

The ISO Definition of the Dynamic Range of a Digital Still Camera

Douglas A. Kerr

Issue 2
February 6, 2008

ABSTRACT

The dynamic range of a digital camera can be simplistically defined as the ratio of the maximum and minimum luminance that a camera can “capture” in a single exposure. But when we try to quantify this property, we find that the establishment of an explicit definition is much more complicated than it seems on the surface. International Standard ISO 15739-2003 gives an explicit definition of dynamic range for a digital still camera and a procedure for determining it. This article explains the basic concept of dynamic range and discusses some of the complications in defining it. Then, the definition given by ISO 15739-2003 is discussed in detail.

BACKGROUND

Dynamic range

The dynamic range of a digital camera is often (very simplistically) defined as the ratio between the maximum and minimum luminance which, in a single image, can be successfully “captured”. Clearly, since many scenes exhibit a large range of luminance, having a sufficiently-large dynamic range is desirable in attempting to completely and accurately record such scenes.

What do we really seek in our quest for “adequate” dynamic range? In general, from a standpoint of photographic technique, what we are really trying to assure is that, within the same “shot”:

- Detail carried by small variations in luminance about the highest local average luminance (that is, the “highlight detail”) is captured by the camera, and, simultaneously
- Detail carried by small variations in luminance about the lowest local average luminance (that is, the “shadow detail”) is captured by the camera

Insufficient dynamic range for the luminance range in the scene will mean that we cannot simultaneously achieve both of those objectives (again, in a single image). (We can generally achieve either one by prudent choice of exposure.)

Photometric exposure and scene luminance

The physical property to which photographic film or the elements of a digital sensor responds is photometric exposure (symbol: H). At a particular spot on the image, this is the product of the illuminance on the film or sensor and the time it persists (*i.e.*, the exposure time).

However, our basic concern is with variations in scene luminance. Because in the discussion of dynamic range we are ordinarily making comparisons with different parts of the same image, the aperture used, related factors such as lens transmission and bellows factor, and the exposure time—the factors that control the relationship between scene luminance and focal plane photometric exposure—are fixed.

Thus, the ratio between (for example) two values of photometric exposure in an image will be the same as the ratio between the corresponding two luminance values in the scene. So for convenience, when discussing the response of a digital camera sensor (since again it is only ratios that are of importance), we can speak of the effects of a certain luminance rather than of a certain photometric exposure.

Explicit definitions of dynamic range

Often, discussions of dynamic range do not revolve around a well-formulated definition of exactly what ratio is meant. Many times, the definition that is implied (if not enunciated) has to do with the effect of the “quantization” of image luminance at various stages in the processing of the image, starting with the digitization of the analog voltage by which an individual photodetector reports the “photometric exposure” (the illuminance-time product) to which it has been subjected. Later impacts of quantization occur after we have deduced the specific color of each pixel in the image (*color* here of course being a property that embraces both luminance and chromaticity) from the suite of photodetector data.

For example, if at the stage of the process where we examine the digital image, we find that the largest luminance of a scene spot that can be digitally represented has a relative luminance of 4095 units on the applicable digital scale, and the smallest luminance that receives a non-zero digital representation has a relative luminance of 1 unit, we may be tempted to conclude that the dynamic range of the system (**as seen at that point in the image processing chain**, an important factor) is 4095 (sometimes stated as 4095:1).

But this doesn't follow our notion of the “base” luminances for which detail, carried by luminance variations about the base value, can be perceived.

So perhaps we need to consider the “maximum” base luminance (the numerator of our ratio) to be 4094 (allowing detail to be recorded that varies from 4093 units to 4095), and the “minimum” base luminance (the denominator of our ratio) to be 2 units (allowing detail to be recorded that varies from 1 unit to 3). Now, our numeric dynamic range would seem to be 2046.5:1. So we see that this little piece of “hair-splitting” cuts our assessment of dynamic range about in half—an apparent degradation of “one stop” to the photographer.

So perhaps this whole approach to defining dynamic range is ill advised, since the result we get varies so much with how we decide to split the hairs.

The role of noise

Another matter we must consider before settling on a definition of dynamic range is that of noise, the random variation in the reported luminance of pixels compared to their actual luminance values. We recognize that in fact the image, for areas below a certain base luminance, may be so “corrupted” by noise that we cannot honestly say that detail carried by small variations of luminance about that base luminance is “captured”, at least not in a way that is usable.

Thus, we may eventually wish to think in terms of the dynamic range for the camera as being defined as the ratio of the highest luminance at which small luminance differences are recorded to the lowest luminance at which details are recorded with less than some arbitrarily-established level of noise.

In fact, it is a specific definition along this line that we will discuss in detail here, as soon as some further philosophical preliminaries are taken care of.

Where do we look?

We often speak of the way in which luminance values are “captured”, when it is “captured and delivered” that is really meant. Thus, in discussing the matter of the dynamic range of a camera, we must be clear as to which form of the image we will be examining as we seek to determine the dynamic range of the camera in a laboratory or field test.

Looking at the raw photodetector data

Sophisticated photographers may be most interested in the dynamic range as manifested in the raw¹ camera output. This output is just a verbatim transcription of the digitized values of the outputs of all the individual photodetectors across the camera frame.

In most digital cameras, there is one photodetector for each pixel location in the image. These are not “tristimulus” detectors, capable of reporting the color of the light they receive from some tiny patch on the image. (Note again that “color” here means the property that embraces both luminance and chromaticity.)

Rather, these are “monochromatic” photodetectors that have been given a response centering on certain wavelengths of the visible spectrum by placing “color filters” in front of them, generally of three different kinds, applied in a repeated pattern across the entire frame (an arrangement often called a “color filter array”, or CFA).

Such an array obviously cannot tell us the color of the light falling at any particular pixel location (either in luminance or chromaticity). But we can deduce the color

¹ This is often called the “RAW” output, although there is no reason for that designation, the word not being an acronym but merely a metaphor for “unprocessed”.

each pixel “probably has” by analyzing the “cloud” of photodetectors in that neighborhood, a process called “demosaicing”.²

When we ask the camera to deliver to us an actual (“developed”) image, this demosaicing is done inside the camera. But for much serious work, the photographer takes the raw output, delivered by the camera in a special file format, and applies the demosaicing later in “raw conversion” software. By doing so, the photographer can control various parameters of the conversion in order to produce the “best image”, especially in cases where the exposure utilized was not ideal, where there are severe problems with the chromaticity of the incident light, and so forth.

One result of the CFA arrangement as usually implemented is that, for any given setting of the “ISO sensitivity” of the camera, the maximum recordable luminance will vary between the photodetectors of the different color groups.

This situation gives some complications in deciding how to assess the dynamic range of the camera from this perspective. Do we consider only the response of the camera to “white” light, perhaps considering the maximum luminance to be that which just pushes the “most sensitive” color group to its limit, and then consider the luminance at which that same group meets our criterion (noise-based or otherwise) as the “denominator” of the ratio?

I’m going to abandon at this point further pursuit of this line of thought, since the dynamic range definition we will really be discussing here is implicitly predicated on working with a “developed image”. This is in no way meant to trivialize the importance of a good concept of dynamic range for the context of a photographer operating with the raw data output of a camera. This is just a very complex issue that is a proper topic for a separate essay, and which is not treated at all in the ISO standard that is the excuse for this article.

Looking at the developed image

Of course, often the use of digital cameras is with an actual image being delivered by the camera, often in the form of a JPEG file. It is interesting to follow the path of the image on its way to that file.

1. Each photodetector develops an (analog) voltage that is (ideally) proportional to luminance (assuming a certain shutter speed, aperture, and so forth), but which may contain a random (noise) component.
2. The analog voltages from the individual photodetectors are amplified and then digitized. The magnitude of the amplification can usually be changed by the photographer to control the “ISO sensitivity” of the camera. (Because of quantizing, from this digital form of the image we could never precisely recover the photodetector data as it was in step 1.)

² We sometimes speak of this as “developing” the image, basking in the nostalgia of film photography.

3. The raw photodetector data is demosaiced, resulting into a basic “linear **rgb**” representation of an actual image (and recall that the color values of the pixels have only been “deduced” from the suite of raw photodetector data). The output representation is digital, and it of course differs from the “ideal” result of the process because of quantizing.
4. This image is subjected to various “image processing” actions, such as sharpening, compression/expansion and/or shifting of the “tonal scale” (for brightness or contrast adjustment), and so forth. (Sometimes some of this is actually done as part of step 3.) The result differs from the “ideal” result of the process because of quantizing.
5. That **rgb** representation is converted to an RGB representation, typically in the sRGB color space. Among other things, this involves transforming the **rgb** values into a non-linear form, a process that is often called *gamma precompensation*. (Because of quantizing, from this form of the image we could never precisely recover the image as it was after step 4.)
6. The RGB representation is transformed into a YCbCr representation. (Because of quantizing, from this form of the image we could never precisely recover the image as it was after step 5.)
7. The suite of YCbCr data is “compressed” by a JPEG algorithm to reduce its total bit size. This is a non-reversible transformation (often said to be “lossy”), meaning that we can never again, downstream, recover exactly the image as it was after step 6 (to a more serious extent than from the situations mentioned at the ends of several of the previous steps).
8. The suite of JPEG data, along with other information, is recorded in the image file.

When we decode the file, we get an image in the same form as described in step 5, above, but not precisely the same image (because of both quantizing impact and the implications of the non-reversible JPEG compression).

Note that these things that happen to the data along this chain can impact the measured dynamic range (for any conceptual definition). But, if this is the output the user receives, then it is the assessment of dynamic range based on examination of this image that really describes the performance of the camera for such a user. And thus the ISO standard for determining the dynamic range of a digital camera is predicated on the examination of the delivered JPEG image data.

Going back to the raw data, if we “develop” that into an actual image outside the camera with raw conversion software, and intercept the image after step 4 above,³

³ Note that most “raw conversion” software packages allow us to capture the image at this point into an RGB TIFF file, which means that we can examine it without the further corruption introduced at stages 4 and 5 of our chain.

we can use that output in the determination of dynamic range, getting a value that is perhaps the most meaningful single-number answer for the camera's sensor system itself. This sidesteps the conundrum of defining dynamic range for a CFA array from the photodetector raw data. But this simple "one number answer" doesn't tell the really sophisticated user everything that is needed to be able to predict what the camera can do when the raw data is "custom processed". (No single-number answer could!)

A STANDARD DEFINITION AND MEASUREMENT PROCEDURE

International standard ISO 15739-2003, which covers the definition and measurement of the noise performance of digital still cameras, also provides a definition for dynamic range and an associated measurement procedure, based on the noise outlook we discussed above.

Unfortunately, as is so often the case in standards relating to photography, the standard is a bit careless (often paradoxical) in some of its terminology, notation, and discussion. Accordingly, it can be difficult to understand exactly how the definition, and the measurement procedure, work.

In this article, we will discuss the ISO 15739 dynamic range definition and measurement procedures, hoping to demystify these topics.

THE ISO 15739 DEFINITION OF DYNAMIC RANGE

The concept of the definition

The concept behind the dynamic range definition given by ISO 15739 is based on the ratio of the maximum luminance that receives a unique coded representation (the "saturation" luminance) to the lowest luminance for which the signal to noise ratio (SNR) is at least 1.0. This is based on the very arbitrary assumption that detail recorded with a SNR of 1.0 or above is useful and that recorded with an SNR less than 1.0 is not.

About noise

Before we proceed with the definition, let's talk a little more about noise.

The term *noise* here is borrowed from usage in electrical signal technology. In the context of camera testing, it refers to random variations in the digital representation of the luminance of pixels that are given equal and unchanging photometric exposure.

This noise appears in two forms:

- *Spatial noise* is the variation in digital output among different pixels in the image that have been given the same photometric exposure (that is, for a given shutter speed and aperture, are the images of parts of a uniform-luminance object). The relationship between the (inconsistent) outputs for different pixels, from this phenomenon, does not change between successive images taken of the same

object under the same exposure conditions. Not surprisingly, this noise component is often called “fixed-pattern” noise.

- *Temporal* noise is the variation in the digital output from a given pixel in successive images taken of the same object under the same exposure conditions.

Both types of noise serve to degrade the image. Only temporal noise is taken into account in this noise-based definition of dynamic range.

Noise on a luminance basis

The noise which is spoken of is conceptually defined on the basis of the luminance equivalent of the noise as we observe it in the digital image. That is, if we find that there is a certain random (noise) component in the digital code for a pixel in the digital image, we then determine what luminance variation would produce the same variation in the digital code, and consider this number as describing the “noise” for use in the reckoning of “signal-to-noise ratio”. “Signal”, in this case, refers to the actual luminance of object itself.

The measure of noise

This noise is quantified in terms the standard deviation (sigma) of the luminance implication of the randomly-varying digital code, for a consistent actual object luminance, over the collection of all pixels in the test region over a number of test images, analyzed in such a way that only the temporal aspect is finally considered.⁴ In effect, the number we get is the average, over all the pixels in a test “block”, of the temporal noise exhibited by the individual pixels.

Borrowing notation most commonly used in electrical engineering, we may speak of the standard deviation as the “root mean square (RMS) deviation of the implied luminance about its mean value. There are many reasons for the use of this measure of the variation. One again goes back to the electrical engineering parallel of this topic. There, the power in an electrical signal (such as the noise component of a noisy signal) is proportional to the RMS value of the instantaneous voltage of that signal.

Since human perception of the “potency” of a sound basically follows its power content, we can see why the RMS value of a noise components is a useful measure of the noise. Various rationales allow us to extend this concept to luminance noise in an image. Thus, we use the RMS variation (the standard deviation) of the luminance equivalent of the digital output values as our numeric measure of the amount of noise.

⁴ This is wholly parallel to the practice, in electrical engineering, of using the RMS (root-mean-square) measure of the discrepancy between a voltage signal we have and the ideal (noise-free) signal as a measure of the noise on that signal. In fact here we often speak of the “RMS noise”.

Backing our the nonlinearity

Because the scale of digital codes for pixel luminance normally has a non-linear relationship to luminance proper (usually because of the application of “gamma precompensation”), in order to determine the luminance equivalent of the observed digital noise we must know the slope of the curve of digital code vs. luminance at the luminance for which the noise is determined. In the ISO standard, this slope is called the “incremental gain” of the image coding system at that luminance. The reason that the slope is the relevant factor is that we are starting with a **variation** in digital codes (resulting from noise) and we want to determine what **variation** in luminance would cause that same variation in digital codes.

Conceptually, if we have determined the noise in terms of the digital codes in the image, we can divide that value by the incremental gain for the luminance at which we are working and get the “luminance-basis” noise.

Note that when stating incremental gain numerically, we must be careful about the units of both numerator and denominator.

The matter of incremental gain is discussed further in Appendix A.

The problem with the conceptual definition

We said that the conceptual premise of the definition of dynamic range in the ISO standard is the ratio of the greatest luminance that receives a unique digital representation to the luminance at which the luminance-based SNR is 1.0

But there is problem in actually following that concept. In statistical terms, a situation in which the signal-to-noise ratio is actually 1 corresponds to a mathematical random variable whose *standard deviation* (the measure of the noise) is equal to its *mean* (the measure of the underlying “signal”). If the random variation followed the “normal distribution”—and it won’t necessarily for our camera case, but this is a good illustrative example—then for about 16% of the values, the actual value of the variable would be negative. Of course, in our case, the variable represents luminance, and ordinarily there cannot be a negative value of luminance for any pixel at any time. (It would have no physical meaning, and more to the point, the usual digital coding scheme cannot represent it.)

Thus, in an actual situation with a signal-to-noise ratio of 1, the observed digital noise would have been “truncated” and thus would give a misleadingly-small luminance noise value.

To evade this problem, we proceed as follows:

1. We measure the noise at a certain arbitrary low luminance (one that is still high enough that the “truncation” problem mentioned just above would be negligible, since only a very tiny fraction of the occurrences would now have negative values). This is called the “reference black” luminance. (In the ISO standard, this luminance is in fact 1/100 of the maximum recordable luminance.)

2. We then say that the luminance which is numerically the same as that value of noise is the one that “would have” a signal-to-noise ratio of 1.0.
3. The ratio of the maximum recordable luminance (the “saturation” luminance) to the luminance reckoned in (2) above is reported as the dynamic range of the camera.

THE GORY DETAILS

The actual details of the way in which this definition is represented, and the equations and supporting definitions that formalize the definition, are given in Appendix A.

Appendix B discusses some of the paradoxes, conundrums, and inconsistencies of terminology found in the standard.

#

APPENDIX A

THE GORY DETAILS

INTRODUCTION

In this appendix, we examine standard ISO 15739-2003 in more detail.

As with so many standards in the area of photographic photometry, this document is an absolute horror. Many things are defined in ways which, while rigorously correct, seem almost intended to prevent any possibility that the reader will understand what is going on. Names and symbols for various quantities are assigned in perverse and grossly misleading ways. Its occasional attempts to give tutorial insight fall short of the mark.

Hopefully, the subtitle of this appendix can be, "I let my head hurt so yours won't have to (quite so much)".

THE FORMAL DEFINITION

ISO 15739 describes the basic calculation for determining dynamic range this way (my notes in square brackets):

. . . the camera dynamic range is obtained by measuring the camera signal-to-temporal-noise ratio using a 2.0 density "black reference" [*i.e.*, at a luminance of 1/100 of the maximum recordable luminance]. The measured signal-to-black temporal noise ratio is then multiplied by the ratio of the camera maximum level to the camera level from the "black reference". The result is reported as the ISO DSC [*digital still camera*] dynamic range.

You got that?

What this means is, conceptually:

The noise is measured on a digital basis at a luminance of 1% of the maximum recordable luminance. This is, in effect, converted to a luminance basis. The maximum recordable luminance is divided by that noise value to give the dynamic range of the camera.

Recall that this is, in effect, determining the ratio of the maximum recordable luminance to the luminance for which the signal-to-noise ratio **would be** 1, done in a peculiar way since we actually cannot have a situation where the signal-to-noise ratio is 1 (since it would require some fraction of the pixel values to imply a negative luminance, which is physically impossible).

THE REFERENCE BLACK LEVEL AND SCALES OF MEASUREMENT

The luminance at which the noise is measured for the determination of dynamic range (the "reference black" luminance) is 1% of the maximum recordable

luminance (which I call L_{clip} ⁵). From the perspective of a widely-held model of exposure control via reflected light exposure metering (implicitly embraced by ISO 15739), this would be the recorded relative luminance given by an actual photographic subject with a reflectance of 1.4% in an “ideal” metering scenario (one in which the average luminance of the scene was 18%).

Under this model, an actual object with a reflectance of 100% would, for this ideal metering scenario, be recorded with a relative luminance of 71% (100/140) of L_{clip} (providing the well-known so-called “half stop margin” against overexposure⁶).

As a result of this outlook, the scale used for luminance runs from 0 to 100 units at 71% of L_{clip} , and thus to 140 units at L_{clip} .

The scale of measurement used for the noise as observed in the digital domain just follows the numerical values of the digital codes; a difference of one “count” in the digital representation is considered to be one unit on the scale of digital noise. If we are examining the image on the basis of an RGB representation (perhaps sRGB), this is a 0-255 unit scale.

INCREMENTAL GAIN

Typically, the scale of the digital representation of image luminance used in image files is related by a nonlinear function to the actual luminance represented, an arrangement often spoken of as “gamma precompensation”.

Although we will note the noise in terms of the variability of the digital representations of a constant luminance, we must actually have the noise on the basis of luminance (that is, we must be able to describe the noise as if it came from an actual random variation of luminance) if we are to be able to compare it to the “base” luminance (the “signal”).

To allow this, we use a factor called the *incremental gain* of the gamma precompensation process. This is just the slope of the curve of the digital representation vs. luminance. Since the function is nonlinear, this slope will vary across the range for the variables, so we must know the incremental gain at the particular point on the curve where we are working.⁷

⁵ It would of course make sense to call this L_{sat} , but the standard inexplicably uses this symbol for the luminance at the “margin line” (at 70.1% of the maximum recordable luminance), a concept discussed in the next paragraph.

⁶ It is this margin provision that causes the confusion between the roles of an 18% average reflectance and a 13% average scene reflectance as premises for the “standard” exposure meter calibration; the luminance that is 18% of “70.1% of maximum recordable luminance” is 13% of the maximum recordable luminance itself.

⁷ Actually, since the digital scale has only discrete values, the curve has a staircase nature, comprising only infinite and zero slopes. What is actually done is that a smooth curve is fit to the staircase curve, and its slope at the luminance of interest is considered the incremental gain.

In terms of (the) differential calculus, this incremental gain is defined as:

$$G_i(L) = \frac{dD}{dL} \quad (1)$$

where $G_i(L)$ (read " G_i of L ") is the incremental gain at some luminance, L , and dD/dL is the first derivative of D (the digital representation of luminance) as a function of L (the luminance being represented), at the luminance of interest.

In this, the unit of D is one step in the progression of digital code values, while the unit of L is on the scale that runs from 0 to 100 at the "margin" level, or to 140 at saturation.

THE DYNAMIC RANGE EQUATION

The standard then gives this equation for the definition; I have recast it in slightly tidier notation and simplified it:

$$D = \frac{1.4 L_{sat} G_i}{N_D} \quad (2)$$

where

D is the dynamic range

L_{sat} is the luminance at the "margin level", 70.1% of the maximum recordable (saturation) luminance (yes, L_{sat} is a peculiar symbol for that)

G_i is the incremental gain at the reference black luminance

N_D is the measured digital noise at the reference black luminance

(In the equation as stated in the standard, the constant 1.4 is actually made up of the constants 0.014 and 100, whose individual significances are stated in a very confusing way. You'll see in the next paragraph what the final significance of 1.4 is.)

But L_{sat} is 100/140 (1/1.4) of the actual saturation luminance (which we call here L_{clip}). Thus we can recast equation 2 thus:

$$D = \frac{L_{clip} G_i}{N_D} \quad (3)$$

DERIVING THE EQUATION

The equation for the determination of dynamic range, under the approach of ISO 15739, can be developed from fundamental principles.

For later reference, note that:

$$N_D = G_i N_L \quad (4)$$

where N_D is the noise as seen in the digital image, N_L is the equivalent noise on a luminance basis⁸, and G_i is the incremental gain, the slope of the gamma precompensation function of the digital encoding scheme at the luminance of interest. This can be rewritten as:

$$N_L = \frac{N_D}{G_i} \quad (5)$$

Now, signal to noise ratio is defined by:

$$S = \frac{L}{N_L} \quad (6)$$

where S is the signal to noise ratio, N_L is the observed noise on a luminance basis, and L is the luminance involved (N_L and L in consistent units).

We are interested in the luminance, L_0 , for which the signal to noise ratio, using the noise as measured at reference black luminance, would be 1.0.⁹ Thus:

$$L_0 = N_L \quad (7)$$

where N_L is the noise, on a luminance basis, observed at the reference black luminance (at which the noise is actually tested).

The dynamic range, D , is defined by:

$$D = \frac{L_{clip}}{L_0} \quad (8)$$

Substituting equation 7 into equation 8, we get:

$$D = \frac{L_{clip}}{N_L} \quad (9)$$

Then, substituting for N_L from equation 5 into equation 9, we get:

$$D = \frac{L_{clip} G_i}{N_D} \quad (10)$$

⁸ Note that the standard never speaks of the noise on a luminance basis—that is my notation. It gets to the result we get in a different way.

⁹ Recall that at a luminance of L_0 , or for any luminance in fact, the actual signal to noise ratio cannot be 1.

Note that this equation is consistent with the equation 3, and thus with the equation given in the standard itself (essentially equation 2, from which I obtained equation 3).

#

APPENDIX B

CURIOSITIES OF THE STANDARD

Several curiosities are encountered in the study of ISO 15739-2003. They serve to confound understanding of the definitions and procedures given by the standard.

The definition of L_{sat}

The luminance that is 100/140 of the luminance that produces saturation of the imaging system is inexplicably designated L_{sat} .¹⁰

The multi-patch test target

One procedure for both noise and dynamic range testing involves a special test target with patches having certain prescribed reflectances (actually specified in terms of *density*¹¹). The “lightest” patch (patch 12) has a density of 0.1 (reflectance of 0.79). The procedure prescribes setting the exposure so that the image of patch 12 saturates the imager.

Nevertheless, the luminance at which the noise is to be measured for the determination of dynamic range is the luminance established by patch 1, with density 2.0 (reflectance 0.01). Thus, the reflectance of patch 1 is 0.0125 that of patch 12. If patch 12 gave saturation luminance, then patch 1 would give a luminance 0.0125 of saturation luminance. But the noise is supposed to be measured at a luminance of 0.01 of saturation luminance.

This cannot be resolved by the conjecture that L_{sat} is really intended to refer to the saturation luminance. That would just lead to a different inconsistency!

I have been unable to resolve this conundrum.

Other procedures defined do not use a multiple-patch test target. They involve the use of individual “attenuators” of defined transmission density. They provide for first setting the intensity of the light source, with no attenuator in place, to the level that produces saturation of the imager. The appropriate attenuator is then introduced into the path for the measurement. For the measurement of noise in connection with dynamic range, that is an attenuator with a density of 2 (transmission of 0.01). There is no conundrum in that situation—all the numbers work out properly.

#

¹⁰ It is not easy, from the language of the standard, to discern that this is what L_{sat} represents. The equivalence I state can be confirmed by analysis of the various definitions and equations.

¹¹ Strictly, density applies to a transmissive object, and is defined as the negative logarithm (base 10) of the transmission. The measure has been adapted for reflective objects, where it is defined as the negative logarithm of the reflectance.