details you want to adjust finer or coarser (this maps to sigma_s in the algorithm).

contrast

How strongly the algorithm distinguishes between brightness levels (maps to sigma_r in the algorithm). Increasing the value will result in a more contrasty look.

detail

Add or remove detail. higher values will increase local contrast.

Example

before

After, a little overdone to demonstrate the effect. Use this sparingly to avoid a cheap overprocessed look.

4.6.6. Tonemapping

Overview

This module allows to compress the tonal range of HDR images so they fit into the limits of a normal low dynamic range image. It uses Durand's 2002 algorithm. darktable can import HDR images if they come in OpenEXR, RGBE or PFM format or as a DNG gen-

erated by darktable's HDR creation mechanism (see
chapter on Lighroom mode).
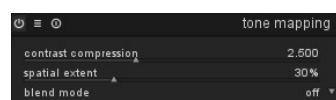
Usage

The underlying algorithm uses a bilateral filter to decompose your image into a coarse
base layer and a detail layer. The contrast of the base layer is compressed while the detail
layer is preserved, then both layers are re-combined.

contrast compression

Sets the compression of the contrast to fit dynamic range.

spatial extent

Sets the spatial extent of the bilateral filter. Lower values will cause contrast compression
to have stronger (often adverse) effects also on image details.

4.6.7. Global tonemap

Overview

This modules implements another approach to com-
press the tonal range of HDR input images into the
limited tonal range of a typical LDR output file. It of-
fers several implementations of global tonemap op-
erators.

Usage

Global tonemapping processes each pixel of an HDR image without taking the local sur-
rounding into account. This is generally faster than local tonemapping as implemented
in module *tonemapping* but might lead to less convincing results with very high dynamic
scenes. As an enhancement to the original operators, darktable's implementation allows
to preserve detail of the input image and transfer it back into the output image.

operator

Reinhard, Filmic and Drago global tonemap operators are available for use. Depending
on the selected operator different parameters can be adjusted. Some operators are fully
self-adjusting and do not require specific controls.

bias

Only offered for operator *Drago*. This parameter influences the contrast of your output
image. It is an essential parameter for adjusting the compression of high values and the
visibility of details in dark areas. According to the original paper a value of 0.85 is recom-
mended as a starting point.

target

Only offered for operator *Drago*. This is a scale factor to adjust global image brightness
to the brightness of the intended display. It is measured in cd/m^2 and should match the
according value of your output device. Higher values lead to a brighter, while lower values
lead to a darker image.

detail

Offered as an addition to all operators. This parameter controls how much detail is pre-
served and transfered back into the output image after tonemapping.

44

4.7. Color group

This group contains the modules to work with colors of your image.

4.7.1. Velvia

Overview

The velvia module enhances the saturation in the image; it increases saturation on lower saturated pixels more than on high saturated pixels.
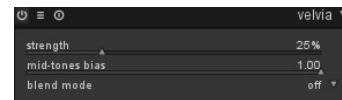
Usage

strength

This slider controls the strength of saturation for the effect.

mid-tones bias

This slider controls saturation preservation in mid-tones. It is expressed as the lightness repartition between mid-tones and shadows/highlights. The greater the first number, the stronger the effect.

4.7.2. Channel mixer

Overview

This module is a powerful tool to manage channels. As entry, it manipulates red, green and blue channels. As output, it uses red, green, blue or gray or hue, saturation, lightness.
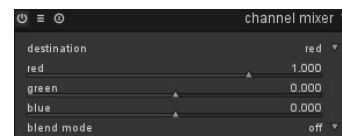
Usage

Select your output channel and then set the amount of each channel to feed the output channel.

Examples

This is the original image use it as reference for the changes below...

Here we use gray as output and the following rgb mix: 0.9, 0.3, -0.3. What this does is to reduce blue colors in the mix which usally represents some details. We use some more red than green as the skin does contain a very even red tone and the results out of this is a pretty smooth skin of the portrait.

Due to the images nature, I needed to lighten it up with 0.1 EV

In this example we use an rgb mix of 0.4, 0.75, -0.15. The difference with the previous example is that we use more green than red, which brings some features back. We still reduce the blue channel in the mix to dampen artefacts.

Table of mixing values for some b/w films

Select gray as output mixing channel, and try out the values descibed below for your like...

Film Type	Red	Green	Blue
AGFA 200X	0.18	0.41	0.41
Agfapan 25	0.25	0.39	0.36
Agfapan 100	0.21	0.40	0.39
Agfapan 400	0.20	0.41	0.39
Ilford Delta 100	0.21	0.42	0.37
Ilford Delta 400	0.22	0.42	0.36
Ilford Delta 3200	0.31	0.36	0.33
Ilford FP4	0.28	0.41	0.31
Ilford HP5	0.23	0.37	0.40
Ilford Pan F	0.33	0.36	0.31
Ilford SFX	0.36	0.31	0.33
Ilford XP2 Super	0.21	0.42	0.37
Kodak T-Max 100	0.24	0.37	0.39
Kodak T-Max 400	0.27	0.36	0.37
Kodak Tri-X 400	0.25	0.35	0.40
Normal Contrast	0.43	0.33	0.30
High Contrast	0.40	0.34	0.60
Generic B/W	0.24	0.68	0.08

4.7.3. Output color profile

Overview

This module manages the output profile both for dis-
play and for export.
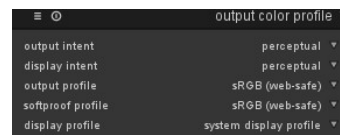
Usage

output intent

Sets the rendering intent for output/export. You can easily override this setting whenever
you do exports from lighttable mode.

perceptual	Suited to pictures. It scales one gamut to fit into the other while maintaining the relative position of colors.
relative colorimetric	Out-of-gamut colors are converted to colors having the same lightness, but different saturation. Other colors remain unmodified.
saturation	Saturation is kept but lightness is slightly changed.
absolute colorimetric	Keep the white point.

display intent

Sets the rendering intent for your display. See above for available options.

output profile

Sets the color profile for output/export. darktable comes with pre-defined internal pro-
files: sRGB, AdobeRGB, XYZ and linear RGB. Additional profiles to be used are picked up
from $DARKTABLE/share/darktable/color/out and $HOME/.config/darktable/color/out.

$DARKTABLE is used here as a placeholder for darktable's installation directory and
$HOME is your home directory.

softproof profile

Sets the color profile for softproof. Profiles to be used are picked up from $DARK-
TABLE/share/darktable/color/out and $HOME/.config/darktable/color/out.

Softproof is toggled on/off by pressing *s*; this function allows you to preview the image
rendered using a printer profile, to see how colors will end up on the final print.

Likewise Gamut Check is toggled on/off by pressing *g*; this function marks all pixels, which
are out of gamut with respect to the selected softproof profile, by showing them in cyan.
Softproof and Gamut Check are mutually exclusive modes.

Softproof and Gamut Check can be activated at any place in darkroom mode. Only if mod-
ule "Output color profile" is in focus, you will get the status displayed left below the cen-
ter image. Especially for Softproof this may raise doubts, if this mode is still active or not.
If in doubt press *g*, which will switch to Gamut Check and will be clearly discernible by

cyan marked pixels. Press *g* again and you are back to normal display without Softproof or Gamut Check.

display profile

Sets the color profile for the display. Profiles to be used is picked up from $DARK-TABLE/share/darktable/color/out and $HOME/.config/darktable/color/out. The additional option "system display profile" is taken directly from your system's color manager or from your X Server. When working with a calibrated display this is the preferred option to choose.

4.7.4. Color contrast

Overview

Color contrast filter changes the contrast between green versus magent and blue versus yellow.
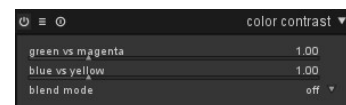
Usage

Higher values increase color contrast, lower values decrease it. The effect of this module is similar to a steepened or flattened a- or b-curve in module tonecurve.

green vs. magenta

Changes color contrast for green versus magenta.

blue vs. yellow

Changes color contrast for blue versus yellow.

4.7.5. Color correction

Overview

This module can be used to modify the global saturation or to give a tint.
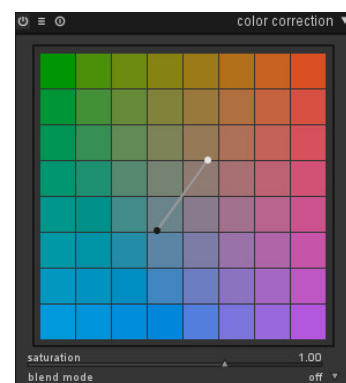
Usage

color board

You can drag a white dot and a dark dot independently to give global tint to highlights and shadows. Set dots to the same color for a homogeneous global tint or to different colors for split toning.

48

saturation

Use the saturation slider to correct the global saturation.

4.7.6. Monochrome

Overview

This module is a quick way to convert image in black and white. You can simulate a color filter in order to modify your conversion. That filter can be changed in size and color center.
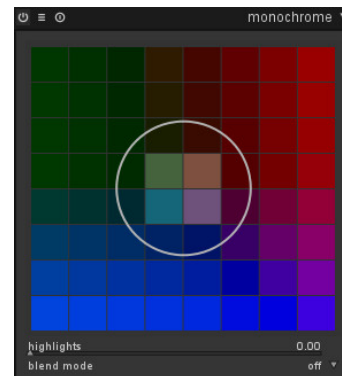
Usage

You can change the filter size by scrolling with your mouse wheel.

4.7.7. Color zones

Overview

This module allows to selectively modify the colors in your image. It is highly versatile and allows every transformation possible in the LCh colorspace.
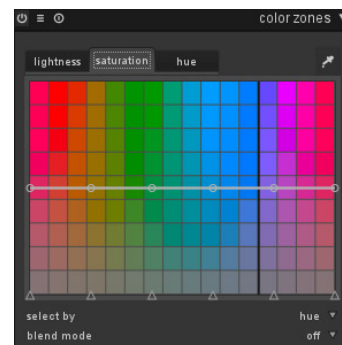
Usage

The horizontal axis represent the different values you can modify. The vertical axis shows the changes you can achieve.

For both horizontal and vertical axises you can work on lightness, saturation or hue. A color picker is activated by pressing ⬀ and will show the picked values in the diagram.

You can click on any of the six nodes on the curve and drag to adjust it vertically. A circle indicates how strong adjacent nodes will be affected. Use the scroll wheel of your mouse to change the circle diameter. You can also use the six controlpoints (triangles which define the vertical value of the nodes) at bottom to adjust the curve.

4.7.8. Color transfer

Overview

This module allows to transfer on an image the colors found on another.
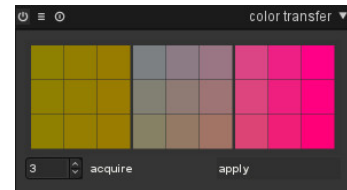
Usage

number of clusters

Set the number of color clusters to use.

acquire

Acquire the colors of the current image.

apply

Apply the colors (shown above) previously acquire on another image, on the current image.

4.7.9. Vibrance

Overview

This modules affects the vibrance in the image. Vibrance is a widely used term and the endresult differs from program to program. Vibrance in darktable does saturate the most saturated pixels and brings down the lightness of those pixels to make the colors more vivid.
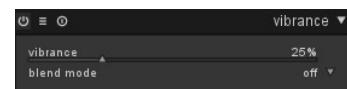
Usage

Vibrance does only have one parameter which controls the amount of vibrance to apply.

vibrance

The amount of vibrance to apply to the image.

4.7.10. Input color profile

Overview

This module sets up the input color profile.
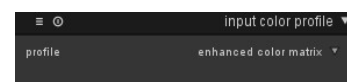
Usage

profile

Choose the profile or color matrix to apply. darktable offers many widespread matrices along with an enhanced matrix for various camera models. The enhanced matrices were processed by the darktable team in order to provide a look close to the manufacturer one.

You can also supply your own input ICC profiles and put them into $DARKTABLE/share/darktable/color/in or $HOME/.config/darktable/color/in. $DARKTABLE is used here as a placeholder for darktable's installation directory and $HOME is your home directory. One common source of ICC profiles is the software that is shipped with your camera; it often contains profiles specific to your camera model.

If your input image is a low dynamic range file like JPEG it might already contain an embedded ICC profile. darktable will use this as a default on those images. If you know what you are doing, you can overrule darktable and select a different one. Select 'embedded icc profile' to restore the default.

4.7.11. Unbreak input profile

Overview

This module allows to correct an input profile.
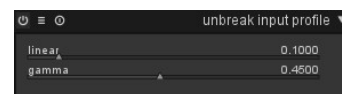
Usage

linear

Set the linear value.

gamma

Set the gamma value to correct your input profile.

4.8. Correction group

The correction group contains the modules that will correct typical problems in an photo such as hotpixels, spot removal, noise, lens correction among others.

4.8.1. Sharpen

Overview

This is an standard UnSharp Mask tool for sharpening the details of an image.

Usage

This module works by enhancing the contrast around edges and thereby enhances sharpness impression of an image. In darktable this module is only applied to the L-channel in Lab color space.

radius

USM applies a gaussian blur to your image as part of its algorithm. This controls the blur radius which in turn defines the spatial extent of edge enhancement. Too high values will lead to ugly over-sharpening.

amount

This controls the strength of sharpening.

threshold

Contrast differences below this threshold are excluded from sharpening. Use this to avoid amplification of noise.

4.8.2. Equalizer

Overview

This versatile module can be used to achieve a variety of effects, such as: bloom, denoising, and local contrast enhancement. It works in the wavelet domain, and parameters can be tuned for each frequency band separately.

Usage

The parameters can be tweaked for each frequency band. In particular, you can adjust contrast boost and denoise threshold splines for both lightness and chromaticity, as well as the sharpness of the wavelet basis on each frequency scale.

Each spline can be dragged with a proportional edit approach, use the mouse wheel to adjust the radius in which your changes will have an effect. The transparent area indicates where you would drag the spline with the current mouse position and radius. Also the little triangles at the bottom can be moved to alter the x-position of the spline nodes.
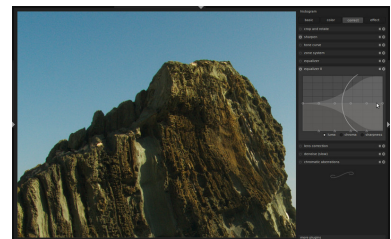
Drag the upper line (bright circles, here for the lightness channel) to affect local contrast. Pulling it up, as shown here, will result in a contrast boost for that frequency band. Higher frequencies, i.e. smaller details, are to the right of the grid. Pulling it down works, too.
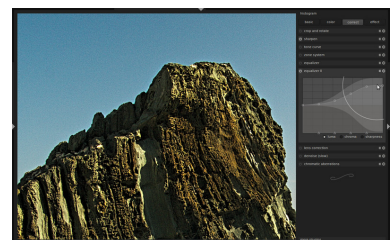
The bottom spline (black circles) is used to perform denoising. It adjusts the wavelet shrinkage threshold for each frequency band. Pull it up to see the effect. In this example, the noise which has been amplified by local contrast enhancement is removed.

52
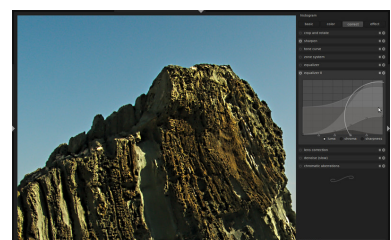
This screen shows the effect of the sharpen parameter. It is here pulled down to zero for all bands. This is effectively a regular à trous wavelet, without edge detection, and results in the characteristic halos around sharp edges in the image.

This image is the other extreme. The wavelet basis now oversharpens, which results in ugly gradient reversals near the ridge of the rock.

Note that the sharpness parameter only affects the wavelet basis, not the image directly. So you have to change some denoise/contrast boost parameters to see an effect if fiddling with the sharpness parameter.

Also have a look at the presets. All parameters are changed in there, so these might be a good starting point to gain an intuitive understanding about them. Note that the sharpness of the wavelet basis has been set to zero to achieve the bloom effect.

4.8.3. Denoise - Non local means

Overview

This is a denoise algorithm, which will work on chroma and/or luma.
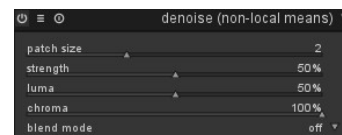
Usage

This module reduces noise in your image but preserves structures. This is accomplished by averaging a pixel with other pixels in the image. The weight of averaging depends on the similarity of the pixels surroundings with the surrounding of the one pixel to be denoised. A patch with a certain size is used to measure that similarity. As denoising is a resource hungry process, it slows down pixelpipe processing significantly; consider to activate this module late in your workflow.

patch size

The radius of the patch.

strength

The strength of the denoise. Higher values lead to a stronger effect.

luma

Amount of denoise to apply on luma. Select carefully in order not to lose too much structure.

chroma

> Amount of denoise to apply on chroma. You can be much more aggressive with this parameter compared to luma.

4.8.4. Denoise - bilateral

Overview

This module is used to denoise high iso pictures. It is flagged as a slow module due to its high resources consumption. Depending of your hardware, this module can currently slow down darktable or even crash it. Quite counter-intuitively, the greater the values for sliders, the lesser resources.
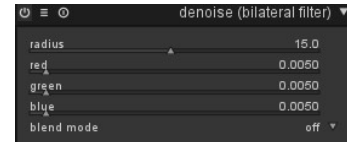
Usage

This module reduces noise in your image but preserves sharp edges. This is accomplished by averaging pixels with their neighbors, taking into account not only the geometric distance but also the distance on the range scale, i.e. differences in color intensities. As denoising is a resource hungry process, it slows down pixelpipe processing significantly; consider to activate this module late in your workflow.

radius

> Set the spatial extent of the gaussian blur.

red

> Blur intensity for red channel.

green

> Blur intensity for green channel.

blue

> Blur intensity for blue channel.

4.8.5. Lens correction

Overview

This module is able to correct certain lens flaws, namely distortions, transversal chromatic aberrations (TCA) and vignetting. It relies on the external library lensfun, which comes with correction profiles for many (but not all) common cameras end lenses.
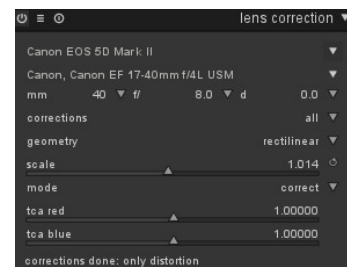
Usage

In order to perform lens corrections the module uses EXIF data of your image to identify the specific camera/lens combination and collects the needed correction parameters from a profile in lensfun's database.

54

camera

> The camera make and model as determined by EXIF data. You can overrride this manually and select your camera from a hierarchical menu.

lens

> The lens make and model as determined by EXIF data. You can overrride this manually and select your lens from a hierarchical menu. This is mainly needed for pure mechanical lenses.

photometric parameters

> Corrections additionally depend on certain photometric parameters that are read from EXIF-data: focal length (needed for distortion, TCA, vignetting), aperture (needed for TCA, vignetting) and focal distance (needed for vignetting). Many cameras do not record focal distance in their EXIF data; most likely you need to set this manually. You can override all automatically selected values.
>
> If no controls for the three photometric parameters are displayed, this means that no correction profile could be found for the automatically identified camera/lens combination. You may try to find the right profile yourself by searching for it in the menu. If there is no matching profile, visit lensfun's home page at http://lensfun.berlios.de and learn how to generate your own set of correction parameters. Don't forget to share your profile with the lensfun team!

corrections

> This combobox gives you a choice about which corrections (out of distortion, TCA and vignetting) darktable shall apply. Change this from its default "all", if your camera has already done some internal corrections (e.g. of vignetting), or if you plan to do certain corrections with a separate program.

geometry

> In addition to correction of lens flaws, this module can change the projection type of your image. Set this combobox to the aimed projection type.

scale

> This slider allows you to adjust the scaling factor of your image. Pressing the auto scale button (right to the slider) will let darktable find the best fit to avoid black corners.

mode

> The default behavior of this module is to correct lens flaws. Switch this togglebutton to "distort" in order to simulate the behavior of a specific lens (inverted effect).

TCA red

> This slider allows to override the correction parameter for TCA. You can also use this slider to manually set the parameter in case the lens profile does not support TCA correction. Look out for colored seams at features with high contrast edges and adjust this parameter and the following one to minimize those seams.

TCA blue

> This slider allows to override the correction parameter for TCA. You can also use this slider to manually set the parameter in case the lens profile does not support TCA correction.

corrections done

You will sometimes observe that for a given camera/lens combination only part of the possible corrections (distortion, TCA, vignetting) are supported by lensfun's profiles. This message box will tell you what corrections have actually been applied.

4.8.6. Spot removal

Overview

Spot removal allows you to correct a zone in your image by using another zone as model.
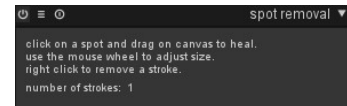
Usage

click on the area to be healed then drag to a clean area. The circle with line inside is the clean one.

use mouse scroll-wheel over the stroke to change size.

right-click on a stroke to delete.

Examples
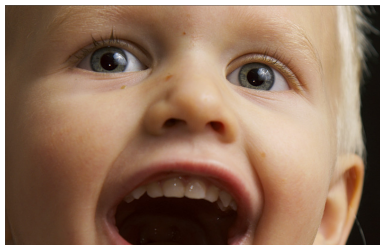
Let's use this portrait as example; we want to remove some dirt and unwanted catchlight from camera pop-up strobe.

I have marked all the spots that I want to remove from the image. I first start by pressing the spot and drag to an area that should be used as reference to remove the spot. The handle (circle) which has a line that ends in the center of the circle indicates reference source.

And here is the result image of the spotremoval.

4.8.7. Raw denoise

Overview

Raw denoise allows you to perfom denoising on pre-demosaic data. It is ported from dcraw.
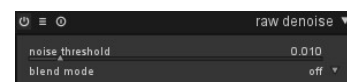
Usage

noise threshold

Set the threshold for noise detection.

4.8.8. hotpixels

Overview

This module is able to automatically detect and eliminate hotpixels. Hotpixels are pixels which failed to record light level correctly. Detected hotpixels are replaced by an average value of their neighbors.
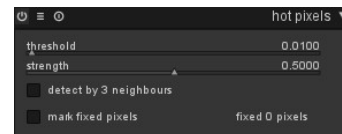
Usage

Control the detection sensitivity with the threshold parameter and the level of elimination with the strength parameter.

threshold

The threshold of the detection, i.e. how strong a pixel's value needs to deviate from its neighbors to be regarded as a hotpixel.

strength

The strength of blending hotpixels with their surrounding.

detect by 3 neighbours

Will extend the detection of hotpixels. Even regard a pixel as hot if a minimum of only three (instead of four) neighbor pixels deviate by more than the threshold level.

mark fixed pixels

This options will mark the pixels that have been corrected. It also displays the count of detected and fixed pixels.

4.8.9. Chromatic aberrations

Overview

This module allows you to correct chromatic aberrations.
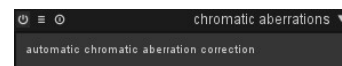
Usage

This module has no option. On activation it will automatically try what it can in order to reduce visible CA's.

4.9. Effect group

In the effect group you will find modules with a more artistic touch.

4.9.1. Watermark

Overview

The watermark module provides a way to render a
vector-based overlay onto your image. Watermarks
are standard SVG documents and can be designed us-
ing Inkscape.
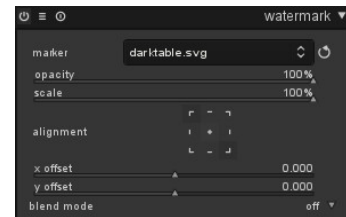
The SVG processor of darktable also substitutes strings within the SVG document, which
gives the oppurtunity to include image dependent information in the watermark such as
aperture, exposure time and other metadata.

User-designed watermarks are placed into the directory $HOME/.config/darktable/water-
marks. Once in place, use the reload button at the right of the watermark file name to
update the list of available watermarks to use.

Here follows a list of available variable strings that is supported for substitution within
the svg document.

$(DARKTABLE.NAME)	The application name
$(DARKTABLE.VERSION)	The version of darktable
$(IMAGE.ID)	The uniq image id within current library
$(IMAGE.FILENAME)	The iamge filename
$(IMAGE.EXIF)	The image exif string
$(EXIF.DATE)	The image date
$(EXIF.MAKER)	The maker of camera model
$(EXIF.MODEL)	The camera model
$(EXIF.LENS)	The specific lens used
$(XMP.DC.CREATOR)	The creator string
$(XMP.DC.PUBLISHER)	The publisher string
$(XMP.DC.TITLE)	The title of the image
$(XMP.DC.DESCRIPTION)	The description of the image
$(XMP.DC.RIGHTS)	The rights assigned to the image

Usage

marker

Choose the watermark of interest, use the reload button next to the combobox to update
list with newly added watermarks.

opacity

Set the opacity of the render of watermark.

scale

Scale the watermark pixelindependent.

58

alignment

Use these controls to align the watermark to any edge or center of the image.

x offset

Pixel-independent offset realtive to the choice of alignment on x-axis.

y offset

Pixel-independent offset relative to the choice of alignment on y-axis.

4.9.2. Framing

Overview

This module is an artistic feature to generate a frame around your image. The frame consists of a border with a user defined color and a frame line inside that border, which has another user defined color. There are various options for you to control the geometry of your frame.
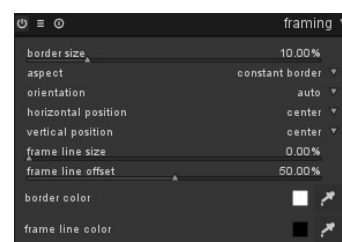
Usage

border size

This slider controls the size of the frame in percent of the underlying full image.

aspect

With this combobox you can choose between different aspect ratios for the final output of this module, i.e. underlying image plus frame.

orientation

If you select a non-square aspect ratio this combobox defines the orientation - portrait or landscape. Set to "auto" if you want darktable to select the most reasonable orientation based on the underlying image.

horizontal position

Select from a set of pre-defined ratios where you want your underlying image be positioned on the horizontal axis. You can also right click and enter your own ratio as "x/y".

vertical position

Select from a set of pre-defined ratios where you want your underlying image be positioned on the vertical axis. You can also right click and enter your own ratio as "x/y".

frame line size

The percentage of the frame line size relative to the border size at its smallest part.

frame line offset

Where the frame line is positioned relative to the underlying image. Select 0 for a frame line touching the image, 100% for a frame line touching the outer border limits.

border color

> Clicking on the colored field will open a color selector dialog which allows to define a color for the border in HSL or RGB color space. You can also activate a color picker by pressing ⬈ and take a color probe from your image.

frame line color

> Clicking on the colored field will open a color selector dialog which allows to define a color for the frame in HSL or RGB color space. You can also activate a color picker by pressing ⬈ and take a color probe from your image.

Examples

Example image with a user defined frame.

4.9.3. Splittoning

Overview

darktable's splittoning method creates a two color linear toning effect where the shadows and highlights are represented by two different colors. In the example image below you can see an original black and white image and one where a splittoning effect are applied with blue in shadows and yellow-ish color in highlights.
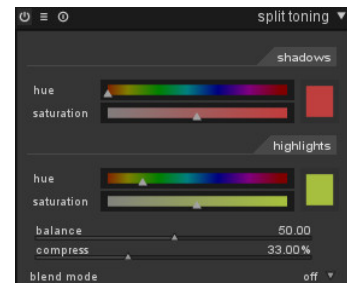
darktable splittoning module has more parameters that impacts on the default above behavior- We have 'balance' which offsets the 50% gray level in image and 'compression' that does compress toning in the shadows and highlights which leaves a gap which are untouched by the effect.

The splittoning module does not convert images to black and white and has limited benefits on color images. So if you want to do traditional splittoning, use the monochrome module to make the image black and white before playing around with splittoning effect.

Usage

shadows and highlights color

> These controls are used to set the color of the splittoning effect, you select the desired color and saturation for both shadows and highlights, you can also click the color preview box to bring up a common color picker dialog.

balance

The balance is presented as a ratio and defaults to 50/50, which means that 50% of light-ness range in image are used for shadows toning and the other 50% for highlights toning.

compression

Compression is a percentage of total lightness range that is not affected by color toning. Default value is set to 33%; this is not the default behaviour of an original splittoning which would be 0% compression. The choice of 33% as a default is to invite you expriment-ing with these parameters and how it extends the original splittoning method.

Examples
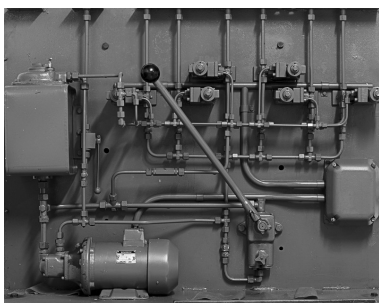
Original black and white image.
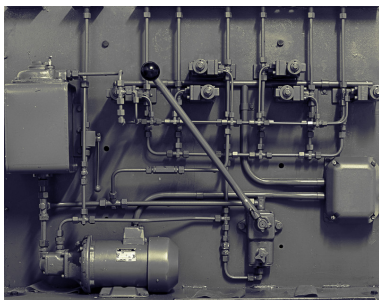
Splittoning with blue shadows and yellow highlights.

4.9.4. Vignetting

Overview

This module is an artistic feature which creates vi-gnetting (modification of the brightness/saturation at the borders).

Usage

The vignetting module has an extensive set of parameters to precisely tune its effect. It also will display graphical controls in center view if the module is in focus. Give it a try to get a feeling how it works. On-screen controls and parameter sliders will stay in sync.

scale

Set radius of the vignetting area.

fall-off strength

Set the progressiveness of the fall-off. Higher values will cause a steeper transition.

brightness

Sets the intensity of brightening (positive values) or darkening (negative values).

saturation

Controls how strong colors get desaturated or saturated in the darkened or brightened vignetting area.

horizontal center

Shifts the center of the vignetting area horizontally.

vertical center

Shifts the center of the vignetting area vertically.

shape

Influences the shape of the vignetting area. The default value of 1 causes a circular or elliptical area. Smaller values will shift the shape into a more square one; higher values turn it into a cross-like shape.

automatic ratio

Click this button to automatically adjust the width/height ratio of the vignetting area to the aspect ratio of the underlying image. The vignetting area will typically become elliptical.

width/height ratio

Manually adjust the width/height ratio of the vignetting area.

dithering

With this combobox you can activate dithering which prevents banding artifacts. Those artifacts are generated when darktable's continuous 32-bit float data are quantized into an 8-bit or 16-bit output format. Select "8-bit output" to prevent banding on monitor display and for JPEGs. When set to "16-bit output", only a little dithering will be applied, just strong enough to compensate for banding on the fine grained 16-bit level.

Examples
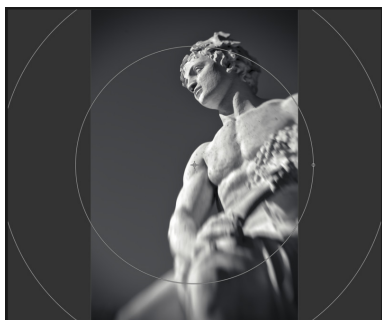
An image with vignetting and with graphical vignetting controls.

4.9.5. Soften

Overview

This module is an artistic feature that creates a soft-ened image also commonly known as the Orton ef-fect.
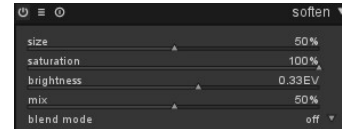
Usage

Michael Orton achieved such results on slide film by using two exposures of the same scene: one well exposed and one overexposed; then he used a darkroom technique to blend those into a final image where the overexposed image was blurred.

This module is almost a copy of Orton's analogue process into the digital domain. You can control brightnesss and blur with the provided parameters; we also add a control for saturation of the overexposured image for more play.

size

Set the size of blur of the overexposed image in the process, the bigger the softer.

saturation

Set the saturation of the overexposed image.

brightness

Expressed in [EV], the brightness slider selects the increase in brightness.

mix

Controls the mix of the overexposed image and the overall effect.

Examples

This is the original image, use it as reference for the changes below...

In this image I used the default values, and added 0.33EV to brightness for a little more light in the soft layer.

This version is the same as above but with 25% saturation.

4.9.6. Grain

Overview

This module is an artistic feature which simulates the grain of a film.
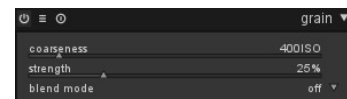
Usage

The grain is processed on the L* channel from CIELAB.

coarseness

Set the grain size, which has been scaled to simulate an ISO number.

strength

Set the strength of the effect.

4.9.7. Highpass

Overview

Highpass acts as a high pass filter. The primary usage for this filter is in combination with a blending operator. Try out blend mode "soft light" to get high-pass sharpening. Use the opacity slider to adjust the strength of the effect or even conditional blending to limit the effect to only parts of your image.

Usage

sharpen

Set the sharpen. The higher, the more details.

contrast boost

Set the contrast boost.

4.9.8. Lowpass

Overview

A lowpass filter (eg. gaussian blur) with additional control of outcome as contrast and saturation. The primary usage for lowpass filter is in combination with a blending operator, try out the preset named
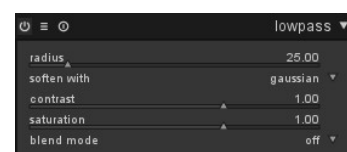
"local contrast mask" with a overlay blending operation.

Usage

This module offers enormous artistic potential, albeit with difficult to predict results sometimes.

radius

Set the radius of the blur.

soften with

This combobox defines the blur algorithm; you can chose between Gaussian blur (default) and bilateral filter. The latter leads to an edge preserving blur. Gaussian will blur all image channels: L, a and b. Bilateral will only blur L channel.

contrast

Changes the contrast. Negative values result in an inverted negative image. Higher absolute values increase contrast; lower absolute values reduce contrast. A value of zero leads to a neutral plane.

saturation

Changes the color saturation. Negative values result in complementary colors by inverting the a/b-channels. Higher absolute values increase color saturation; lower absolute values reduce color saturation. A value of zero leads to a desaturated black&white image.

Examples

The original image, already heavily processed. The boat is almost a silhouette.
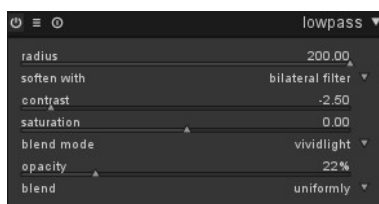
Bilateral blur with high radius. Desaturated, inverted and with high contrast.

Resulting image from lowpass filter ...

... and final image after this was applied with blend mode "vividlight".

4.9.9. Low light

Overview

Low light module allows to simulate human lowlight vision, thus providing ability to make lowlight pictures look closer to reality. It can also be used to perform a day to night conversion.

The idea is to calculate scotopic vision image (http://en.wikipedia.org/wiki/Scotopic_vision) which is perceived by rods rather than than cones in the eye under low light. Scotopic lightness then is mixed with photopic value (regular color image pixel) using some blending function. Also this module is able to simulate the Purkinje effect by adding some blueness to the dark parts of the image (http://en.wikipedia.org/wiki/Purkinje_effect).
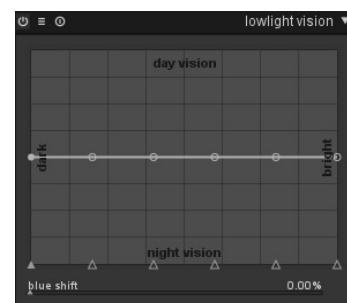
Usage

curve

The horizontal axis is about pixel lightness from dark (left) to bright (right). The vertical axis represents the kind of vision from night vision (bottom) to day vision (top)

blue

Set the blue hint in shadows (Purkinje effect).

Examples

This is the original image.

With low light module on.
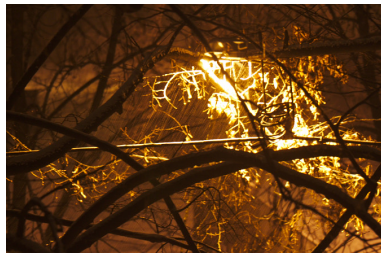
This is the original image.
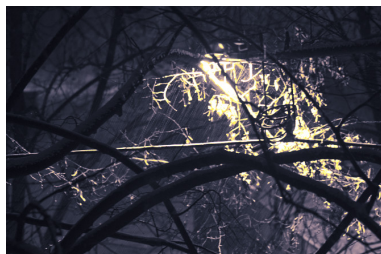
With low light module on.

4.9.10. Bloom

Overview

This module boost highlights and softly blooms them over the image hence the name of the effect, there are numeros ways to use this module dependent on the image actual scenary lightning.

Usage

Keep the default values, and start off by changing the strength value for a pleasant value for your image, then change the size to control the spread of light.

size

Represents the spatial extent of the bloom effect.

threshold

Set the threshold turning on the increase in brightness.

strength

Set the strength of overlightning for the effect.

Examples

This is the original image, use it as reference for the changes below...
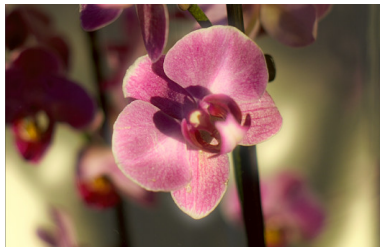
Here I have choosen to use the size 10% which is a rather small radius for the soft light spread, and I also boosted up the strength to 50% for a more exaggerated effect.

4.9.11. Colorize

Overview

This module is an artistic feature that adds a solid layer of color to your image.
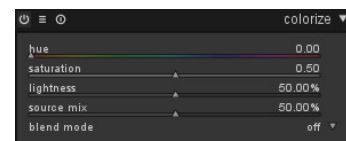
Usage

Several parameters control the effect of this module. Much more versatility can be reached if you apply blending or even conditional blending.

hue

Select the hue tone of the color layer.

saturation

Select the color saturation of shadow tones.

lightness

Select the lightness of the color layer.

source mix

This slider controls how the lightness of input image is mixed in. If you set this to zero a uniformly colored plane will result.

4.9.12. Graduated Neutral Density

Overview

This module aims at simulating a neutral density filter, in order to correct exposure and color in a progressive manner.
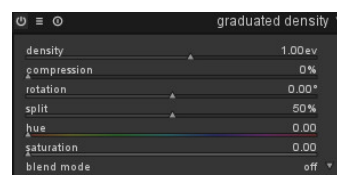
Usage

The module uses a gradient to modify the exposure and the color cast of the image in a non-homogeneously manner.

density

Set the density of the filter in [ev]. A low value underexposes slightly whereas a high value creates a strong filter.

It is expressed as [ev] that is equivalent to *f-stops*. Lens filters are often referred as ND2, ND4, ND8 and so on. Each time you add an [ev] you double the ND. So ND2 is 1 ev, ND4 is 2 ev, and so on. You can also express it in optical density or transmittance. The table below sums up the different approach for the most common filters:

ND	[ev] or f-stop	absorbance	transmittance
ND2	-1	0.3	50%
ND4	-2	0.6	25%
ND8	-3	0.9	12.5%
ND400	-9	2.7	0.195%

compression

Set progressiveness of the gradient. A low value creates a smooth transition, whereas a high value makes the transition abrupt.

split

Set split to offset the filter, i.e. to move the transition zone above or below the centerline. Expressed as a percentage, values above 50% set the transition zone above the centerline.

rotation

Set the rotation so the gradient acts along an axis different from the vertical one. Expressed in degrees, negative values turn clockwise.

hue

Set the hue to add a colorcast to the gradient.

saturation

Set the saturation to add a colorcast to the gradient.

TIP: If you know you intend to use the GND filter before actually making a shot with your camera you might want to underexpose by one or two thirds of an f-stop to make sure de-

tail remains in the highlights. When all detail has truly been blown out the GND filter cannot produce a pleasing results, this is a limitation that is inherent to digital postprocessing. For instructions on how to intentionally underexpose, please consult your camera's manual, look for 'exposure compensation'.

Examples

This is the original image with a pretty overexposed sky, use it as reference for the changes below...

And now we have added a neutral ND8 filter which does a pretty good job on the image..
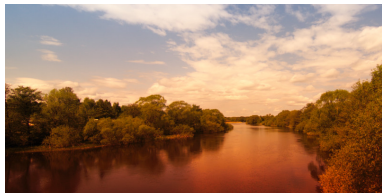
And at last, i added an orange colored filter rotating it -180 degrees for applying it on water/trees for a more artistic use of the filter.

5. Examples

5.1. Converting to Black and White

5.1.1. Overview

Black and white conversion can be achieved in several ways with darktable. Indeed, darktable comes with a lot of modules, especially for color manipulation. In this manual, I will show you 4 ways to perform a black and white conversion.

5.1.2. The obvious way: monochrome module

To perform such conversion, just activate the monochrome module. You can then simulate a color filter, by dragging the circle above the colours you want to filter. Filter size can be modified thanks to wheel scrolling.

5.1.3. The simple way: color correction module

To perform such conversion we use the color correction module.

1. Activate the color correction module

2. Use the bottom slider to set saturation to zero

5.1.4. The artistic way: color zones module

To perform such conversion we use the color zone module.

1. Activate the colour zones module

2. By default, the first radio-buttons row is set to "colorness" whereas the second is set to "hue" wich means that color are selected according to their hue (horizontal scale) and you can change for each hue its "colorness" (vertical scale). You simply need to set all points to the minimum of the vertical scale to de-saturate every hue.

3. But now if you want, you can keep some hues a little bit saturated, so your image will be all black and white but some hue. A classical use for portrait is to keep red hue saturated in order to make the lips standing out.

   You can also use presets that perform black and white conversion, keeping some hues saturated.

5.1.5. The sophisticated way: channel mixer module

To perform such conversion we use the channel mixer module.

1. Activate the channel mixer module

2. Select the gray output channel

3. Set the proportion of each color, the sum having to equal 1 if you want to keep your global lightness.

5.2. Cross-processing

5.2.1. Overview

Cross-processing is a analog processing technique where slide film (normally developped thanks to an E6 solution) is processed in chemicals used for processing print film (C41). The resulting images get skewed colors usually a cyan hue and increased contrast and saturation.

The standard way for doing digital cross-processing is to use a channel curve tool but darktable lacks this tool for the moment and another way to accomodiate the effect is used.

5.2.2. Procedure

This procedure uses tonecurve, channelmixer and splittoning modules.

1. Image preparation

   Prepare the image for the cross process steps by adjusting the base settings such as exposure,whitebalance etc. for a correctly looking image.

2. Boost contrast

   Select the medium contrast curve preset for tonecurve module to boost the overall contrast in the image. You later might get back here to tune the curve for better result.

3. Color cast

   This step does change the colorcast as the base for the effect using the channel mixer. You might later get back to this and finetune the colorcast of the final result.

   a. Enable the channel mixer module

   b. Select blue channel and set blue color value to 0.8

   c. Select red channel and change blue color value to 0.1

   d. Select green channel and change blue color value to 0.1

4. Splittoning

   We use splittoning to add some more coloring to the result for cyan/blue shadows and yellow highlight.

   a. Enable the splittoning module

   b. Select a cyan/blue tone for shadows and set saturation around 50%

   c. Select a yellow/orange tone for highlights and set saturation around 70%

   d. Set compression to 10%

   e. Use the balance slider to tune the splittoning effect due this differs on every image due to it's exposure, motive etc.

5.3. Cyan toned image

5.3.1. Overview

Cyan is a nice color touchup for black and white im-
ages, this example guides you through how to make
this with darktable and how to control the tone. You
can choose any tone of your like for this tutorial...

5.3.2. Procedure

This procedure uses tonecurve, channelmixer and splittoning modules.

1. Image preparation

   Prepare the image for the cyan toned steps by adjusting the base settings such as
   exposure,black level, contrast etc. for a correctly looking image.

2. Black and white

   Enable the monochrome module to make the image black and white.

3. Add color tone

   This step selects the basetone of the image using channel mixer, we are going for cyan
   tone but you can choose the tone of your like here.

   a. Enable the channel mixer module

   b. Select red channel and set red color value to 0.7

   c. Select green channel and red color value to 0.150

   d. Select blue channel and red color value to 0.150

   As you notice we mix blue and green color to get a cyan tone, we substract 0.3 from
   red channel and add them to blue and green.

4. Splittoning

   The result of previous step does also add a colorcast on highlight that we actually want
   to have white for a prettier result. We also wants to add some blue color cast to the
   shadows to emphasis them.

   a. Enable the splittoning module

   b. Select a blue/cyan tone for shadows and set saturation around 50%

   c. Set highlights saturation to zero, to remove saturation on highlights.

   d. Set compression to zero

   e. Use the balance slider to tune the effect, our example uses a balance of 70/30

Chapter 4. Tethering

The tethering view is where you remote control and capture images from your connected camera directly into darktable.
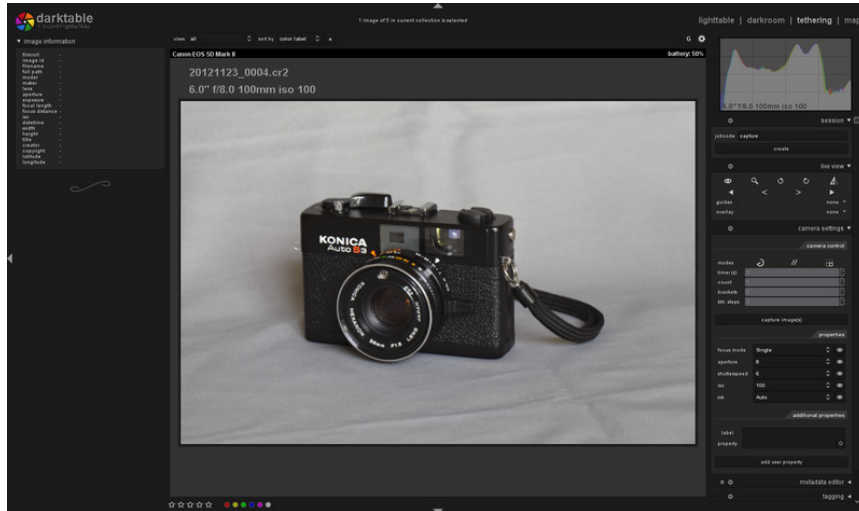
1. Overview

To use the tethering feature you need to connect your camera to your PC using an USB cable. Your desktop environment might ask a question of what you want to do. In that case *unmount/eject* is the answer. This is required to unlock the camera so darktable can lock it for usage.

When your camera is connected, head over to the import panel in lighttable. If your camera is not visible in this panel, hit the 'scan devices' button and it will appear with two functions: "import from camera" and "tethered shoot", the latter button is the one you should press to enter tethering mode.

darktable uses gphoto2 to interface with your camera. If you have problems to find the connected camera as described above, do check the troubleshoot section in this chapter in order to verify that your camera is supported for use with tethering.

1.1. Tethering

In the center view images are shown while you capture them. You can get an exposure by either using darktable's userinterface or manually triggering a capture on your camera. If you are using LiveView it will be shown in darktable's center view.
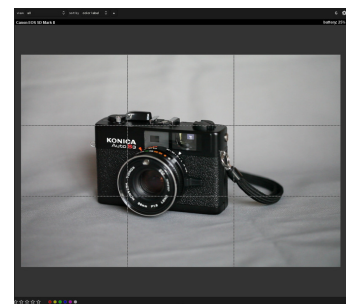
When entering tethering view, a filmroll will be created using the same structure as defined when you import from the camera. Job code will be predefined as "capture".

If you want to group your captures into different filmrolls, you should use the session panel in right side. When entering a new name and hitting enter, a new filmroll will be created and captured images will go into this new filmroll.

darktable provides some nifty tools to setup a capture using the user interface. You can setup timelapse captures and brackets for HDR creations. The configuration is so dynamic that you can create sequential capture of brackets - go figure... For more information read the documentation about capture panel and the examples in this chapter.

2. Panels

This section contains documentation for panels that are specific to the tethering view.

2.1. Session

A session is a sequence of exposures taken in tethering mode and going into a single filmroll. A new session equals a new filmroll. A filmroll is created with the same storage structure that is used when you import images from camera into darktable.

It's a bit akward, but configuring this storage structure is done in the camera import dialog for now.
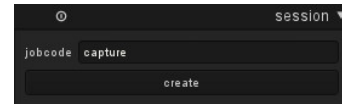
2.2. Live view

This panel gives you access to control your camera's live view mode. Functionality such as focus setting, rotation, adding guides and overlays are supported.
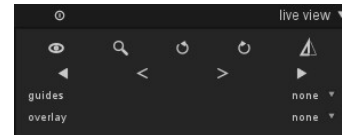
2.3. Camera settings

In the camera settings you have access to functionality as setting up a capture job which includes sequences, brackets and delays. Also this panel includes access to camera settings that can be controlled remotely from within darktable.
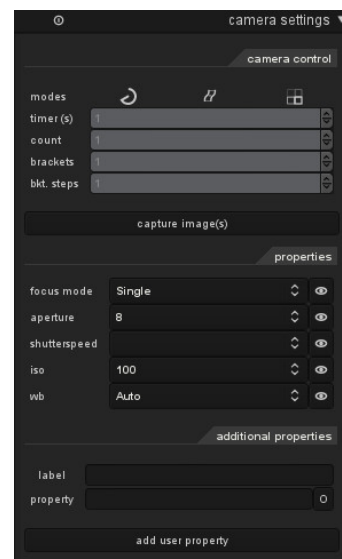
3. Examples

This section contains examples of typical usages of tethering.

3.1. Studio setup with screening

This is a pretty common use case. You have your studio and subject rigged, camera is connected to your computer and tethering view is active in darktable. You work behind the camera and take images. Whenever you want you can screen the image directly on your computer monitor instead of using the camera LCD for validation.

This workflow is effective, as you don't need to take 20 captures and then later - while out of the rigg - review them and hope that you got "the one". Just do it up front. If you're shooting a model this is a pretty nice way to pre-show the captures to the client instead of fumbling around with your camera.

This workflow is kind of out-of-the-box when entering tethering view. You just need to set a session name - preferable subject name of the shoot - then just start to take images. They will fall into the right filmroll for the session for easy review on-site.

3.2. Capturing a timelapse

A timelapse is a video clip composed out of images taken on a interval of X seconds. A typical example is to take a timelapse of cityscapes where you capture clouds and traffic etc.

To setup a timelapse capture you should as always create a new session as descibed before. Then you need to decide if you want to shoot in manual or auto mode. Only go for auto in situations were the ambient light will strongly change over time of the timelapse. Eg. shooting a timelapse over 24 hours might give you easier control of lights in that kind of captured sequence.

Head over to the camera settings panel where you define delay and sequence. Sequence will give you the oppertunity to choose how many images you want to capture and delay will set the time in seconds between capture.

To start the capture just hit the capture button in the same panel and watch the filmstrip fill up with images. The latest captured image is always previewed in center view.

4. Troubleshoot

4.1. Verify that your camera is supported

This troubleshoting guide will give you steps how to verify that your camera can be used with tethering. This is mainly done using the gphoto2 commandline tools as this is what darktable uses to interface with your camera.

1. Verify that camera is detected

   The following command will provide which cameras are detected and connected to the computer. Find your camera port name to use it in the following tests below. Usally port **usb:** will be enough and therefore used in these examples.

   ```
   env LANG=C gphoto2 --auto-detect
   ```

2. Verify camera driver abillities

   Execute the following command and verify that the "Capture choices" ability supports *Image* and "Configuration support" is *yes*. darktable will check these two abilities and decide if "tethered shoot" button should be shown or not.

   ```
   env LANG=C gphoto2 --port usb: --abilities
   ```

3. Verify camera remote capture

   This step will verify that your camera can be remotely controlled; that it can capture an image and then download it to your computer by software.

   ```
   env LANG=C gphoto2 --port usb: --capture-image-and-download
   ```

4. Verify camera tethered capture

   And this last step tests if your camera supports events which darktable heavily relies on. Running this command will make the gphoto2 process wait for an image capture event which you must manually trigger on your camera. On success an event is fetch and an the image will download.

   ```
   env LANG=C gphoto2 --port usb: --capture-tethered
   ```

4.2. Soo now what?

If any of the steps above failed, there are problems with your specific camera and driver. Please report the issues to gphoto2 mailing list for further help. Add the following flags to the command above that failed for better support and attach the log output to your mail:

```
--debug --debug-file gphoto2_debug.log
```

If you successfully went through all the tests above, your camera will likely be supported by darktable. If you anyway stumble into a problem in darktable please file a bug at http://www.darktable.org/redmine. Please attach outputs of the steps above and the log file output while running darktable with the following command.

```
darktable -d camctl 2>1 >camctl.log
```

Chapter 5. Map

The Map view is were you geotag your images.
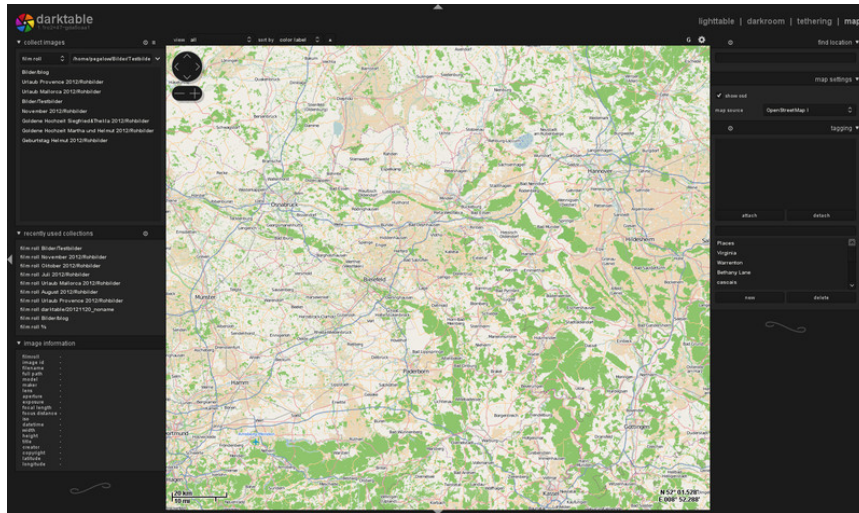
1. Overview

Map view will show you a world map and pin images that are geotagged into the map. There are several processes to get this happen. Some cameras have an internal GPS that tags images with location data; such as smartphones. Other cameras such as DSLR needs some additional GPS hardware to do this.

Even if you don't have hardware that supports this: darktable is able to apply GPX track data onto images. While you are on a walk and you either have a GPS tracker recording your movements or you use the right app on your smartphone with GPS functionality you can record a GPX track. When you are taking images with your camera on your walk, the timestamps of images are stored in EXIF data. They can later be matched with the GPX track data and therefore a location can be estimated. Matching image timestamps with GPX track data can be done from within lightable mode.

1.1. Center map view

In the center of the map view you will see a map.
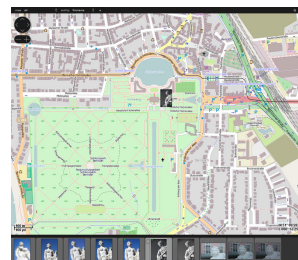
Map data are taken from open map sources on the internet. New map data are only available if you are connected to the internet. darktable keeps a disk cache of already loaded map data.

You can navigate in your map using the mouse. Left-click will drag the map; using the scroll-wheel will zoom in or out.

Some on-screen controls and displays assist you to find your way. There is a navigation area on top left. Use it as an alternative to mouse-dragging and scrolling. The scale of your map is displayed on bottom left. On bottom right you see the geographical coordinates of the center of the map.

Images that already have geo location attributes in their metadata are displayed as small icons on the map.

In order to assign geo coordinates to an image, activate the film-strip on the lower panel (press *Ctrl-f*). You can simply assign a geo location to an image by dragging the image icon from the film-strip and position it onto the map. darktable will record the new geo location (longitude and latitude) as part of the image metadata. Exported images will inherit these settings.

Left and right to the central map there are panels for additional control.

2. Panels

This section contains documentation for panels that are specific to the map view.

2.1. Left panels

The panels on the left side we already know from lighttable mode.

You can define image collections with the "collect images" panel where you can define sets of selection rules. Recently used collections can be reached with their respective names in separate panel. You can also get an overview of the metadata of the image under your mouse cursor in panel "image information".
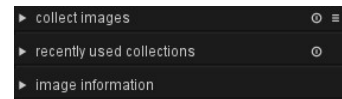
2.2. Find location

Use this modules to search for a place on map. You need to be connected to internet to use this feature.

To use, just name a place and hit enter and a list of results will be shown. Click on one of the resulting items and the map will zoom into the place; you can then start to drop images into that location for geotagging.
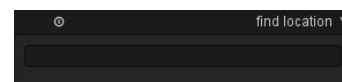
2.3. Map settings

In this panel you control the map settings. You can select the source for map data from different providers and what layers they bring such as satelite view etc.

You can also toggle the view of some overlay info.
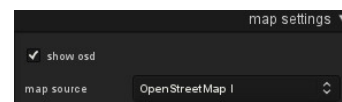
2.4. Tagging

This panel gives you access to darktables tagging features. It is equivalent to the one you reach from within lighttable mode. See chapter 2 for more details.
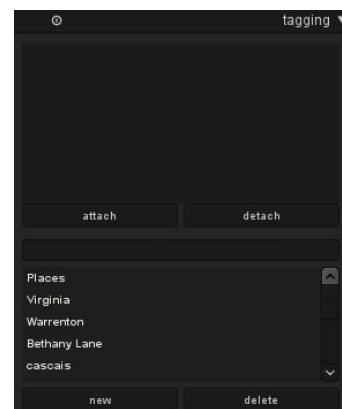
Chapter 6. Preferences and Settings

darktable comes with a number of settings that can be configured by users. You reach the configuration system by clicking ⚙.
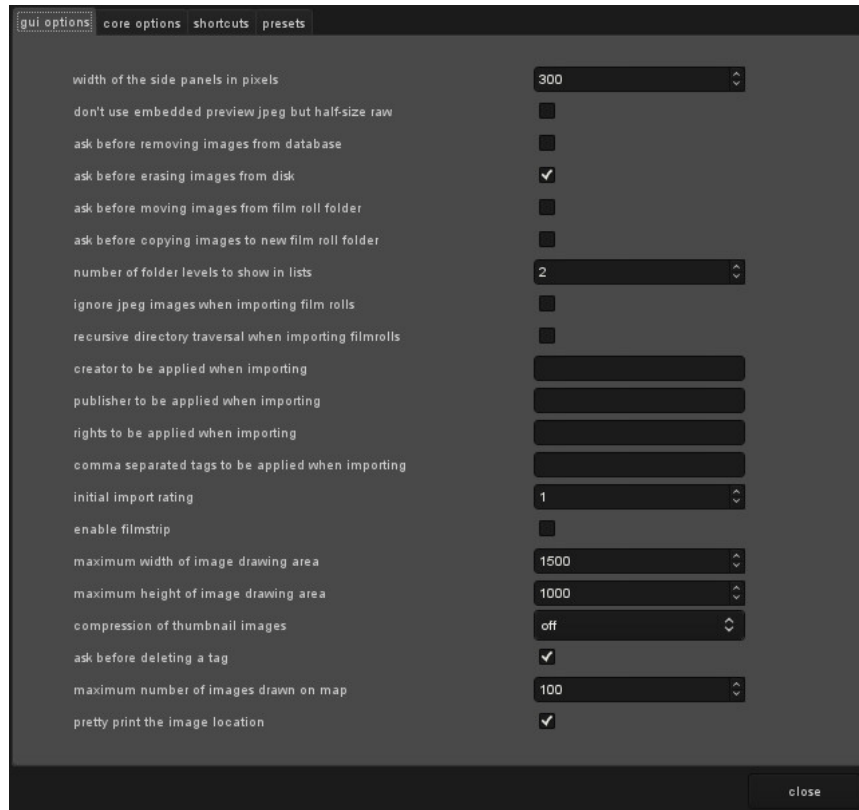
1. Gui options

These options control the look and feel of darktable.

width of the side panels in pixels

this controls the size of size panels in pixels; side panels are found left and right to the center view (default 300).

don't use embedded preview jpeg but half-size raw

check this option to not use the embedded jpeg from the raw file but process the raw data; this is slower but gives you color managed thumbnails (default off).

ask before removing images from database

always ask the user before any image is removed from db (default on).

ask before erasing images from disk

always ask the user before any image file is deleted (default on).

ask before moving images from film roll folder

always ask the user before any image file is moved (default on).

ask before copying images to new film roll folder

always ask the user before any image file is copied (default on).

number of folder levels to show in lists

the number of folder levels to show in film roll names, starting from the right (default 1).

ignore jpeg images when importing film rolls

when having raw+jpeg images together in one directory it makes no sense to import both; with this flag one can ignore all jpegs found (default off).

recursive directory traversal when importing filmrolls

not only import images from the directory selected but recursively go through all subdirectories as well (default off).

creator to be applied when importing

if given automatically add this string as a creator tag when importing images (default none).

publisher to be applied when importing

if given automatically add this string as a publisher tag when importing images (default none).

rights to be applied when importing

if given automatically add this string as a copyrights tag when importing images (default none).

initial import rating

initial star rating (from 0 to 5) for all images when importing a filmroll (default 1).

enable filmstrip

enable the filmstrip in darkroom, tethering and geomapping modes (default on).

maximum width of image drawing area

maximum width of the image drawing area in darkroom mode. adjust to your screen; needs a restart and will invalidate current thumbnail caches (default 1300).

maximum height of image drawing area

maximum height of the image drawing area in darkroom mode. adjust to your screen; needs a restart and will invalidate current thumbnail caches (default 1000).

compression of thumbnail images

controls the compression of thumbnail images in memory and on disk; options: "off", "low quality (fast)", "high quality (slow)"; (default off).

ask before deleting a tag

always ask user before deleting a tag from an image (default on).

maximum number of images drawn on map

the maximum number of geotagged images drawn on the map. increasing this number can slow drawing of the map down; needs a restart if changed (default 100).

pretty print the image location

show a more readable representation of the geo location in the image information module (default on).

2. Core options

These options control some of the internals of darktable.

memory in bytes to use for mipmap cache

in order to speed-up display of filmrolls, darktable stores thumbnails in a cache on disk and loads it into memory at startup; this controls the cache size in bytes; needs a restart if changed (default 536870912).

number of background threads

this controls how many parallel threads are used to create thumbnails during import; on 32bit systems it is strongly recommended to set this to 1; needs a restart if changed (default 2).

export multiple images in parallel

set this variable to the desired number of threads if you want multithreaded export to process multiple images at a time; be warned: every thread will need about 1GB of memory; setting this to 1 switches on per-image parallelization, which darktable does very efficiently and which normally obsoletes parallel exports; 32bit systems need to set this to 1 (default 1).

host memory limit (in MB) for tiling

in order to manage large images on systems with limited memory darktable does a tile-wise processing; this variable controls the maximum amount of memory (in MB) a module may use during image processing; lower values will force memory hungry modules to process image with increasing number of tiles; setting this to 0 will omit any limits; values below 500 will be treated as 500; on a 32bit system you should set this to 500; needs a restart if changed (default 1500).

minimum amount of memory (in MB) for a single buffer in tiling

if set to a positive, non-zero value this variable defines the minimum amount of memory (in MB) that darktable should take for a single tile; on a 32bit system you should set this to 8; 64bit systems can live with higher values; needs a restart if changed (default 16).

write sidecar file for each image

these redundant XMP files can later be re-imported into a different database, preserving your changes to the image; it's strongly recommended to have this option activated so you don't lose data in case of a database corruption; backing up your RAW file plus the accompanying XMP file will allow you to fully restore your work (default on).

activate opencl support

darktable can use your GPU to speed up processing significantly; interface OpenCL requires a suited hardware and matching OpenCL drivers on your system; if one of those is not found the option is greyed out; can be switched on and off at any time and takes immediate effect (default on).

always use littlecms2 during export

if this option is activated, darktable will use system library littlecms2 instead of its own routines; this is about 28x as slow as the default but might give more accurate results in some cases (default off).

do high quality resampling during export

the image will first be processed in full resolution, and downscaled at the very end. this can result in better quality sometimes, but will always be slower (default off).

demosaicing for zoomed out darkroom mode

interpolation when not viewing 1:1 in darkroom mode: "always bilinear (fast)" is fastest, but not as sharp. "at most ppg (reasonable)" is using ppg + interpolation modes specified below, "full (possibly slow)" will use exactly the settings for full-size export (default "at most ppg (reasonable)").

pixel interpolator

pixel interpolator used in rotation, lens correction, up- and downscaling; options are "bilinear", "bicubic", "lanczos2", "lanczos3" (default "lanczos3").

password storage backend to use

the storage backend for password storage; options: "none", "kwallet", "gnome keyring" (default none).

run indexer

this will index images in background, so you can search them by similarity (default off).

3. Shortcuts

darktable has a big set of keyaccelerators that with the work of Robert Bieber now is user configurable through the preference pane.

When you open the shortcuts menu you are presented with a hierarchical list of all actions that can receive a keyaccelerator. Go to the action you want to change and double click. You are then prompted to press the new key combination to be mapped to the selected action. In order to remove an existing keyaccelerator, click on the action and press back-space.

You can export your mappings to a file or import mappings from a file. Press "default" to reset all keyaccelerators to their default state.

Here follows a table with the default keybindings of actions available in darktable.

<darktable>/capture/toggle film strip	<Control>f
<darktable>/darkroom/export	<Control>e
<darktable>/darkroom/image back	BackSpace
<darktable>/darkroom/image forward	space
<darktable>/darkroom/plugins/clipping/commit	Return
<darktable>/darkroom/plugins/clipping/undo	<Control>z
<darktable>/darkroom/plugins/colorout/toggle softproofing	s
<darktable>/darkroom/toggle film strip	<Control>f
<darktable>/darkroom/zoom/close-up	<Alt>1
<darktable>/darkroom/zoom/fill	<Alt>2
<darktable>/darkroom/zoom/fit	<Alt>3
<darktable>/film strip/color/blue	F4
<darktable>/film strip/color/green	F3
<darktable>/film strip/color/purple	F5
<darktable>/film strip/color/red	F1
<darktable>/film strip/color/yellow	F2
<darktable>/film strip/history/copy	<Control>c
<darktable>/film strip/history/discard	<Control>d
<darktable>/film strip/history/paste	<Control>v
<darktable>/film strip/rating/1	1
<darktable>/film strip/rating/2	2
<darktable>/film strip/rating/3	3
<darktable>/film strip/rating/4	4
<darktable>/film strip/rating/5	5
<darktable>/film strip/rating/desert	0
<darktable>/film strip/rating/reject	r
<darktable>/film strip/scroll back	Left
<darktable>/film strip/scroll forward	Right
<darktable>/global/capture view	t
<darktable>/global/darkroom view	d

<darktable>/global/decrease brightness	F9
<darktable>/global/decrease contrast	F7
<darktable>/global/increase brightness	F10
<darktable>/global/increase contrast	F8
<darktable>/global/leave fullscreen	Escape
<darktable>/global/lighttable view	l
<darktable>/global/quit	<Control>q
<darktable>/global/switch view	period
<darktable>/global/toggle fullscreen	F11
<darktable>/global/toggle side borders	Tab
<darktable>/lighttable/color/blue	F4
<darktable>/lighttable/color/green	F3
<darktable>/lighttable/color/purple	F5
<darktable>/lighttable/color/red	F1
<darktable>/lighttable/color/yellow	F2
<darktable>/lighttable/navigate/down	<Shift>g
<darktable>/lighttable/navigate/page down	Page_Down
<darktable>/lighttable/navigate/page up	Page_Up
<darktable>/lighttable/navigate/up	g
<darktable>/lighttable/plugins/image/delete from disk	
<darktable>/lighttable/plugins/image/remove from collection	Delete
<darktable>/lighttable/preview	z
<darktable>/lighttable/rating/1	1
<darktable>/lighttable/rating/2	2
<darktable>/lighttable/rating/3	3
<darktable>/lighttable/rating/4	4
<darktable>/lighttable/rating/5	5
<darktable>/lighttable/rating/desert	0
<darktable>/lighttable/rating/reject	r
<darktable>/lighttable/scroll/center	apostrophe
<darktable>/lighttable/scroll/down	Down
<darktable>/lighttable/scroll/left	Left
<darktable>/lighttable/scroll/right	Right
<darktable>/lighttable/scroll/up	Up
<darktable>/lighttable/zoom/in	<Alt>2
<darktable>/lighttable/zoom/max	<Alt>1
<darktable>/lighttable/zoom/min	<Alt>4
<darktable>/lighttable/zoom/out	<Alt>3

4. Presets

This menu gives you an overview of the presets that are defined for darktable's modules. darktable already comes with a set of pre-defined presets for several modules. In addition you can define your own presets from within each module in darkroom mode. In this dialog you can select whether a certain user defined preset shall be auto-applied to matching images.

Pre-defined presets are shown with a lock symbol. Their auto-apply properties can not be changed.

Double clicking on a user-defined preset will open a menu.
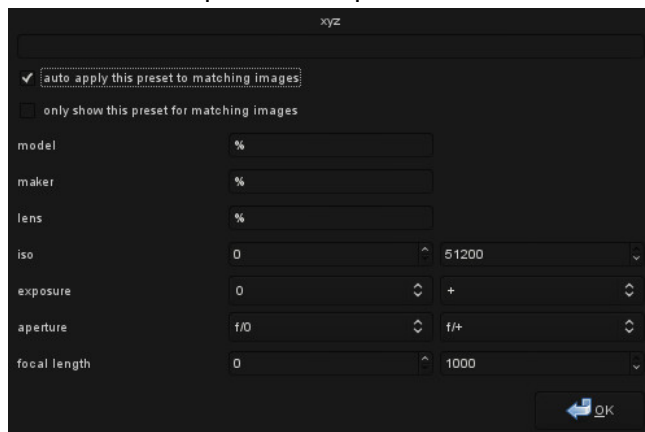
auto apply this preset to matching images

activate this checkbutton to make the preset be auto applied to newly imported images; a set of fields is displayed where you can define patterns to be matched against EXIF data.

only show this preset for matching images

activate this checkbutton to hide the preset in darkroom mode if it does not match the defined patterns.

model

a pattern to be matched against the EXIF field that describes your camera model; use "%" as wildcard.

maker

a pattern to be matched against the EXIF field that describes the maker of your camera; use "%" as wildcard.

lens

a pattern to be matched against the EXIF field that describes your lens; use "%" as wildcard.

iso

only apply the preset if the iso value of your image lies within the given range.

exposure

> only apply the preset if the exposure time of your image lies within the given range; set "+" as the upper value to match arbitrarily long exposures.

aperture

> only apply the preset if the aperture of your image lies within the given range; set "f/0" as the lower value to match arbitrarily open apertures; set "f/+" as the upper value to match arbitrarily closed apertures.

focal length

> only apply the preset if the focal length of your image lies within the given range (from 0 to 1000).

Chapter 7. Special Topics

darktable's home page at www.darktable.org has an extensive blog section. It covers many technical details, background information as well as tips and tricks. We replicate a few of our blog articles here as they might be helpful to get things running.
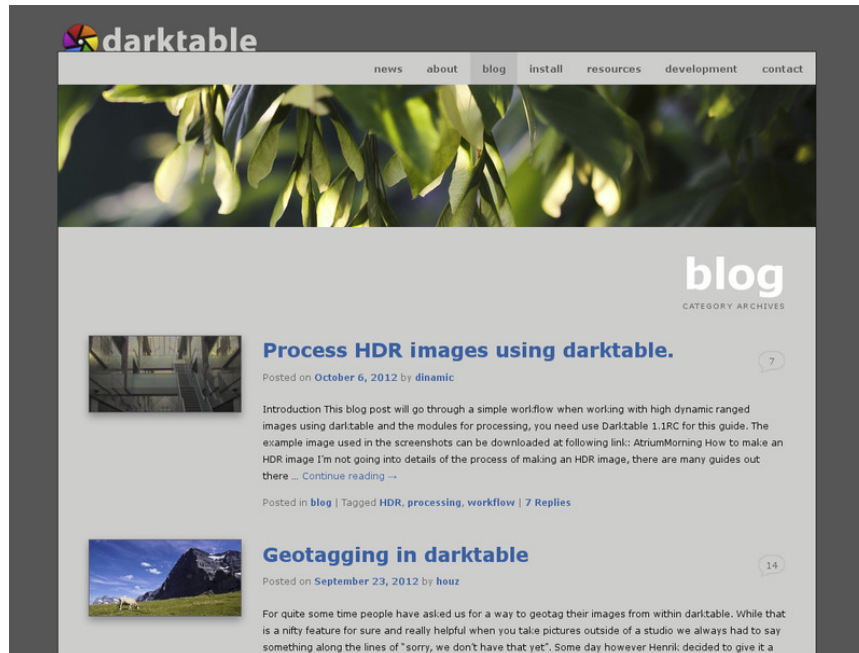
1. darktable and memory

darktable's memory requirements are high. A simple calculation makes this clear. If you have a 20MPx image, darktable for precision reasons will store this internally as a 4 × 32-bit floating point cell for each pixel. Each full image of this size will need about 300MB of memory. As we want to process the image, we will at least need two buffers for each module – one for input and one for output. If we have a more complex module, its algorithm might additionally require several intermediate buffers of the same size. Without further optimization, anything between 600MB and 3GB would be needed only to store and process image data. On top we have darktable's code segment, the code and data of all dynamically linked system libraries, and not to forget further buffers where darktable stores intermediate images for quick access during interactive work (mip map cache). All in all darktable would like to see a minimum of about 4GB to run happily.

1.1. Total system memory

From what I said before, it is evident that your computer needs a sane memory setup to properly run darktable. We suggest that you have a least 4GB of physical RAM plus 4 to 8GB of additional swap space installed. The latter is required, so that your system can swap out temporarily unneeded data to disk in order to free physical RAM.

Theoretically you could also run darktable with lower amounts of physical RAM and balance this with enough swap space. However, you should be prepared that your system could then heavily go into "thrashing" action, as all too many memory accesses require your system to read or write data pages from/to hard disk. We have positive reports that this functions well for several users, but it still might get extremely slow for others...

1.2. Available address space

Besides the total amount of system memory there is another limiting factor: the available address space of your hardware architecture. How much memory can be addressed by a process depends on the number of address bits your CPU offers. For a CPU with 32-bit address registers, this is 2^{32} bytes, which makes a total of 4GB. This is the absolute upper limit of memory that can be used by a process and it constitutes a tight situation for darktable as we have seen above.

darktable's escape route is called tiling. Instead of processing an image in one big chunk, we split the image into smaller parts for every processing step (module). This will still require one full input and output buffer, but intermediate buffers can be made small enough to have everything fit into the hardware limits.

1.3. Memory fragmentation

Unfortunately this is not the full story yet. There is an effect called memory fragmentation, which can and will hit software that needs to do extensive memory management. If such a program allocates 5 times 300MB at a time and frees it again, that memory should normally be available for one big 1.5GB allocation afterwards. This however is often not the case. The system's memory allocator might no longer see this area as one contiguous 1.5GB block but as a row of 300MB areas. If there is no other free area of 1.5GB available, the allocation would fail. During a program run this mechanism will take away more and more of the larger memory blocks in favor of smaller ones. darktable 1.0 introduces a caching algorithm to address this problem. It pre-allocates blocks of memory and makes them available on request.

1.4. Further Limitations

And as if this were not challenging enough, there are further things that might limit your access to memory. On some older boards you need to activate BIOS option "memory remapping" in order to have all physically installed memory enabled. In addition if you are on a 32-bit OS you will probably need a kernel version that has "Physical Address Extension" (PAE) enabled. This is often but not always the case for Linux. Many distributions deliver different kernels, some with and some without PAE activated; you need to choose the right one. A typical sign that one of these topics are not setup correctly is system command free (1) not reporting all of your installed physical RAM; for example you have 4GB on your board, but your kernel is only seeing 3GB or less. You need to consult the manual of your BIOS and the information about your Linux variant for further help.

1.5. Setting up darktable on 32-bit systems

As we've seen 32-bit systems are difficult environments for darktable. Still many users are successfully running darktable on them, if the basic requirements in terms of total system memory and the topics mentioned in the paragraphs above are addressed properly.

There are several adjustment parameters to get it running. If you make a fresh install, darktable will detect your system and set conservative values by default. However, if you upgrade darktable from an older version (e.g. coming from 0.9.3 and going to 1.0), chances are you have unfavorable settings in your preferences. The consequences might be darktable aborting due to allocation failures or – very typically – darktable not being able to properly import a new film roll. As a frequent symptom you get skulls displayed instead of thumbs for many of your pictures.

Take a minute to optimize the preference settings in this case. You will find them under "core options" in darktable's preference dialog. You might also find these parameters as configuration variables in $HOME/.config/darktable/darktablerc and edit them there.

Here is a short explanation of the relevant parameters and their proposed settings.

number of background threads

This parameter defines the maximum number of threads that are allowed in parallel when importing film rolls or doing other background stuff. For obvious reasons on 32-bit systems you can only have one thread eating resources at a time. So you need set this parameter to 1; anything higher will kill you. For the same reason you also must set the number of parallel export threads to 1.

host memory limit (in MB) for tiling

This parameter tells darktable how much memory (in MB) it should assume as available to store image buffers during module operations. If an image can not be processed within these limits in one chunk, tiling will take over and process the image in several parts, one after the other. Set this to the lowest possible value of 500 as a starting point. You might experiment later whether you can increase it a bit in order to reduce the overhead of tiling.

minimum amount of memory (in MB) for a single buffer in tiling

This is a second parameter that controls tiling. It sets a lower limit for the size of intermediate image buffers in megabytes. The parameter is needed to avoid excessive tiling in some cases (for some modules). Set this parameter to a low value of 8. You might tentatively increase it to 16 later.

memory in bytes to use for mipmap cache

> This controls how many thumbnails (or mip maps) can be stored in memory at a time. As a starting point set this to something like 256MB (give the number in bytes). To avoid the problem of memory fragmentation during longer runs of darktable, the new caching scheme frontloads the memory costs and allocates this cache once at the beginning. Some Linux kernels use over-committing memory allocation, which means you don't immediately pay for all of the memory in terms of RSS (resident set size, the non-swapped physical memory), but in any case you pay for the address space. As explained before, this poses a problem for 32-bit systems and will at first sight appear as a regression over the (terribly slow, but that's another story) 0.9.3-style cache. In the long run however, this is all memory that's ever going to be allocated for thumbnails, so if we can successfully grab this portion once, we are relieving a lot of pressure on fragmentation for long sessions.

1.6. darktable on 64-bit systems

> There's not much to be said here. Of course also 64-bit systems require a sufficient amount of main memory. So the 4GB plus swap recommendation holds true. On the other hand, 64-bit architectures do not suffer from the specific 32-bit limitations like small address space and fragmentation madness.

> Most modern Intel or AMD 64-bit CPUs will have an available address space in the range of several Terabytes. The word "modern" is relative in this context: all AMD and Intel CPUs introduced since 2003 and 2004, respectively, offer a 64-bit mode. Linux 64-bit has been available since many years.

> All relevant Linux distributions give you the free choice to install a 32-bit or a 64-bit version with no added costs. You can even run old 32-bit binaries on a 64-bit Linux. The only thing you need to do: invest some time into the migration. In the end we strongly recommend to move to a 64-bit version of Linux. There isn't much good reason, not to make this step.

2. darktable and OpenCL

darktable can use GPU acceleration via OpenCL to improve performance.

2.1. The Background

Processing high resolution images belongs to the more demanding tasks in modern computing. Both, in terms of memory requirements and in terms of CPU power, getting the best out of a typical 15, 20 or 25 Megapixel image can quickly bring your computer to its limits.

darktable's requirements are no exception. Our decision to not compromise processing quality, has led to all calculations being done on 4 × 32bit floating point numbers. This is slower than "ordinary" 8 or 16bit integer algebra, but eliminates all problems of tonal breaks or loss of information.

A lot of hand optimization has been invested to make darktable as fast as possible. If you run a current version of darktable on a modern computer, you might not even notice any "slowness". However, there are conditions and certain modules where you feel (or hear from the howling of your CPU fan) how much your poor multi-core processor has to struggle.

That's where OpenCL comes in. OpenCL allows us to take advantage of the enormous power of modern graphics cards. It has been gamer's demand for more and more highly detailed 3D worlds in modern ego shooters, that has fostered GPU development. ATI, NVIDIA and Co had to put enormous FPU processing power into their GPUs to meet these demands. The result is modern graphics cards with highly parallelized GPUs to quickly calculate surfaces and textures at high frame rates.

You are not a gamer and you don't take advantage of that power? Well, then you should at least use it in darktable! For the task of highly parallel floating point calculations modern GPUs are much faster than CPUs. That is especially true, when you want to do the same few processing steps over millions of items. Typical use case: processing of megapixel images.

2.2. How OpenCL works

As you can imagine, hardware architectures of GPUs can vary significantly. There are different producers, and even different generations of GPUs from the same producer may differ clearly. At the same time GPU manufacturers are normally not willing to disclose many hardware details of their products to the public. One of the known consequences is the need to use proprietary drivers under Linux, if you want to take full advantage of your graphics card.

Fortunately an industry consortium lead by The Khronos Group has developed an open, standardized interface called OpenCL. It eases the use of your GPU as a numerical processing device. OpenCL offers a C99-like programming language with a strong focus on parallel computing. An application that wants to use OpenCL will need to bring along a suited OpenCL source code that it then hands over to a hardware specific OpenCL compiler at run-time. This way the application can use OpenCL on different GPU architectures (even at the same time). All "hardware secrets" are hidden in this compiler and are normally not visible to the user (or the application). The compiled OpenCL code is loaded onto your GPU and - with certain API calls - it is ready to do calculations for you.

2.3. How to activate OpenCL in darktable

Using OpenCL in darktable requires that your PC is equipped with a suitable graphics card and that it has the required libraries in place. Namely modern graphics cards from

NVIDIA and ATI come with full OpenCL support. The OpenCL compiler is normally shipped as part of the proprietary graphics driver; it is reachable as a dynamic library called "libOpenCL.so". This library must be in a folder where it is found by your system's dynamic linker.

When darktable starts, it will first try to find and load libOpenCL.so and – on success – check if the available graphics card comes with OpenCL support. A sufficient amount of graphics memory (1GB+) needs to be available to take advantage of the GPU. If that is OK, darktable tries to setup its OpenCL environment: a processing context needs to be initialized, a calculation pipeline to be started, OpenCL source code files (extension is .cl) need to be read and compiled and the included routines (called OpenCL kernels) need to be prepared for darktablele's modules. If all that is done, the preparation is finished.

Per default OpenCL support is activated in darktable if all the above steps were successful. If you want to de-activate it you can do so in "core options" by unchecking "activate opencl support". This configuration parameter also tells you if OpenCL initialization failed: it is grayed out in that case.

You can at any time switch OpenCL support off and on; this will take immediate action. Depending on the type of modules you are using, you will notice the effect as a general speed-up during interactive work and during export. Most modules in darktable can take advantage of OpenCL but not all modules are demanding enough to make a noticeable difference. In order to feel a real difference, take modules like "shadows and highlights", "sharpen", "lowpass", "highpass" or as an extreme case "equalizer".

If you are interested in profiling figures, you can call darktable with command line parameters -d opencl -d perf. After each run of the pixelpipe you will get a detailed allocation of processing time to each module plus an even more fine grained profile for all used OpenCL kernels.

Besides the speed-up you should not see any difference in the results between CPU and GPU processing. Except of rounding errors, the results are designed to be identical. If, for some reasons, darktable fails to properly finish a GPU calculation, it will normally notice and automatically (and transparently) fall back to CPU processing.

2.4. Possible Problems and Solutions

darktable will detect OpenCL run-time problems automatically. It will then reprocess everything on CPU; only speed is affected, the final result should not be endangered.

There can be various reasons why OpenCL could fail during initialization phase. We depend on hardware requirements and on the presence of certain drivers and libraries. In addition all these have to fit in terms of maker model and revision number. If anything does not fit, e.g. your graphics driver (loaded as a kernel module) does not match the version of your libOpenCL.so, OpenCL support is likely not available.

In that case, the best thing to do is start darktable from a console with

```
darktable –d opencl
```

This will give additional debugging output about the initialization and use of OpenCL. First see if you find a line that starts with "[opencl_init] FINALLY ..." This should tell you, if OpenCL support is available for you or not. If initialization failed, look at the messages above for anything that reads like "could not be detected" or "could not be created". Check if there is a hint about where it failed.

Here are a few cases observed in the past:

darktable might tell you that no OpenCL aware graphics card is detected or that the available memory on your GPU is too low and the device is discarded. In that case you might need to buy a new card, if you really want OpenCL support.

darktable might find your libOpenCL.so but then tell you that it couldn't get a platform. NVIDIA drivers will often give error code -1001 in that case. This happens because libOpenCL.so is only a wrapper library. For the real work further libraries - specific to vendor, device and driver - need to be loaded. This failed for some reason. There is a structure of files in /etc/OpenCL on your system that libOpenCL.so consults to find these libraries. Check if you find something fishy in there and try to fix it. Often the needed libraries cannot be found by your system's dynamic loader. Giving full path names might help.

darktable might also tell you that a context could not be created. This often indicates a version mismatch between (loaded) graphics driver and libOpenCL. Check if you have leftover kernel modules or graphics libraries of an older install and take appropriate action. In doubt, make a clean reinstall of your graphics driver. Sometimes, immediately after a driver update, the loaded kernel driver does not match the newly installed libraries: reboot your system in that case.

darktable might crash in very rare cases directly during startup. This can happen if your OpenCL setup is completely broken or if driver/library contains a severe bug. If you can't fix it, you can still use darktable with option --disable-opencl, which will skip the entire OpenCL initialization step.

darktable might on some systems fail to compile its OpenCL source files at run-time. In that case you will get a number of error messages looking like typical compiler errors. This could indicate an incompatibility between your OpenCL implementation and our interpretation of the standard. In that case visit us at darktable-devel@sourceforge.net and report the problem. Chances are good that we can help you. Please also report in case you see significant differences between CPU and GPU processing of an image!

There also exist a few on-CPU implementations of OpenCL. These come as drivers provided by INTEL or AMD. We observed that they do not give us any speed gain versus our hand-optimized CPU code. Therefore we simply discard these devices.

2.5. Setting up OpenCL for AMD/ATI devices

While NVIDIA devices will most often run out of the box, there is more to do for AMD/ATI graphics cards. This already starts with the fact that AMD/ATI will only report to darktable half of the total GPU memory. For a 1GB device this amounts to 512MB, a value which darktable in its standard configuration will refuse as not being sufficient for its tasks. Consequence: the device will not be used.

On the net you might find as a tip to set environment variable GPU_MAX_HEAP_SIZE=100 in this case. Indeed this will cause the AMD/ATI driver to report the full installed memory to darktable. However, there is a problem. On many (most?) cards this will cause buffers not to be allocated on your device but on host! In this case all memory transfers will need to go through the slow PCIe bus. This will cost you a factor of 10 or more in performance and will render OpenCL useless for you, especially when exporting files.

Another environment variable which changes driver behavior is GPU_MAX_ALLOC_PERCENT. You could set this to 100 in order to allow memory allocations as high as 1GB on your AMD/ATI card. Problem is that this simply tends not to work and cause darktable to crash sooner or later.

Our recommendation is to leave these settings untouched. Most often your card will be recognized with 512MB memory and a maximum allocation size of 128MB. There are three configuration parameter which you set in file $HOME/.config/darktable/darktablerc to get things running. Here are the details:

opencl_memory_requirement

Set this parameter to 500 so that darktable will accept your 512MB graphics memory as being sufficient in memory.

opencl_memory_headroom

This parameter controls how much graphics memory (out of the reported one) darktable should leave untouched for driver and display use. As for AMD/ATI devices we anyhow only can get half of the available RAM it's safe to set this to zero. So all of the 512MB can be used by darktable.

opencl_avoid_atomics

Atomic operations in OpenCL are a special way of data synchronization. They are only used in few kernels. Unfortunately some (most?) AMD/ATI devices are extremely slow in processing atomics. It's better to process the affected modules on CPU rather than accepting an ultra-slow GPU codepath. Set this parameter to TRUE if you experience slow processing of modules like shadows&highlights, monochrome, local contrast or global tonemap or if you even get intermittent system freezes.

2.6. OpenCL still does not run for me!

As has been said before OpenCL systems come with a huge variety of setups: different gpu manufacturers, different gpu models, varying amounts of gpu memory, different drivers, different distributions etc. Many of the potential problems will only appear with a very specific combination of those factors.

As we developers of darktable on our computers only have access to a small fraction of those variations, please understand that we might not be able to fix your specific problem. There is not much we can do, if there is no way for us to reproduce.

If nothing else helps, the best option might be to start darktable with

```
darktable --disable-opencl
```

In the end there is nothing in darktable which only runs on gpu. Don't let OpenCL discourage you; also darktable's cpu code is highly optimized and performant!