

Welcome to Luminance HDR!

Luminance HDR is an open source workflow tool for HDR imaging.

Its highlights are:

- streamlined user interface in a dozen of languages
- support for all major HDR file formats
- batch tonemapping
- works on Linux, Windows and Mac OS X
- free to distribute among family, friends and enemies

If you are new to HDR imaging, please read the [introductory chapter](#). If you already know what's good in HDR for you, you might like reading [detailed features overview](#). Finally, if you know your way around Luminance HDR, you can learn [what's new](#) in this version.

HDR Imaging Basics

There will be much more in this chapter, but for now we just list basic terminology:

HDR

Stands for "High Dynamic Range". An HDR image is an image which presents more than 8 bit per color channel. Most CRTs, LCDs and printers only have a limited dynamic range, and can display only LDR images (see below). Thus various methods of "converting" HDR images into a viewable format have been developed, generally called "tone mapping".

LDR

Stands for "Low Dynamic Range". The most common image formats, such as JPEG, PNG, GIF, ... have 8 bits per color channel, LDR is just another umbrella definition.

Tone mapping

A method of converting an HDR image into a LDR image. Various algorithms exist for this purpose, and in this context they are also known as "tone mapping operators", or in this manual simply as "operators". You can choose an operator from a list in the top of tone mapping options sidebar.

TMO

Shorthand for "Tone Mapping Operator".

Raw

Another umbrella definition for several (minimally processed) image formats. Raw files can have 12 or 14 bits per color channel, but noise usually cuts down the available dynamic range to something like 1000:1, roughly 10 bits. For all intents and purposes they are HDR files.

Antighosting

The HDR creation process involves merging a stack of images. An object changing position in the image set creates a strange effect in which the object is partially visible (like a ghost) in the final HDR image. This problem can be corrected by automatic or manual procedures.

Creating HDR images interactively

You can access the wizard that will guide you through the process of creating a new HDR image via the *File > New HDR image...* menu item.

On the first page the wizard will ask you to select the set of images (of the same scene, but taken at different exposures) that are going to contribute to the final hdr (supported input: jpeg, raw and tiff -8 and 16 bit-).

Even if Luminance HDR doesn't find the required Exif data (Shutter Speed & Aperture) in your image set you can still proceed creating an HDR. To do so you have to insert **manually** the EV (exposure values) or stop difference values for the images in your the set.

The first page of the wizard enables the user to apply an automatic alignment step to the images in the set. It is possible to use one of two alignment options (or "engines"): `align_image_stack` and `MTB`.

The first option is usually a good choice (`MTB` works only on LDR images and has a simpler model that does not take into account rotation).

If your image set consists of LDR images (JPEG, 8 bit TIFF, or RAW files converted to tiff without the `-4` switch) the second step of the wizard will show you a window that can be used to perform some "pre merging" editing activities.

The first feature will enable the user to add to each image in the set additional horizontal/vertical offsets.

The user is able to shift each image in set by one pixel by first selecting an image in the left (editable) list and then clicking one of the arrow buttons in the bottom-left panel. The same effect can also be obtained by using one of the `WSAD` or `IKJL` key combinations. Pressing `Shift` will shift by ten pixels, pressing `control` will shift by 50 pixels and pressing both `Shift` and `Control` will shift by 100 pixels.

By dragging a rectangle with the mouse, the user can also select a crop area for all the images in the set. Clicking the "Crop All Images" button will effectively crop the images in the set.

This window contains also the interactive anti-ghosting feature.

The user switches to anti-ghosting mode by pressing the "Anti Ghosting" button. It is then possible to draw a mask over the regions of the images that contain moving objects.

The user can select the size of the brush, its strength and whether the brush is adding or removing the mask.

The same window contains a button that enables the user to save the set of images after all the cropping has taken place and with all the anti-ghosting masks and shifts applied. The images will be saved in TIFF format, and the mask is saved as a 0-valued alpha channel in the image set.

The following page of the wizard will also ask you which kind of configuration you want to use to create the hdr from the source images: usually the first choice leads to good results, so feel free to accept the default "Config 1" option by pressing next.

If you think that the results can be improved you can try one of the other predefined configurations.

If you want to fiddle with settings you can check the "custom" option and another page of the wizard will show up with a complete set of configurations. It is worth mentioning that you should use the custom option as a last resort because the predefined configurations (which are a sub-set of the all possible configurations you can have in the custom configuration panel) are known to yield good results.

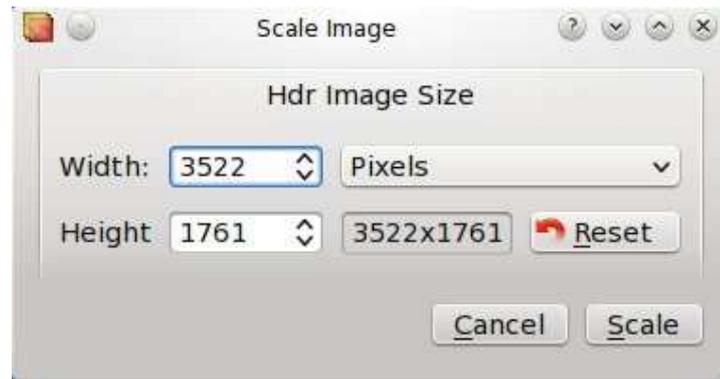
Editing HDR images

After creating an new HDR file or opening an existing one you can do several things to them except tonemapping.

Resizing

You can access this feature via the *Image > Resize the Hdr...* menu item.

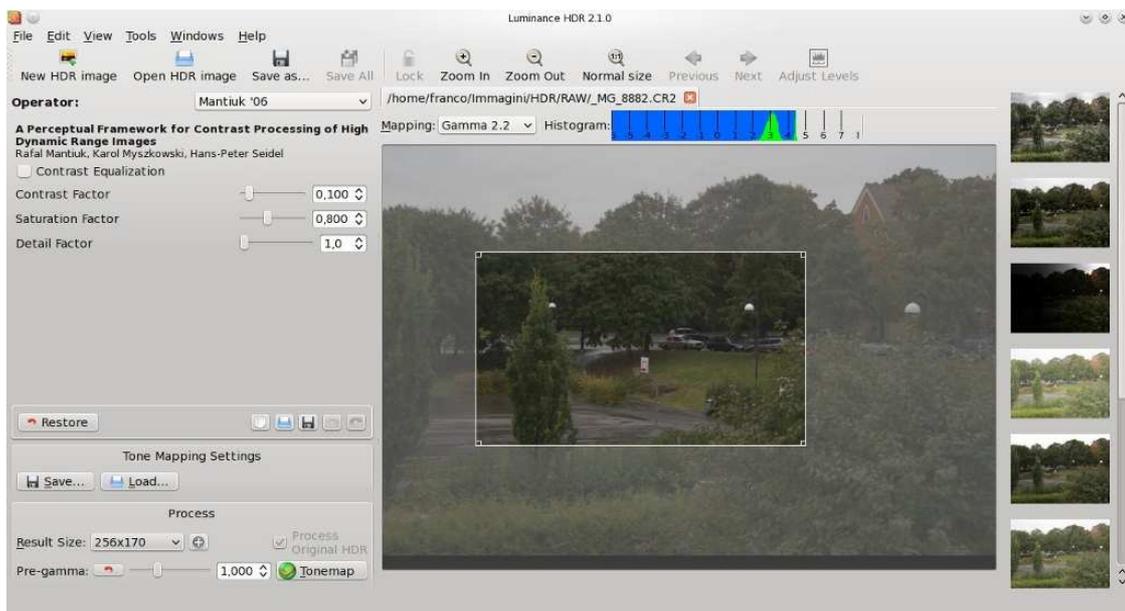
Luminance HDR can resize an HDR image to a given pixel size of percentage value counting from the original. If you use percentage, the final size in pixels will be calculated and displayed to the right from *Height* entry field.



Clicking **Scale** button will resize the HDR image.

Cropping

To crop an HDR file to some area first you need to select this area. Click somewhere on an image, drag the mouse pointer to a side and release it. You will see something like this:



You can further edit the frame you created by dragging its edges or corners. You can also move the frame around by clicking inside it and dragging mouse pointer (that will change from an arrow to a hand icon).

When the frame is placed correctly, choose *Image > Crop to selection* in menu or use the relevant button in the toolbar. Luminance HDR will create a new unsaved HDR image that contains cropped version of the original image.

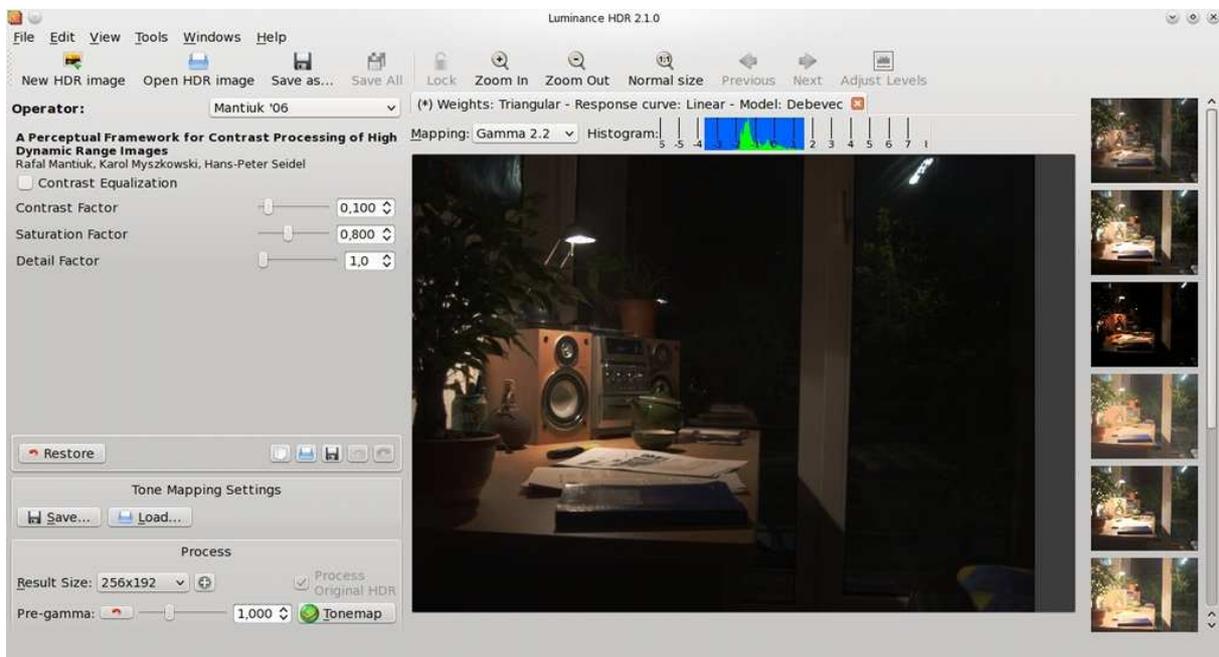
To get rid of the selection frame simply single-click anywhere outside the frame or use the *Image > Remove Selection* menu item.

Rotating

You can rotate an HDR image to 90 degrees a step, using *Image > Rotate Counter-Clockwise* and *Image > Rotate Clockwise* commands or **<** and **>** shortcuts respectively. Unlike setting an Exif orientation tag this will physically modify the HDR image.

Interactive Tone Mapping

To start tonemapping an HDR file click the "Tonemap" button in the tonemapping panel or use **Ctrl+T** shortcut.



The tone mapping step first applies to the HDR the gamma value specified by the pregamma value, then the tone mapping operator is applied.

Adjusting pregamma is an optional step, so you can safely start with picking an operator. The available operators are listed in a combobox on top of Tone Mapping Panel. Most useful operators are listed first. The last two operators, Ashikhmin and Pattanaik, have less practical use for photography, but can be of interest for robotic vision specialists.

Having chosen the operator, tweak its settings, choose resulting size and click **Tonemap** button in the bottom to tonemap an HDR image to an LDR image. Please see the "Tonemapping operators reference" chapter for details on using a particular operator.

To tonemap only a part of the HDR for a quick preview you can select a rectangle of the image using the mouse and click the **Tonemap** button. Alternatively you can click on the small preview images on the right panel. Because three of nine available operators (both Mantiuks and Fattal) are gradient domain operators that cannot provide consistent results on all result sizes by design, the preview function is not very accurate for those operators, so the selection method is preferred.

If you wish to compare several results with different settings, click the **Lock Images** button in the toolbar and use navigation widget in the lower right corner of an image to pan around. Select another image with the **Previous** and **Next** buttons and note that contents of every image window will be synced.

Sometimes the tonemapped LDR image is either too bright or too dark. You can adjust brightness using Levels dialog (**Adjust Levels** button in the toolbar) that works just like in GIMP, Krita, Photoshop or any similar raster graphics editor with the only exception that it works on luminance channel only.

If you have a lot of pictures shot in same lighting conditions or just wish to save current tonemapping settings for future use, you can save them to a text file or in a centralized database. Use "Save..." button in the Tone Mapping Panel under Tonemapping Settings for save to a file or the small "save" button above to save in the database. You then can add a comment that describe your settings. Use the "Load" buttons to reload the previously saved settings.

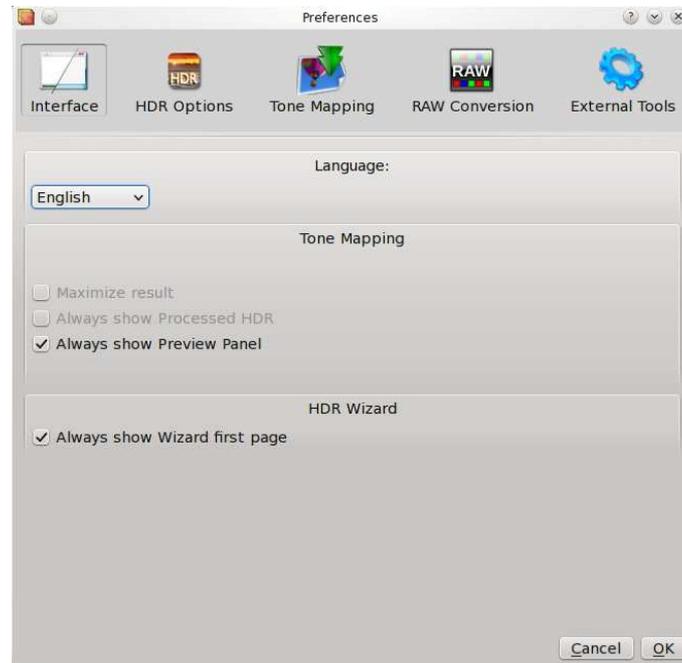
The text file contains:

- the currently selected tone mapping operator and its related settings.
- the values of the pregamma settings.

Please note that since version 1.8.4 the application (at the time named Qtpfsgui) uses a new file format which doesn't contain the postgamma setting.

Setting up

You can configure behaviour of Luminance HDR using Preferences dialog available via *Tools > Preferences...* menu item.



Interface

Here you can choose language for the user interface and help system. If no help system for your native language is available, English version will be displayed.

You can also set if you want to always show the preview panel in the main window and hide the first page of the HDR Creation Wizard

HDR TIFF File format

Here you can choose which TIFF-HDR format Luminance HDR will use to save an HDR image.

LogLuv TIFF is the appropriate format for most cases because it stores the hdr data in a reasonable amount of space (i.e. file size). Even if it is a 16 bit format, it can store floating point data.

The Float TIFF format has to be used only if you want to load your tiff in another application that cannot open the LogLuv format.

HDR Visualization

Click on the buttons to pop-up a window that lets you choose which color to use for the nan/inf and for the negative values.

Tone Mapping

Cache. Luminance HDR needs a directory to save temporary files to. By default home directory is used, but you can specify some other directory.

You can also specify two options for Batch Tone Mapping feature:

File format for LDR images. This is the file format to which tonemapped LDR images will be saved to. JPEG, PNG, PPM, PBM and BMP are your options.

Number of threads. How many threads you want to use. Each thread gets allocated to a different CPU or a core of a multicore CPU.

Tonemapping operators reference

Luminance HDR is based on pfstmo backend that implements algorithms from published scientific papers on HDR imaging. Most HDR tools are based on these papers, but Luminance is kind of different in a way that it actually uses terminology from the papers. Which means that if you are not a PhD in color science, you need a good explanation, otherwise you are lost.

This section helps understanding underlying algorithms and explains how to better use a particular tonemapping operator.

Ashikhmin



ashikhmin02 eq2 local=0.646465

When to use this algorithm

Pictures can be very high in contrast and very colourful (if all sliders are moved quite far to the left side) or very washed out. The washed out reminds on a mix of water colours and charcoal drawing.

Parameters in detail

Simple (Default not set)

It looks like that using the simple button makes it easier to achieve more realistic pictures. It seems as well that "simple" mode is quite close to "Equation 4" mode if not the simple button is set.

Drago

The Drago tone mapping operator is implementation of the "Adaptive Logarithmic Mapping For Displaying High Contrast Scenes" scientific paper written by F. Drago, K. Myszkowski, T. Annen and N. Chiba. The original purpose of the algorithm is to provide a high quality tone mapping technique to display high contrast images on devices with limited dynamic range of luminance values.

To achieve this goal the operator calculates luminance of every pixel and maximum luminance of the whole image, then divides them by average luminance and finally multiplies the result by user defined exposure factor that is called bias. The bias power function is performed to the whole image splitted to a set of 3×3 px tiles (for efficiency sake).

The only user defined option is bias. Smaller bias values produce significantly brighter pictures. Values between 0.7 and 0.9 seem to be most useful, but 0.85 is reported to be optimal according to a small survey held by developers of the algorithm. Thus 0.85 is used by default.



drago03 bias=0.85

When to use this algorithm

This algorithm is intended to imitate the human eye's response, and is useful when a true tone result is desired. It is a global spatially uniform operator: at the beginning, it calculates the average luminance of the image and, using this value and the external parameter "bias", it creates a non-linear logarithmic function that is applied to each pixel separately, without considering the neighbouring pixels. This way the algorithm is fast and it will not change the overall look of the image, but it is less adaptive and the details will be less visible compared to other algorithms. It can be used on noisy images without any particular problem because details are not amplified more than other parts of the image.

Parameters in detail

See the web page at <http://www.mpi-inf.mpg.de/resources/tmo/logmap/> for technical details.

Bias (default 0.85) +++

The Bias expands or contracts the dynamic range of the tone mapped image. It changes the base of the logarithm that is used in the tone-mapping process. If it is smaller, the dynamic range will be compressed more and the overall look of the LDR image will be brighter. After several attempts, the authors have found 0.85 to be the best value to achieve realistic and clear results.

Durand

The Durand tone mapping operator is implementation of the "Fast Bilateral Filtering for the Display of High-Dynamic-Range Images" scientific paper written by Frédo Durand and Julie Dorsey. The original purpose of the algorithm is to adjust local contrast to bring out details of high dynamic range scenes. To do so the operator decomposes every image into a base layer and a layer with fine details, then compresses contrast of the base layer.

The user defined options are:

- Spatial Kernel Sigma. The larger the values is, the larger amount of details is preserved during decomposition.
- Range Kernel Sigma. Larger values tend to result in slightly darker images with less homogenous distribution of pixels across the tonal range and much increased halo around edges.
- Base contrast. Adjustment of contrast is performed only to the base layer to preserve details. The lower value is, the higher contrast compression and the darker the resulted picture is. Thus increasing this value a bit will render a brighter picture.



durand02 spatial=8 range=0.4 base=5

When to use this algorithm

This algorithm produces the most realistic pictures. No extreme effects, but very nice output with lots of details in the picture. This algorithm needs a lot of time to calculate the output.

Parameters in detail

Base Contrast (Default 5)

Found this value usually too high and reduce it between 2 to 3. The lower the value the lower is the contrast, the picture gets darker. This can be important since it seems that light spots are too bright.

Range Kernel Sigma (Default 0.4)

Increasing this value creates a halo around lights, somewhat similar to soft focus filter.

Post tone mapping gamma adjustments (Default 1)

The higher the value here the lower the contrast will be. Seems to have a similar effect like the base contrast setting.

Fattal



fattal02 alpha=0.1 beta=0.8 saturation=1

When to use this algorithm

This algorithm is the one used by most users to produce pictures with a stunning effects. Usually these pictures are not realistic, but can be very impressive. Note that, if your images contain noise (e.g. you took them using a high ISO value), it will be heavily amplified (they are considered details), becoming much more evident. You'd better not to use this algorithm on noise images.

Parameters in detail

The algorithm works manipulating the gradient field of the luminance image by attenuating the magnitudes of large gradients. Very large changes occurring on the edges are reduced; on the other hand, smaller details are amplified thus becoming more evident. Changing the value of pre-gamma causes only minor changes to the output.

Alpha (Default 0.1)

The parameter alpha is a threshold: details whose derivative is less than alpha are amplified, those whose derivative is greater than alpha are decreased. In other words, decreasing alpha the quantity of details that are made more evident is decreased. Increasing alpha will make the details more evident, but the picture will look more unrealistic and noisy (because noise will be amplified as well). The

authors said to have used the value of 0.1 in each test with best results.

Beta (Default 0.8)

The parameter beta expresses how much the algorithm will be effective. Setting $\beta=1$, the algorithm will perform no operation on the HDRI, there will be only a linear shrink of the dynamics (i.e. you are just making a gamma correction with $\gamma=1$). Decreasing beta you will increase the effectiveness of the algorithm, in other words you will increase the compression of the dynamics making the details (and noise) much more evident. Low values of beta will create a very bright image with clear details and an unrealistic look, choosing a beta close to 1 you will get more realistic effects. The authors said to have used values between 0.8 and 0.9 with good results.

Color Saturation (Default 1)

As the name says, the color saturation. If the saturation equals 1, then you are not affecting the saturation at all: all the 3 RGB channels will be scaled with the same coefficients, so the original colors should be preserved. It can be seen like the "gamma" in the gamma-correction process, but it is applied to the three RGB components separately during scaling. That's why choosing 1 as saturation will cause no effect. If you choose saturation equal to zero, then you are telling the algorithm to discard all the information about color, it will use only the luminance of the image. The output can still have colors because of noise, but such colors might be completely different from the original. Using low values of saturation (like 0.2) will create a smooth grayscale effect.

Authors said to have used values between 0.4 and 0.6 with good results.

Mantiuk'06

The Mantiuk'06 tone mapping operator is implementation of the "A Perceptual Framework for Contrast Processing of High Dynamic Range Images" scientific paper written by Rafał Mantiuk, Karol Myszkowski and Hans-Peter Seidel. The original purpose of the framework is to provide means to adjust contrast locally without creating visible artifacts. Two ways to handle contrast are available:

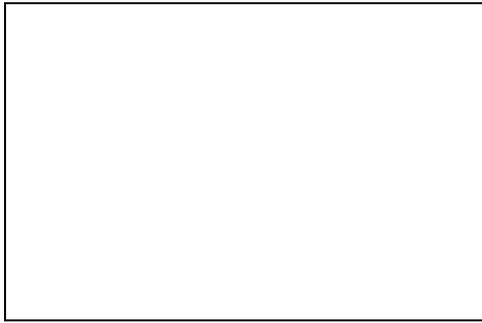
1. *Contrast mapping.* In previous sections we introduce our framework for converting images to perceptually linearized contrast response and then restoring images from the modified response. In this section we show that one potential application of this framework is to compress the dynamic range of HDR images to fit into the contrast reproduction capabilities of display devices. We call this method contrast mapping instead of tone mapping because it operates on contrast response rather than luminance. We try to fit to the dynamic range of the display so that no information is lost due to saturation of luminance values and at the same time, small contrast details, such as textures, are preserved. Within our framework such non-trivial contrast compression operation is reduced to a linear scaling in the visual response space. Since the response $R_{k,j}$ is perceptually linearized, contrast reduction can be achieved by multiplying the response values by a constant I .
2. *Contrast equalization.* Sometimes part of an image has a large contrast than other parts of the same image. By equalizing the histogram of contrast Luminance HDR allocate dynamic range for each contrast level relative to the space it occupies in an image. The resulted LDR image is usually very sharp and quite often looks unnatural.

The user defined options are:

- *Contrast Equalization.* By enabling this checkbox you will switch from contrast mapping to contrast equalization. The *Contrast Factor* option below refers to contrast mapping algorithm only and doesn't affect equalization.
- *Contrast Factor.*
- *Saturation Factor.*
- *Detail Factor.* Currently disabled.

Please note that resulted LDR image always requires gamma correction which you can do using Levels tool.

Pattanaik



pattanaik00 mult=1 cone=0.5 rod=0.5

When to use this algorithm

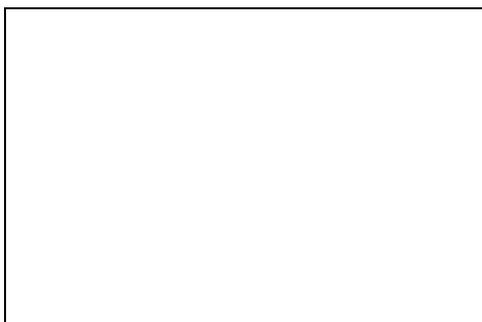
(Thanks to Ron Todd)

The key is to recognize that the ROD and CONE controls might work, to some extent, like the counterparts in our eyes: ROD controls color sensitivity while CONE controls light sensitivity.

My best results have been with high ROD value (0.75 to 0.99) and low CONE value (0.01 to 0.25). It seems that I can keep constant luminance if $ROD + CONE = 1.00$ or close but when CONE is too large the image becomes unsaturated even for large ROD values. The result is still very dependant on input HDR and on pre-gamma setting.

Parameters in detail

Reinhard '02



reinhard02 key=0.18 phi=1

When to use this algorithm

The pictures created with this algorithm will have an almost flat histogram, this means you are fully using the (limited) dynamic a standard 8-bit format can offer. The output will look realistic, keeping the details clear at any contrast.

Parameters in detail

Reinhard '04



reinhard04 brightness=-10 saturation=0.99

When to use this algorithm

It's a relatively simple algorithm with few processing on the original image. It will shrink the dynamics keeping the look very realistic. It does not do anything to stress the importance of details, to the general look will be quite blurred (you can limit this effect using a sharpening filter with Gimp or any other LDR program). It works better than other algorithms if the input image hasn't many details and it contains noise (e.g. night pictures taken with high ISO value)

Parameters in detail

The algorithm gives the possibility to change only two parameters whose meaning is clear, providing a good compromise between adaptability and easy-to-use.

The algorithm is basically this: each pixel whose value is "I" is replaced with a value given by $I/(I+S)$, where "S" is a parameter that changes in each pixel and is computed according to the surrounding pixels. Smaller "S" will cause the pixel to be brighter, simulating a gamma correction with $\gamma > 1$.

Brightness (default -10)

The meaning is obvious, the greater the brightness, the brighter the image. From a mathematical point of view, increasing the brightness will decrease the average value of "S" in the previous relationship, thus increasing the overall brightness of all the pixels.

Saturation (default 0.99)

It obviously affects saturation. If it equals zero, the output image will be more similar to a grey-scale one (even if it will still have colours); increasing its value will give more saturated colours.