

## Filmic and the curse of numbers

📅 April 30, 2020 (<https://darktable.fr/2020/04/filmique-et-la-malediction-du-nombre/>) 👤 Aurélien PIERRE  
 (<https://darktable.fr/author/aurelienpierre/>) 📁 Tutorials (<https://darktable.fr/category/tutoriels/>)

### Introduction

People think of math as the science of numbers. In fact, it is almost a coincidence that we use numbers in math. Moreover, from a certain level, we no longer even use numbers but letters. All this because math is not the science of numbers, but that **of the relations between abstract objects** .

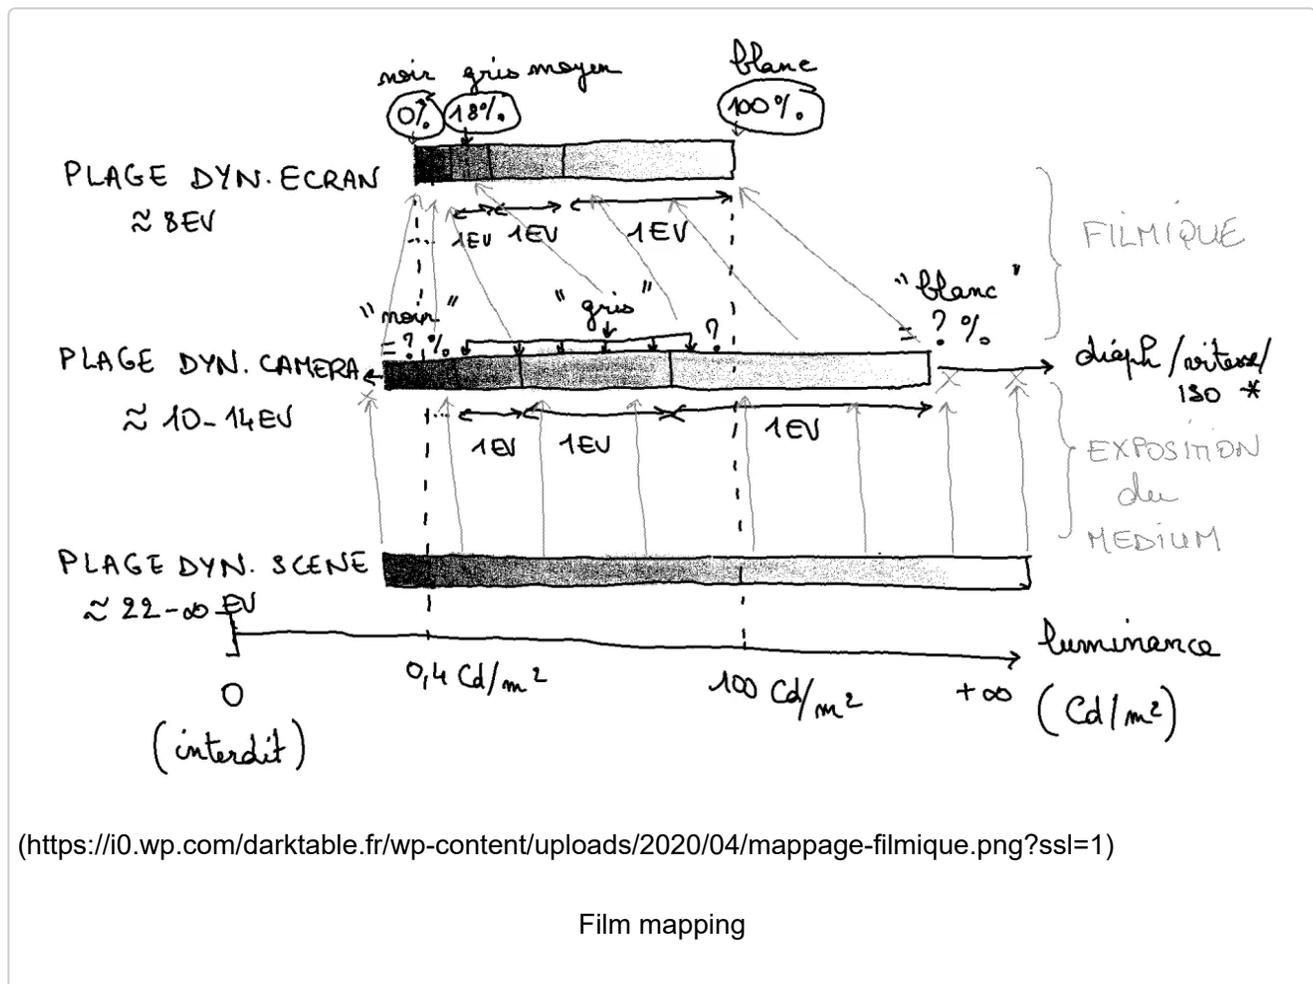
Numbers are one of the many forms these objects can take. For the part of math that interests us in filmic, that is to say analysis, we are interested in the links between variables. These links are formalized by equations. A variable is an abstraction that can contain a special value, so a number, or a whole set of numbers, in general.

We therefore seek to see how one variable varies according to another. This abstraction simplifies our work since it will allow us to consider all the possible values for these variables, in general, while freeing ourselves from individual particularities. This abstraction is often frightening, and so math seems complex, but in reality, it allows us to simplify a lot of problems that would simply be insoluble without it (and which have remained so for a long time). ^

In filmic, we try to project an input dynamic range (the maximum contrast of the image, in a way) onto an output dynamic range. The output dynamic range is known: it is related to the RGB space of the screen, where 0% encodes the blackest black available, 100% encodes the whitest white possible, and 18% encodes the gray said "way". The input dynamic range is unknown, and relatively unpredictable, which is why the user must enter it by hand. It is unpredictable because it depends on:

1. the contrast and the average luminance of the original scene,
2. the exposure settings made on the camera when the picture was taken,
3. the properties of the sensor and its conditioning electronics, and their corrections / adjustments made by the raw photo decoding library (Rawspeed),
4. settings made at the start of darktable's pipeline, because filmic comes to the end and inherits all changes made by previous modules (exposure, levels, color balance, etc.).

Filmique operates a dynamic range projection, from an unknown input to a known output, therefore. It uses a 3-point projection for this: pure black, medium gray and pure white. The medium gray allows to anchor the general exposure, the black and the white make it possible to wrap all the dynamic range accordingly around the gray, in order to compress the contrast to get into the nails while avoiding the pure and simple clipping.



In the diagram above, we see the general principle of image production, from the scene to the screen. The camera has a certain dynamic range, dictated by its physical and electronic properties. Most often, but not always, the scene has a greater dynamic range. The exposure adjustment (\*), on the body, via aperture, speed and sensitivity, slides the dynamic range window from left to right, on the luminance axis, in order to choose at which **s** value **s** of luminance we clip (both left and right). The transfer of luminance from scene to sensor is linear (shown by the vertical arrows).

The principle of film mapping (highly non-linear) is evidenced by the fact that the arrows representing the transfer are clearly non-vertical. The 3-point mapping allows you to have 2 independent compression / expansion rules: one rule for the black-gray area, and another for the gray-white area. We can therefore choose to expand or compress one or the other or both. As a general rule, we compress the gray-white area and expand the black-gray area, but nothing forces it.

Note that I have plotted the perceptual EV (or IL) increments on a linear scale, to make it clear that the first EV below white is much larger in luminance than the last EV before black. On a perceptual scale, they would all have the same width (fatally).

Since the positioning of the sensor dynamic range relative to the dynamic range of the scene is arbitrary, via the exposure triangle settings on the camera body, it is impossible to predict whether the 18% sensor encodes a medium gray on the scene, or whatever. Likewise, we have no idea if the 100% sensor is a white semi-reflective surface (therefore coincides with the white of the paper, on the print), or an incident light source (therefore an "HDR specular white"). , for example.

---

(\*) The problem with exposure is that it does not have the same effect on dynamic range, in the camera or in software.

If you consider its influence on an isolated luminance value, add a stop on the camera or multiply the value by 2 in the software, it's the same thing: you just added 1 EV or EV to this value.

Dynamic range spans a range of luminances, physical, between a certain non-zero value and another non-zero value. These values cannot be zero because light being a form of energy, zero energy is not possible on Earth, and supposes to be at absolute zero (- 273 ° C). If I multiply the entire dynamic range by one value, I shift the entire dynamic range left or right, without affecting its amplitude. 12 EV of dynamic range multiplied by 2 (+1 EV) gives a dynamic range of 12 EV shifted by one EV to the right.

But...

In darktable, one of the first modules in the pipeline is the "RAW black / white point". Its function is to manipulate the encoding of the RAW image so that the bottom of the dynamic range is set to 0, and the maximum to 1. In this way, you get really black black and really white white. Indeed, for a screen, 0% means "completely extinguishes the diodes", and for a printout, "saturates with ink thoroughly", which is not absolute black, but the denser black than the medium can render.

To set the dynamic range minimum, we subtract from the RGB values a special value, called the *dark current* ([https://en.wikipedia.org/wiki/Dark\\_current\\_\(physics\)](https://en.wikipedia.org/wiki/Dark_current_(physics))) , which corresponds to the current of the amplifier that powers the sensor. This means that your sensor never records between 0% and 100% of its

encoding range but between 3-5% and 95-100%. We will therefore subtract 3 to 5% of the signal and divide it by 92 to 100%. It's just a mathematical trick called normalization, and we have the right to apply it because RGB values are only a digital encoding, to represent a physical reality.

Without *dark current* correction, blacks would be systematically milky. The value of the *dark current* is measured for each image by the software, on a part of the image that you never see because it is cropped, and which is never illuminated by the rays formed by the objectives. At all times, the pixels in this region are in the dark and therefore have exactly the value of the *dark current*.

But...

Zero multiplied by anything equal to zero. Once the dynamic range is set to zero on the left, exposure compensation no longer translates it along the luminance axis, but **widens it**. 12 EV of dynamic range multiplied by 2 (+1 EV) gives a dynamic range of 13 EV, increased by one EV to the right, but which has not moved to the left. Remember ... mathematics does not care about numbers, it is only concerned with the links between variables, and these variables exist in an interval that must be considered as a whole.

Whether the dynamic range is widened is not a problem, in itself, since the widening is uniform. On the other hand, this zero on the left will pose a problem for filmic, which uses a logarithmic mapping, because the logarithm of zero is not defined. Filmic therefore clips arbitrarily at - 16 EV below 100%, ie at 0.0015%. But it will also seriously bother us to **predict the optimal filmic settings** because the dynamic range which arrives at the input is no longer necessarily that of the sensor, measured and tabulated by DxO Mark (<https://www.dxomark.com/category/camera-reviews/>) or Photons to Photos (<https://photonstophotos.net/Charts/PDR.htm>).

This therefore means that none of the parameters of the scene can be predicted a priori, in filmic. It depends on the scene, the sensor, and what you did in the upstream pipeline.

---

On this, I find my users to guess on the ideal value of the medium gray, like the Greeks who are ecstatic in front of pi or the golden ratio. Then to derive standard values of white and black, as if it was enough to read graphs of dynamic range of the sensors to make presets with the small amounts onions.

I'll say it again, math doesn't care about particular numbers. They appear anecdotally, but they are not the ones that are interesting. We are interested in **variables** and their **connections**.

Mid gray defines the value of the dynamic range pivot. And this value, we do not care. What matters is what it stands for: the perceptual center of the dynamic range.

If I let you paint a gradation of gray going from black to white, continuous and even, giving you only black paint and white paint, and letting you scale the proportions to adjust the tonal transition from black to white of visually homogeneous way, I can already tell you that the proportion will be roughly 80% black and 20% white in the center of the gradient. The reason is linked to different psychophysical effects, which we know quite well without really understanding their origin. So, for a reflective medium (roughly, paper), the average gray value is between 18 and 20% of the pure white luminance of said medium. This is where the standard medium gray value comes from.

But 20% white, how much light is that? It depends on the illumination the sheet receives. Let's imagine that it receives 100 Cd / m<sup>2</sup>. Your average gray is 20 Cd / m<sup>2</sup>.

Now imagine that I add a second sheet to the scene, with the same gray gradient, but lit by another light source, at 400 Cd / m<sup>2</sup>. Your average gray is 80 Cd / m<sup>2</sup>.

Take a photo of this scene: you have two light sources, two white values and two gray values, which is the real value to choose for filmic?

Scientific response: none.

Pragmatic answer: we dabble the noodle with a babouche.

Artistic response: any value that best highlights your subject, that creates the mood you like ...

It is nothing more than the pivot of the transformation, a mathematical artifice, its value contains no truth, it is what it represents that counts: the perceptual center of the dynamic range, which will serve as an anchor to mathematical transformation.

Any image contains at least 2 sources of light: incident light (which comes directly from a primary source: sun, bulb), and reflected light (which comes from a source after bouncing off a surface: ceiling, floor, etc. walls, etc.). Any image therefore contains an infinity of candidates for medium gray, and it is a matter of choice. Just arrange to snap the overall white of the scene to the level of the brightest white in the frame to avoid losing highlight detail.

Yes, indeed, a natural image with moderate contrast will have an average brightness very close to this famous 18%, but this is a statistical and descriptive reality, not a rule to necessarily follow. And in the event of high-key or exposed backlighting for highlights, this description becomes false and you should absolutely not try to stick to it.

Ditto for black. Adjust it to get as much low light as needed without falling into the milky effect, but don't waste time calculating the dynamic range from DxOMark, as there is no guarantee that this is the effective dynamic range at the input of filmic.

---

The first operation carried out in filmic is  $\log_2$  (RGB / gray), i.e. a logarithmic scaling **designed / manufactured / tampered with** so that:

- pixels equal to this gray value are mapped to 0 EV (because  $\log(1) = 0$ ),
- darker pixels are mapped to negative EV values,
- brighter pixels are mapped to positive EV values.

Why ? Because the slope of the logarithm is steeper between 0 and 1 than between 1 and infinity, which makes it possible to dilate the shadows much stronger than the highlights.

If you choose an 18% mid-gray, then divide by 0.18, which is equivalent to multiplying by 5.55, which is equivalent to applying an exposure compensation of +2.47 EV. If you choose a 9% mid-gray, you divide by 0.09, which is the same as multiplying by 11.11, which is equivalent to applying an exposure compensation of +3.47 EV. So raising the exposure in the exposure module or lowering the gray of the scene in filmic is exactly the same, in the context of this logarithmic scaling. ^

But after the logarithmic scaling, comes the S curve, which is a spline that I custom designed to behave like the densitometric curves on silver film. The first constraint of this curve is to map the gray of the scene to the gray of the screen, regardless of their values. If the gray of the scene used by filmic is equal to the gray of the screen, that is, if you have pre-adjusted the exposure in the exposure module to set what you consider to be the midtones around 18% , the medium gray crosses straight through the film mapping, unaffected, which makes the contrast more easily controllable (ergonomically speaking) by side effect, because the medium gray is then relatively centered on the graph.

If the gray of the scene used by filmic is far from 18%, the curve is off-center on the graph, unbalanced to the right. The available contrast is reduced, the curve is less flexible to control because the dynamic range is less symmetrical, and the blacks are less dense.

The two sets of settings below give the same results for midtones and highlights, but the 18% gray variant results in somewhat denser blacks, because the base of the curve is a little lower (and this, even reducing the latitude to avoid clipping at the bottom). At 18%, the dynamic range is almost symmetrical (white  $\simeq$  - black), which is desirable for purely ergonomic reasons for controlling the curve (unrelated to the actual image processing).





(<https://i0.wp.com/darktable.fr/wp-content/uploads/2020/04/filmique-18.png?ssl=1>)



(<https://i2.wp.com/darktable.fr/wp-content/uploads/2020/04/filmic-9.png?ssl=1>)

Another argument for adjusting exposure as early as possible in the pipeline is the input color profile. The matrices used as the input color profile are not perfectly accurate over the full dynamic range, and the  $3 \times 3$  matrix profiling method is itself relatively crude. The matrices are therefore optimized to give better precision, primarily for the mid tones around the skin tones (Caucasian, etc.). The exposure module, in darktable, is applied by default before the input profile and therefore allows the midtones to be pushed into the comfort zone of the matrix, benefiting from a little extra precision when applying the profile.

So why define the gray of the scene in filmic, by default at 9%? **Out of weakness**. Tired of being harassed by the High Priests of the Holy Church of the intuitive UX, who preach for an image conforming to the JPEG box in 2 clicks®, I cooked, to silence them, a filmic default setting that sets the gray to a value

close to what most JPEGs from several manufacturers do, and which gives a poorly viewable result as soon as the module is activated. The !

But I will not be taken back, and it will emerge in filmic v4, which is already finished and in the testing phase (<https://github.com/darktable-org/darktable/pull/4800>) . From now on, if intuitive cannot be reconciled with rational, that won't be my problem anymore, because nothing, but really nothing, in image processing is intuitive anyway. Doing a touch up in a minimum of operations is a laudable workflow goal, as long as it cares about getting the right operation in the right place in the pipeline. When this is not possible, it is the consistency of the colors and the pipeline that must prevail.

---

## Conclusion

Take it from someone who does math all day: drop the calculator, and touch up with your eyes. Math is my problem, not yours. Just understand what the setting refers to **in real life** , what it means, but don't get stuck on numbers that will spend their time changing from one image to another.

In the course of the image processing chain, a dizzying amount of non-intuitive operations take place that you cannot even imagine. I try to explain the ins and outs in detail, even if it gives more ground to grind to those who think that you have to be an engineer to use darktable ("The proof, his videos are full of maths and last 2 hrs - I'm a photographer, damn it, not a full-time geek ").

No, you don't have to be an engineer to **use it** , but you need a certain scientific background to **understand it** . And understanding it will help you use it in problem-solving, without depending on a guru to tell you what to do and give you ready-made recipes. But there's nothing stopping you from continuing to follow recipes if it works for you. Just, in this case, make sure your guru understands what he is causing, which on the Internet is not the norm.

Filmique was designed to be robust, and was intended to be robust to be able to easily create presets for a given camera. You make your home presets, depending on the sensitivity and the device used, that's fine. But make sure to base these presets on correct assumptions, and not on extrapolations of DxO Mark values, which in addition to being probably inflated, do not represent what filmic sees on entry.

The right way to adjust filmic, in the sense that it is the one that will cause the least problems, is to leave the grays (scene and screen) at 18.45% and forget them. Adjust the overall exposure until the image is readable, and its subject is bright enough. Adjust the white and black levels of the scene, in filmic, until the shadows are uncapped, and the highlights not clipped. Remember that these levels define the dynamic range clipping bounds, that is, the threshold values, not the amount of correction applied. As long as the highlights are clipped, the white level is too low. As long as the blacks are clogged, the black level is too high. If blacks are milky, the black levels are too low.

You have the right to ignore my advice as long as you are well aware that you are going to pay for it at one point or another. Playing with medium gray is to complicate life in less pleasant cases, by degrading the ergonomics of the control of the contrast curve. Which doesn't mean it doesn't work, but it adds unnecessary difficulty. ^

**You don't have to avoid clipping at all costs** . Sometimes clipping highlights is the best compromise to maintain correct overall contrast, especially if the highlights have already been clipped by the sensor. Ditto, clipping the low lights makes it possible not to over-amplify the noise. What matters in a photo is that the subject of the photo is readable. For the rest, we work it out. The whole secret is in the smoothness of the valid / clipped transition.

Mais, de grâce, ne cherchez aucune vérité dans des numéros. Vous perdez votre temps, et vous embrouillez vos disciples. Et c'est une faute épistémologique, scientifique, artistique et pratique. Une variable se décrit par sa définition et par sa relation à la réalité, pas par sa valeur. Autrement, c'est de la numérologie que vous faites.

#### Partager :

 Twitter (<https://darktable.fr/2020/04/filmique-et-la-malediction-du-nombre/?share=twitter&nb=1>)

 Facebook (<https://darktable.fr/2020/04/filmique-et-la-malediction-du-nombre/?share=facebook&nb=1>)

 Telegram (<https://darktable.fr/2020/04/filmique-et-la-malediction-du-nombre/?share=telegram&nb=1>)

Diaspora (<https://darktable.fr/2020/04/filmique-et-la-malediction-du-nombre/?share=custom-1489004640&nb=1>)

 Plus

#### Articles similaires



(<https://darktable.fr/2019/09/comment-utiliser-filmique-rgb-et-le-nouvel-egaliseur-de-tons-dans-le-futur-darktable-3-0/>)  
Comment utiliser filmique RGB et le nouvel égaliseur de tons dans le futur darktable 3.0 ?  
(<https://darktable.fr/2019/09/comment-utiliser-filmique-rgb-et-le-nouvel-egaliseur-de-tons-dans-le-futur-darktable-3-0/>)



(<https://darktable.fr/2020/01/darktable-3-rgb-ou-lab-quels-modules-au-secours/>)  
darktable 3 : RGB ou Lab ? Quels modules ? Au secours !  
(<https://darktable.fr/2020/01/darktable-3-rgb-ou-lab-quels-modules-au-secours/>)

 Faut-il étalonner son écran photo ? (Et autres problèmes colorimétriques)

(<https://darktable.fr/2019/12/faut-il-etalonner-son-ecran-photo-et-autres-problemes-colorimetriques/>)  
Faut-il étalonner son écran photo ? (Et autres problèmes colorimétriques)  
(<https://darktable.fr/2019/12/faut-il-etalonner-son-ecran-photo-et-autres-problemes-colorimetriques/>)

**AURÉLIEN PIERRE ([HTTPS://DARKTABLE.FR/AUTHOR/AURELIENPIERRE/](https://darktable.fr/author/aurelienpierre/))**