Dynamic Range

WHITE PAPER

Date: 28th September 2022

Table of Contents

Table of Figures

Introduction

For more than 10 years, ARRI ALEXA cameras have been trusted by the world's top filmmakers. A key reason for their widespread adoption is the high dynamic range of the cameras. The high dynamic range allows images to be rendered with a soft, film-like roll-off in the highlights, and provides the exposure latitude needed in challenging situations. When it comes to these attributes, ALEXA has set the benchmark against which all others are measured.

This document explains how ARRI measures the dynamic range of its cameras. The design of the sensor in the original ALEXA cameras was so advanced that it led the industry in terms of dynamic range for 12 years. During that time, larger format versions of the sensor were introduced that combined the superior dynamic range with greater resolution. It was only when ARRI released the ALEXA 35, with its new 4.6K Super 35 format sensor, that even greater dynamic range was achieved.

Since the introduction of the ARRI ALEXA camera family, filmmaking has changed. Wide gamut and HDR displays have emerged and content for these displays is produced by OTT providers. With an increased range of tonal values in the display, the camera needs to have a greater dynamic range to maintain the creative possibilities for color grading. On the lighting side, LED illumination has appeared on set and everywhere else. ARRI's lighting fixtures can be used to illuminate sets or stages in pure colors. Capturing the highly saturated colors from these fixtures also requires a large dynamic range in the camera.

Dynamic Range in the Scene

The tonal values in an image are, put in physical terms, surfaces in the scene reflecting or emitting more or less light. The brightness of surfaces is measured as →luminance and its SI unit is candela per square meter (cd/m^2). In the US, the synonym term nits is generally used. Another unit is →footlamberts (fl). When the brightest object has a luminance of 4096 cd/m^2 and the darkest object has a luminance of 0.5 cd/m^2, the contrast ratio or dynamic range in the scene is 8192 (4096:0.5). Photographers and cinematographers measure this ratio in →stops, which in this example is 13 (2^{13} = 8192).

The physical quantity a spot meter measures is luminance. The meter, however, will present the result in a combination of aperture and exposure time for the selected ISO speed[1]. The absolute brightness of the scene is of lesser importance, since it can be compensated for by several means: the aperture, the use of ND filters, or the exposure index (EI) setting of the camera.

When the two objects in the example above have twice the luminance, 8192:1 instead of 4096:0.5, the contrast ratio stays the same. Closing the aperture by one stop will yield the same result in terms of tonal range.

The relative luminance of the elements in the scene translates linearly to exposure on the sensor. The photosites receiving the most light will receive 8192 time as much light as the photosites receiving the least light.

[1] The meter may also show an exposure value (EV) for the selected ISO speed. The two luminance values of 0.5 and 4096 cd/m^2 mentioned in the example are an EV of 2 and 15, respectively, at ISO 100. This yields immediately the dynamic range of 13 stops.

Image 1: Bar scene

The contrast ratio between the brightest (1) and the darkest (10) part in this scene is 12500:1 or 13.6 stops, see the table below.

Target	Tonal Value	Luminance in cd/m^2	Description
1	8.2	22777.0	Brightly illuminated lemon
2	7.3	11937.0	Brightly illuminated glass right
3	3.0	592.1	Naomi's right shoulder
4	1.2	181.6	Anton's cheek
5	1.0	151.1	Anton's forehead
6	0.3	96.9	Naomi's forehead
7	0.0	74.2	Gray card
8	-0.8	44.2	Anton's right shoulder
9	-3.1	8.7	Anton's left shoulder
10	-5.4	1.8	Shadow below the rim of the front side of the bar

Table 1: Luminance and tonal values in the bar scene.

The scene was recorded with an exposure time of 1/50 second and an aperture of T/10 (f/8 + 2/3). The camera was set to EI 800. These values yield the nominal correct exposure of the 18% gray card (at position 7). It is customary to represent this exposure as 0 and to express tonal values above and below as positive and negative stops, respectively.

Image 2: Details of bar scene

It's difficult to display a contrast of more than 13 stops in a SDR image without local adjustments. These image sections illustrate tonal details in the shadows in the highlights.

The scene depicted above is an example of a relatively high dynamic range. The contrast range of outdoor scenes with specular reflections of sunlight on metal or water may be even higher. Night scenes in cities with car lights and bright commercial signs are further situations with a large dynamic range. Here, the bright elements are very often colored. A sensor with a high dynamic range isn't only able to capture a large contrast ratio, but also a large ratio between the RGB channels, which represents very saturated colors. A red LED taillight of a car, for example, may yield an almost saturated red channel and a much lower signal in the other channels.

Lastly, we also need to consider the change of brightness over time. The example above shows a high contrast ratio within one frame. In movie scenes, the contrast changes often over time.

Dynamic Range of the Sensor

While the dynamic range of the scene is simply explained, we need to introduce a few details before we can define the dynamic range of the sensor.

Response of the Sensor

Light energy is quantized, it occurs as multiples of small "particles", called →photons. During exposure, the photosites in the sensor convert photons to electrons and accumulate the charges. This yields a voltage, which is converted to a digital signal at the end of the exposure. Each photosite has a maximum capacity for the charges. A higher exposure will not yield a higher signal beyond this point. The signal *saturates*.

The quantum nature of light leads to an effect that is important for imaging. As an illustration, imagine spreading small particles randomly over a checkerboard. One may get, for instance, an average of 10 particles in each field, but by chance some fields will contain some more and others fewer particles. It's a similar situation when the photons arrive at the sensor. Even in an area of uniform exposure, the number of photons arriving at each photosite will randomly fluctuate to a certain degree. This is called →shot noise.

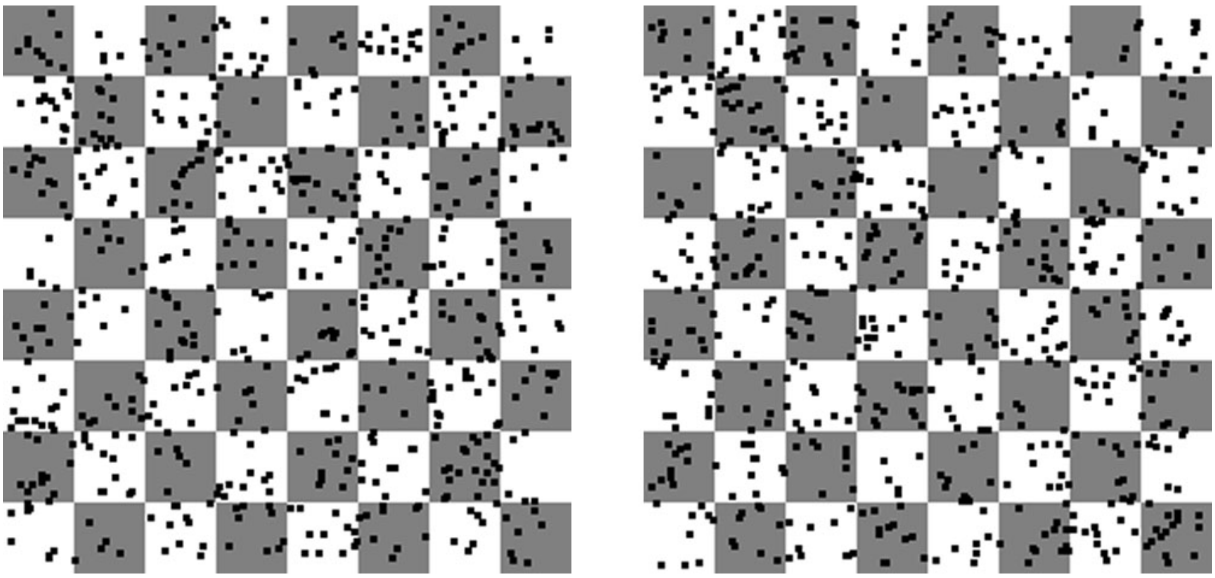

Figure 1: Shot noise

640 "particles" are randomly distributed on a checkerboard. The expected value is 10 particles per tile. Just by chance, some tiles have more particles and others have fewer. Spread out again, the number of particles on the *tiles will again vary by chance. The →standard deviation of the number of particles per tile within the checkerboard and on any one tile across many repetitions is the same. The situation is analogous to photons hitting the photosites of a sensor. We see the variation as spatial noise within one frame and as temporal noise over time.*

Noise is measured as the →standard deviation of the values. The ratio of signal to noise (→SNR) is a measure for the quality of the signal. A signal with a mean value of 10,000 and a noise of 100 is better than a signal with a mean of 100 and a noise of 10. Better means that when those signals appear in images, the former will look much less noisy (see also the images in the next section).

The shot noise is an inherent feature of light and can't be avoided. The only strategy is to collect as many photons as possible, because the SNR will increase with the number of photons. The inverse conclusion is that the visible noise in images will inevitably increase with decreasing exposure.

Another important characteristic is the →dark noise of the sensor. Dark noise means that even without any light falling onto the sensor, the digital numbers generated from the signal of the photosites will fluctuate. At a very low exposure the signal of the sensor will be less than the standard deviation (SNR < 1) because the dark noise is always added. When the exposure and therefore the signal increases, it will at some point reach the level of the noise (SNR = 1). This is called the *sensitivity threshold*.

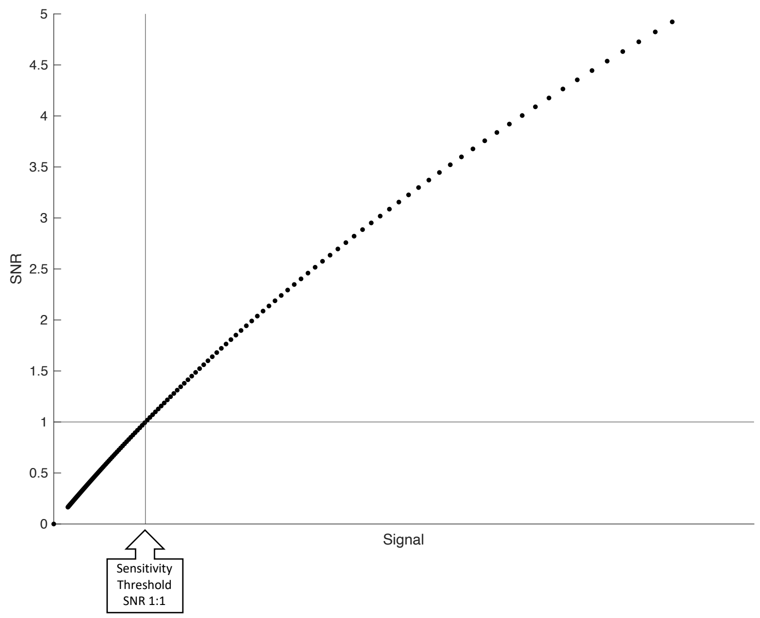

Figure 2: Model simulation of SNR

Model simulation of the SNR of a sensor depending on the signal. The signal increases with exposure, with the signal increases the SNR. The point where the SNR of the signal reaches a value of 1 is called the sensitivity threshold.

Dynamic Range of the Sensor

The ratio of the saturated signal to the sensitivity threshold is the dynamic range of the sensor. This is not an ARRI invention; it is in accordance with general engineering practice as outlined in ISO 15739[2] or in EMVA Standard 1288[3].

---

[2] ISO 15739: Photography — Electronic still-picture imaging — Noise measurements

[3] EMVA Standard 1288: Standard for Characterization of Image Sensors and Cameras.

Below is a series of images exhibiting decreasing SNR values. The first image (SNR = 40) appears almost free of noise. In the second image (SNR = 10) one can clearly see the noise but one can read all text lines. The remaining three images have SNR values of 2, 1, and ½. With decreasing SNR it becomes harder and eventually impossible to read the smaller text. The images illustrates that a SNR of 1 is a sensible lower bound for the signal.

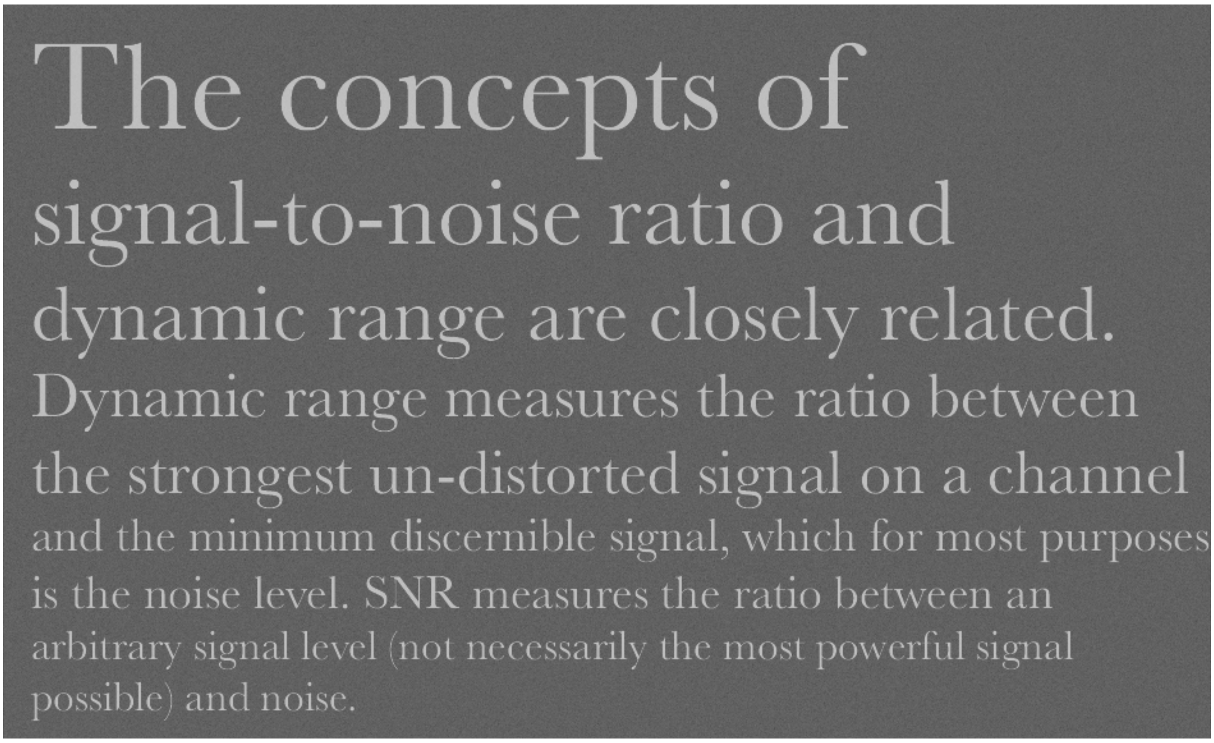

Image 3: SNR 40:1

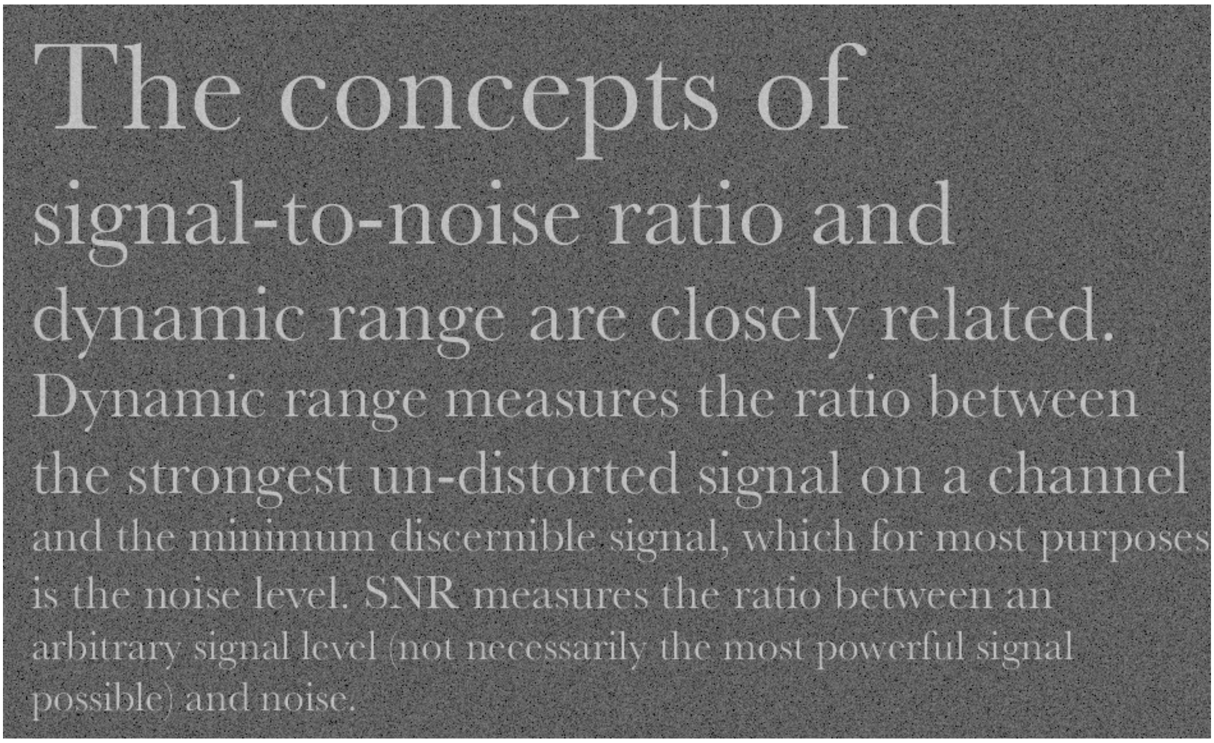

Image 4: SNR 10:1

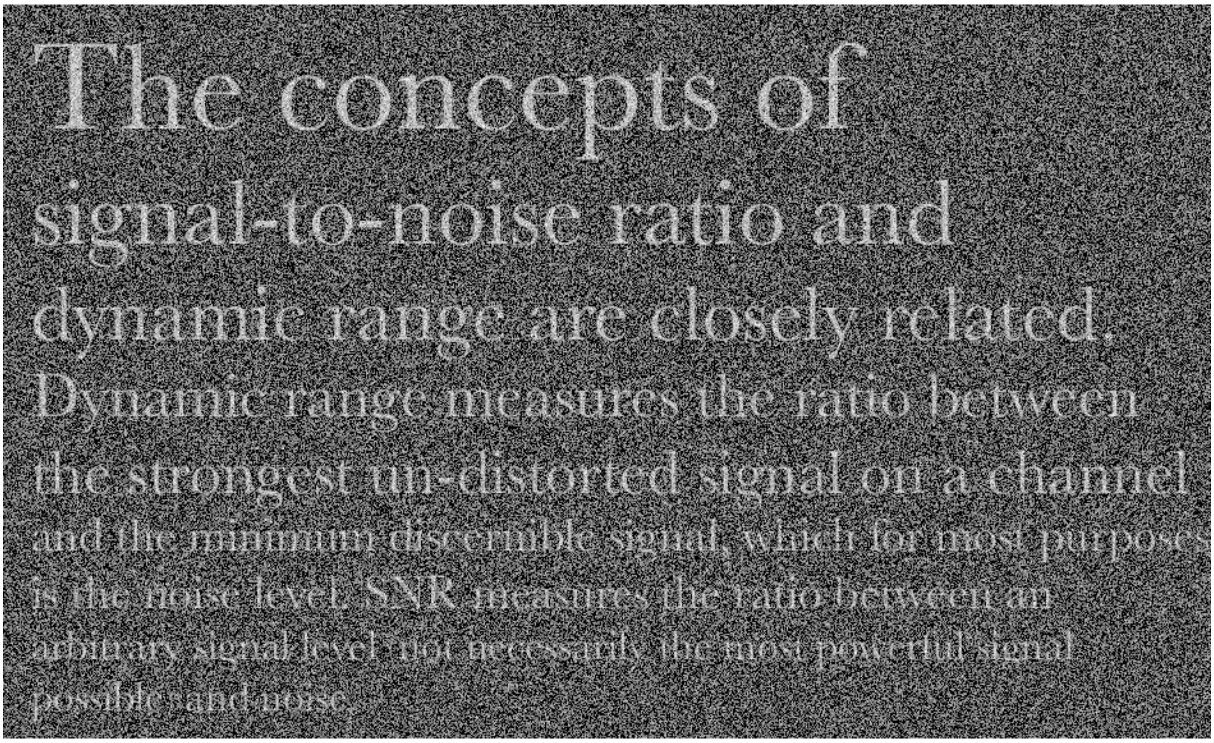

Image 5: SNR 2:1

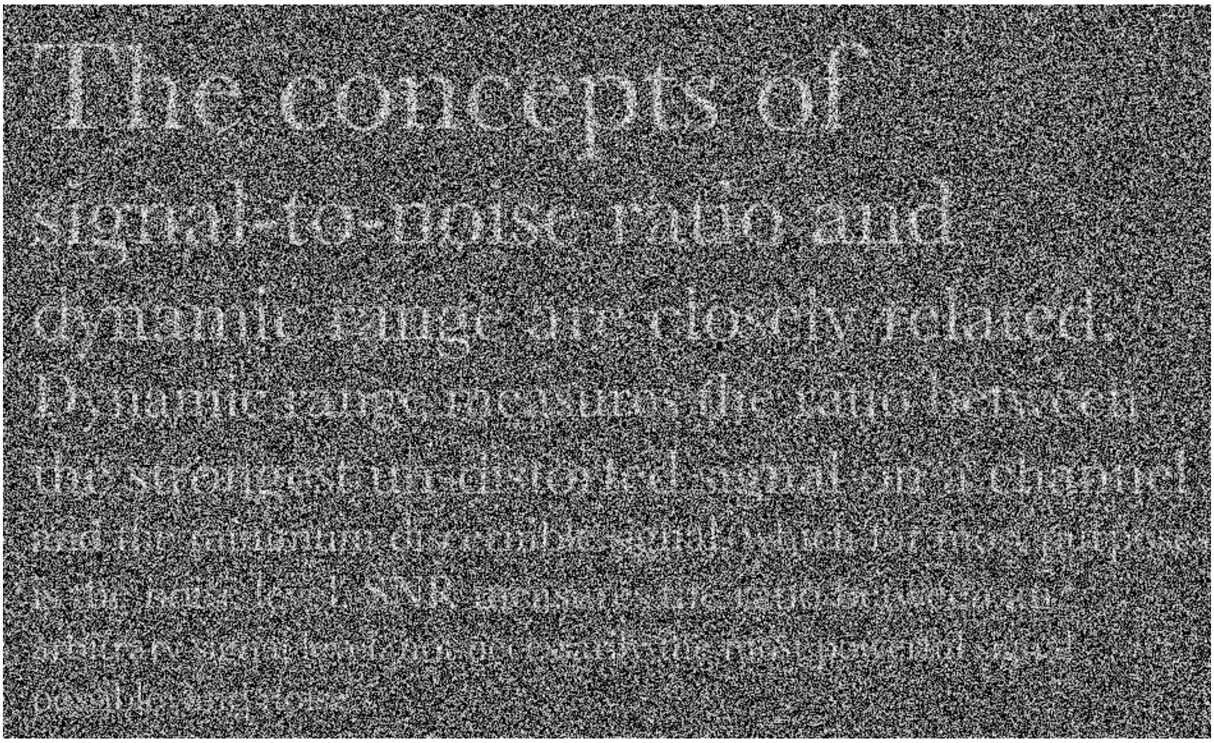

Image 6: SNR 1:1

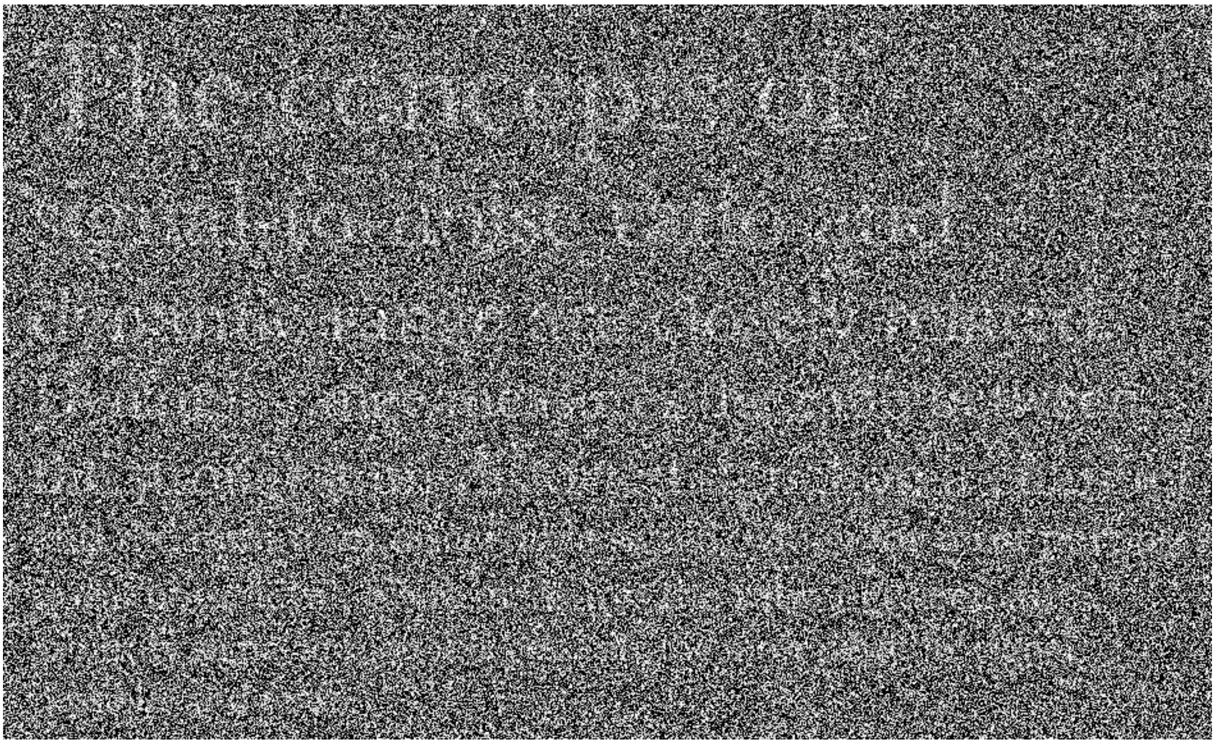

Image 7: SNR 1:2

Measurement Procedure

In theory the procedure to measure the dynamic range is uncomplicated. One captures a series of images starting from black (no light) and increases exposure until the signal is saturated. Details of this procedure are described in the references mentioned above. One can either place the camera in front of an illuminated surface and control its brightness or one can use a test chart like the DSC Labs Xyla chart, which yields a range of exposure spanning 20 stops within one frame.

One should not, however, just try to count the barely visible patches in the dark region of the test chart. As shown in the examples above, one can recognize large patterns like the words in the top row even in a signal with an SNR smaller than 1. But one can't recognize any smaller details.

Image 8: DSC Labs Xyla chart, EI 800

The picture shows the DSC Labs Xyla chart photographed with the ALEXA 35 as a LogC4 image. The image is processed with an →exposure index (EI) of 800. Depending on the monitor and the viewing conditions, one may be able to recognize patch #18 or #19. ARRI specifies a dynamic range of 17 stops for the ALEXA 35.

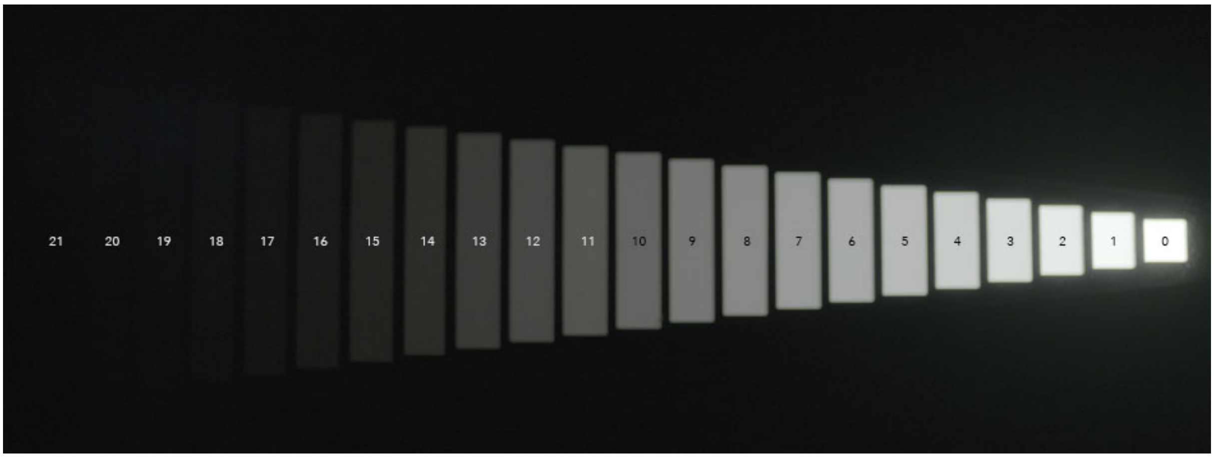

Image 9: DSC Labs Xyla chart, EI 3200

The figure shows the same image processed with an EI of 3200. In this image one may even be able to recognize patch #19 or #20. This illustrates another issue with the visual evaluation. The result will depend on the type of transfer curve applied to the image. The dynamic range of the camera is not changed by increasing the exposure index. The underlying signal from the sensor is the same as above. The camera still has a dynamic range of 17 stops.

If one uses the last, barely visible patch as the lower threshold, one will overestimate the dynamic range of the sensor. It is instead necessary to determine the SNR for each exposure. This will yield data similar to the one illustrated in Figure 2. The dynamic range is the ratio of the exposure where the signal saturates to the exposure where the signal has an SNR of 1. There are commercial applications available that perform the necessary image analysis and computations.

Sensor vs Camera

Engineers measure the dynamic range of a sensor based on the linear raw signal values.

Cinematographers are more interested in the dynamic range of a camera. Therefore, it’s sensible to include the image processing of the camera. Most cameras process images into a log-like RGB encoding (ARRI uses the LogC3 and LogC4 curves). This creates another issue when a visible examination of a test chart is used to derive the dynamic range. All vendors use different encoding curves, even in different versions, which makes results difficult to compare.

For a quantitative measurement of SNR, the resulting images need to be converted into a linear domain. The dynamic range values of ALEXA cameras published by ARRI are measured in this way. The images are processed into LogC3 or LogC4 and then the inverse of the respective curve is used to linearize the image data. With this linear data we calculate the SNR of each exposure and determine the sensitivity threshold and the point of saturation.

ALEXA Mini LF	14.5 stops
ALEXA 35	17 stops

Table 2: Dynamic range of ARRI cameras.

The values are derived according to the method explained in the text. The result for the ALEXA 35 is measured with the ARRI Texture “K445 Default” and without Enhanced Sensitivity mode.

An alternative approach, which delivers very similar results for ARRI cameras, is to process the images into the ACES color space to perform the analysis. This has the advantage that results from different camera systems can be compared.

The ALEXA 35 has two new features that will influence the dynamic range of the camera. ARRI Textures influence the noise in the image in different ways, which will affect the SNR. Some textures will decrease the measured dynamic range by a ¼ stop, others will increase it by a ½ stop. Enhanced Sensitivity (ES) is an option that can be used with high exposure index values of 2560 or greater. It reduces noise in such circumstances, producing cleaner shadows.

The effects of textures and ES mode are dependent on the image content. Recording test charts should not be the only basis for their usage.

Photographic Dynamic Range

The sensitivity threshold is what engineers regard as the lower bound of the signal. As explained above, the noise will be as high as the signal. For a cinematographer, this high relative noise level may not be acceptable. Other characteristics than the dynamic range of the photosites need also to be considered. For example, the camera should not show line noise in the shadows, nor should shadows exhibit any unwanted coloration.

Some people prefer to use an SNR of 2 as the lower threshold when measuring the dynamic range. That yields a value that is at least 1 stop lower than the one reported for a SNR of 1 as is explained below.

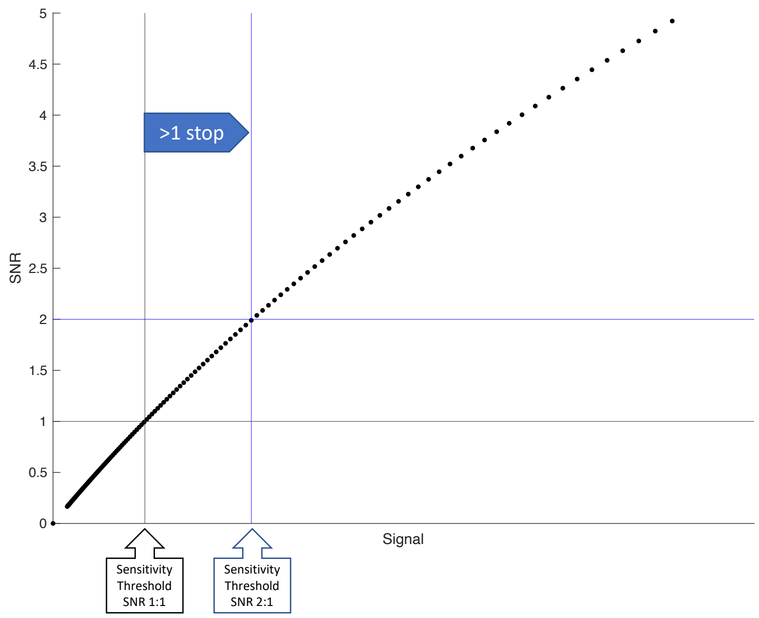

Figure 3: Different thresholds

When an SNR of 2 is used as the lower threshold, the reported dynamic range will be at least one stop lower. If the SNR increased linearly with the signal, doubling the signal would double the SNR. As one can see in the figure the relation isn't linear. Therefore, a bit more than double the exposure is needed to raise the SNR from 1 to 2.

The relation between signal and SNR isn't linear because adding more light increases the shot noise too.

Ultimately, it's the responsibility of the cinematographer to determine the minimum exposure needed to produce the desired quality. Since the upper limit (the signal saturation) is easy to determine and undisputed, it's this decision that defines the usable photographic range of a camera.

Exposure

Based on a clear understanding of the concept of dynamic range, this section provides some insight in the exposure of digital cameras.

False Color Zones

The ARRI ALEXA cameras have a feature that helps to find the optimal exposure. The →Log C image is converted into a grayscale image and certain signal regions are indicated in color. The first purple zone shows where the signal has an SNR of 1 or lower. The table below explains the meaning of the other color zones.

Color	Explanation
Red	1/3 stops below clipping (saturation)
Yellow	2/3 stops below clipping (saturation)
Pink	1 stop over 18% middle gray
Green	18% middle gray
Blue	Edge of shadow detail (SNR <= 3)
Purple	Noise floor (SNR <= 1)

Table 3: ALEXA False color zones

When a gray card is used (or another object in the scene with a similar reflectance level), the exposure can be set such that the card is in the green zone. At the same time, the purple and red zones will indicate when the limits of the sensor are exceeded. There is a subtle difference between the purple and red zones. A sensor has a signal below the saturation threshold; it may not contain useful image details, but it can be used in the false color display. On the other end of the exposure range, a sensor can't discriminate exposure beyond the saturation. Therefore, the red zone shows saturation before it happens.

Exposure Index

The exposure (the combined effect of scene brightness, aperture, ND filter, and shutter time) determines the signal level in the sensor. The exposure index setting in the camera selects which sensor signal becomes →middle gray in the Log C image and consequently in the monitor image. We explain the differences between LogC3 and LogC4 later.

All ARRI ALEXA cameras feature an exposure normalized Log C signal (refer to the section on page 19). This means that when one closes (opens) the aperture and increases (decreases) the exposure index by an equivalent amount, the overall image brightness will be the same. Neither varying the exposure nor the exposure index changes the dynamic range of the sensor.

The following diagram illustrates the relation between the contrast ratio of the scene and the dynamic range of the sensor, and between exposure and exposure index.

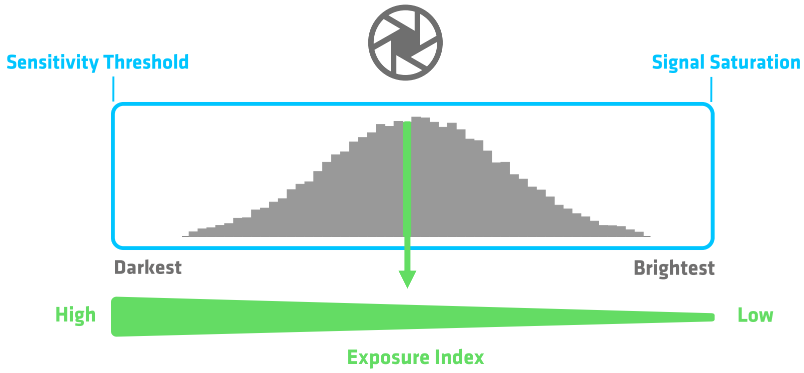

Figure 4: Exposure diagram.

The width of the histogram symbolizes the contrast ratio of the scene. The left end represents the darkest part of the scene, and the right end represents the brightest part. For simplicity, middle gray is shown as the center of the scene contrast range by a vertical green line, even while the tonal range in a scene may not be symmetrically distributed.

The blue rectangle represents the dynamic range of the sensor, which is constant.

The exposure is controlled by aperture, ND filter, and shutter time. The combined effect is represented by the aperture symbol shown on top of the blue rectangle.

The exposure changes the mapping of the tonal values in the scene to the sensor signal. The exposure index selects which sensor signal becomes middle gray in the Log C image and consequently in the monitor image. The exposure index matching the exposure is shown by the arrow that extends the middle gray line in the histogram. The value of the exposure index is inversely proportional to the exposure.

In a situation where the contrast ratio of the scene is less than the dynamic range of the sensor the cinematographer is to a certain degree free to choose the exposure. Changing the exposure index accordingly will maintain the same brightness in the Log C image.

In the situation depicted below the contrast ratio of the scene is less than the dynamic range of the sensor. The cinematographer is to a certain degree free to choose the exposure without exceeding the limits of the sensor. Changing the exposure index accordingly will maintain the same brightness in the Log C image. When the exposure is decreased (top row), the histogram moves left. Therefore, the EI is increased. The resulting image will have a lower SNR overall—it will exhibit more noise. Increasing the exposure is depicted by moving the histogram to the right (bottom row). To keep 18% reflectance at middle gray, the EI is decreased. The resulting image will have a higher SNR overall.

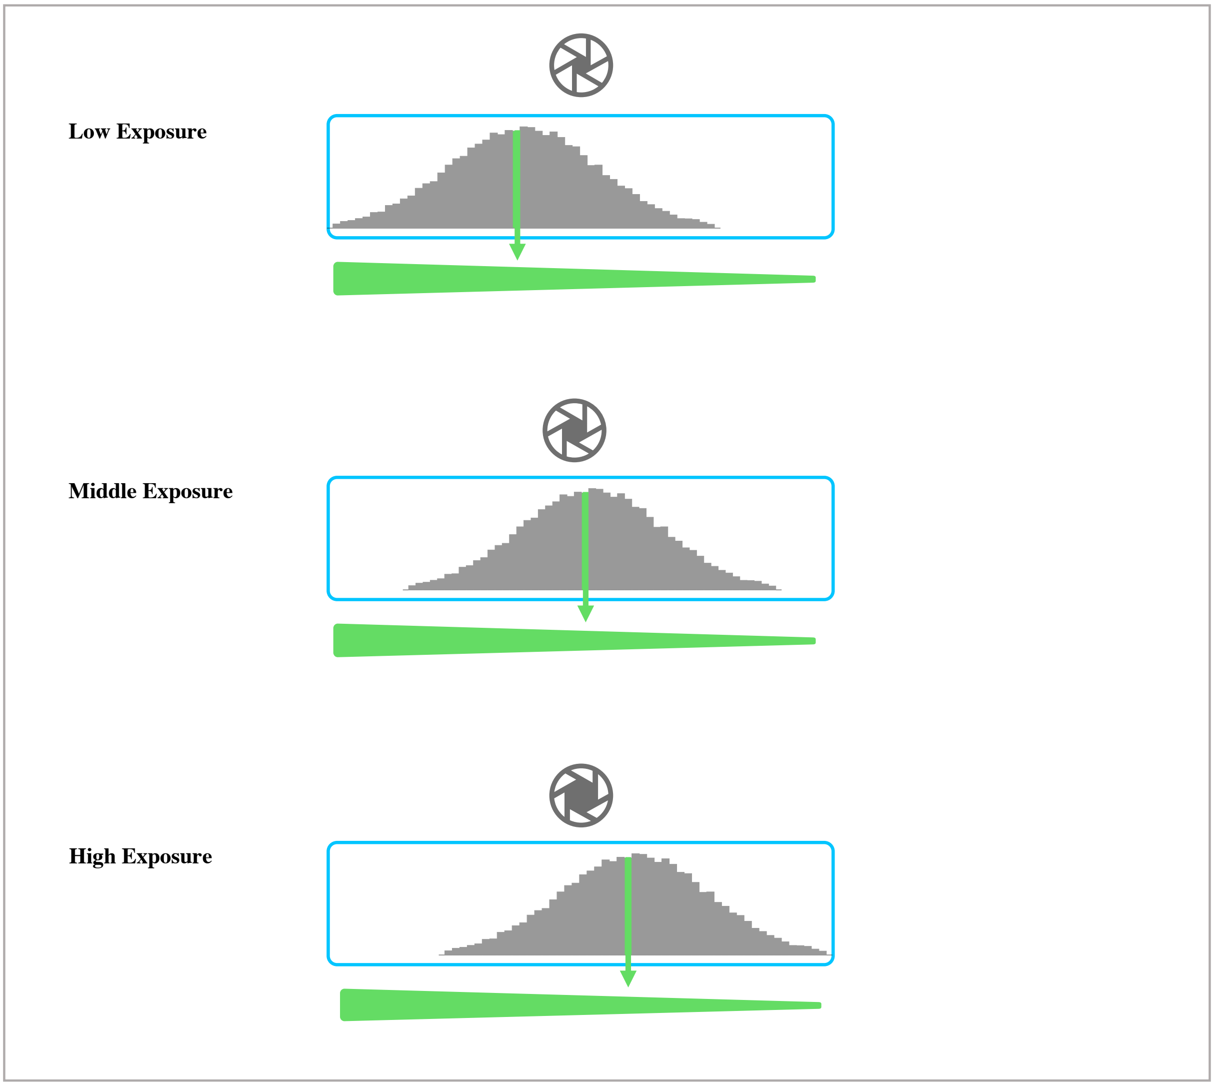

Figure 5: Exposure levels

In other situations, the primary reason to adjust the exposure is not the placement of the 18% gray but the desire to preserve shadows or highlights.

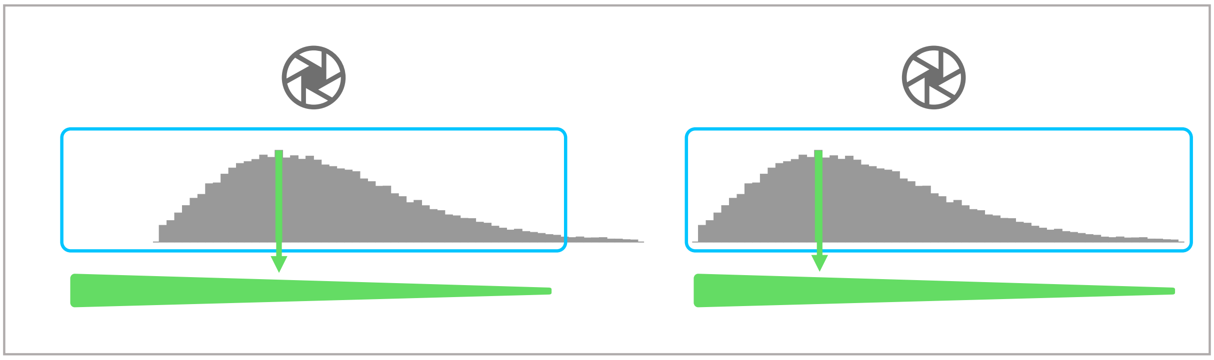

Figure 6: Expose for highlights

In this situation the highlights of the scene exceed the dynamic range of the sensor. To avoid clipping, the exposure needs to be decreased, which is compensated for by increasing the EI.

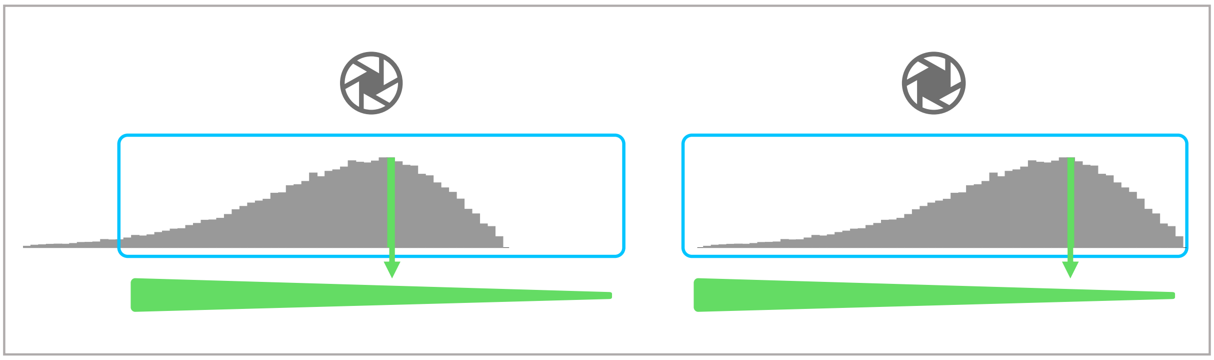

Figure 7: Expose for shadows

In this situation the shadows of the scene exceed the dynamic range of the sensor. To avoid losing shadow detail in the signal noise, the exposure needs to be increased, which is compensated for by decreasing the EI.

A method called “expose to the right” is used to maximize the SNR in each situation. Exposure is changed until the maximum scene brightness is just below the sensor saturation. The exposure index is adjusted to maintain the mid gray.

With an ALEXA camera this can be accomplished by first changing the exposure such that no or just a little red indication is visible in the false color image. Then, the exposure index is adjusted until the 18% gray is indicated green.

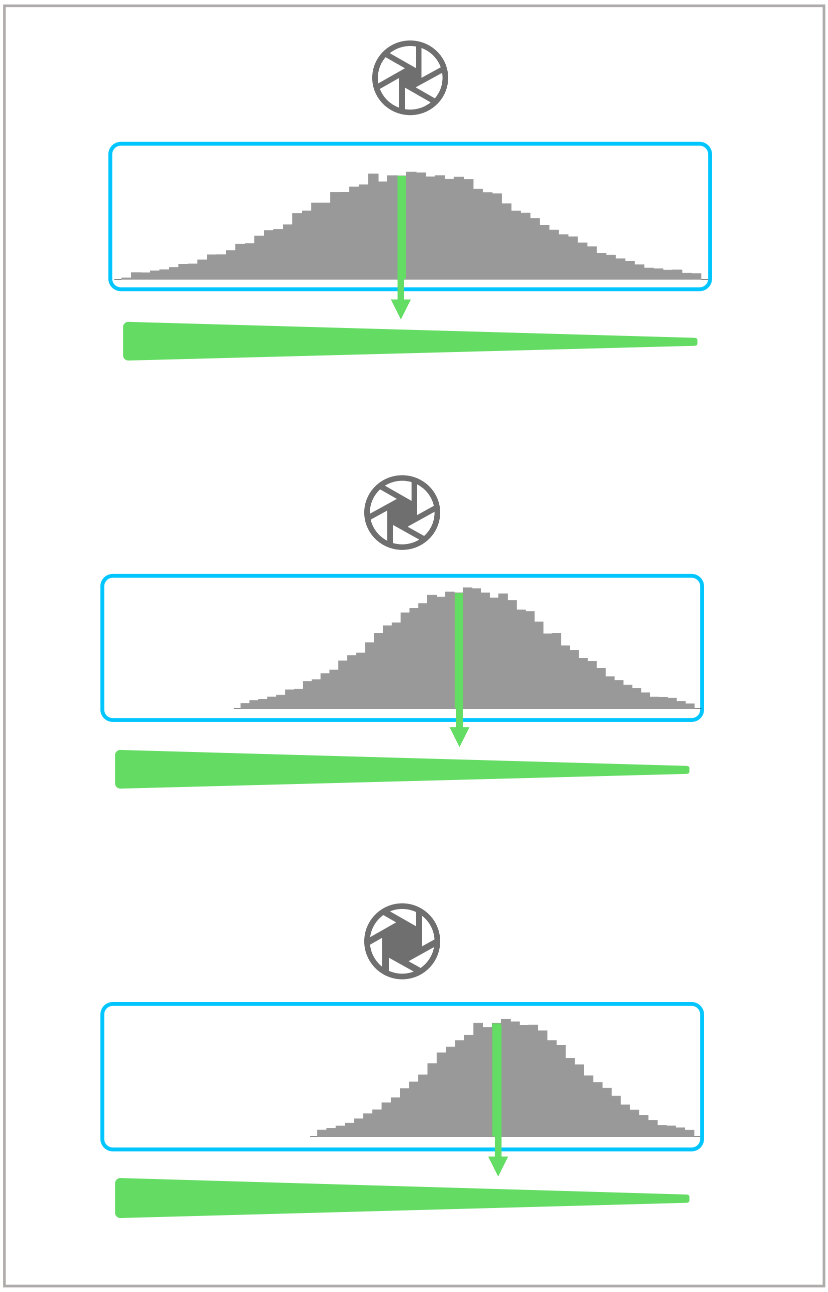

Figure 8: Expose to the right

These diagrams explain the "expose to the right" method. Exposure is changed until the maximum scene brightness is just below the sensor saturation. The exposure index is adjusted to maintain the mid gray.

Exposure Latitude

Since the dynamic range of the sensor is constant, the exposure latitude above and below mid gray changes with the exposure. This can be seen in the diagrams below.

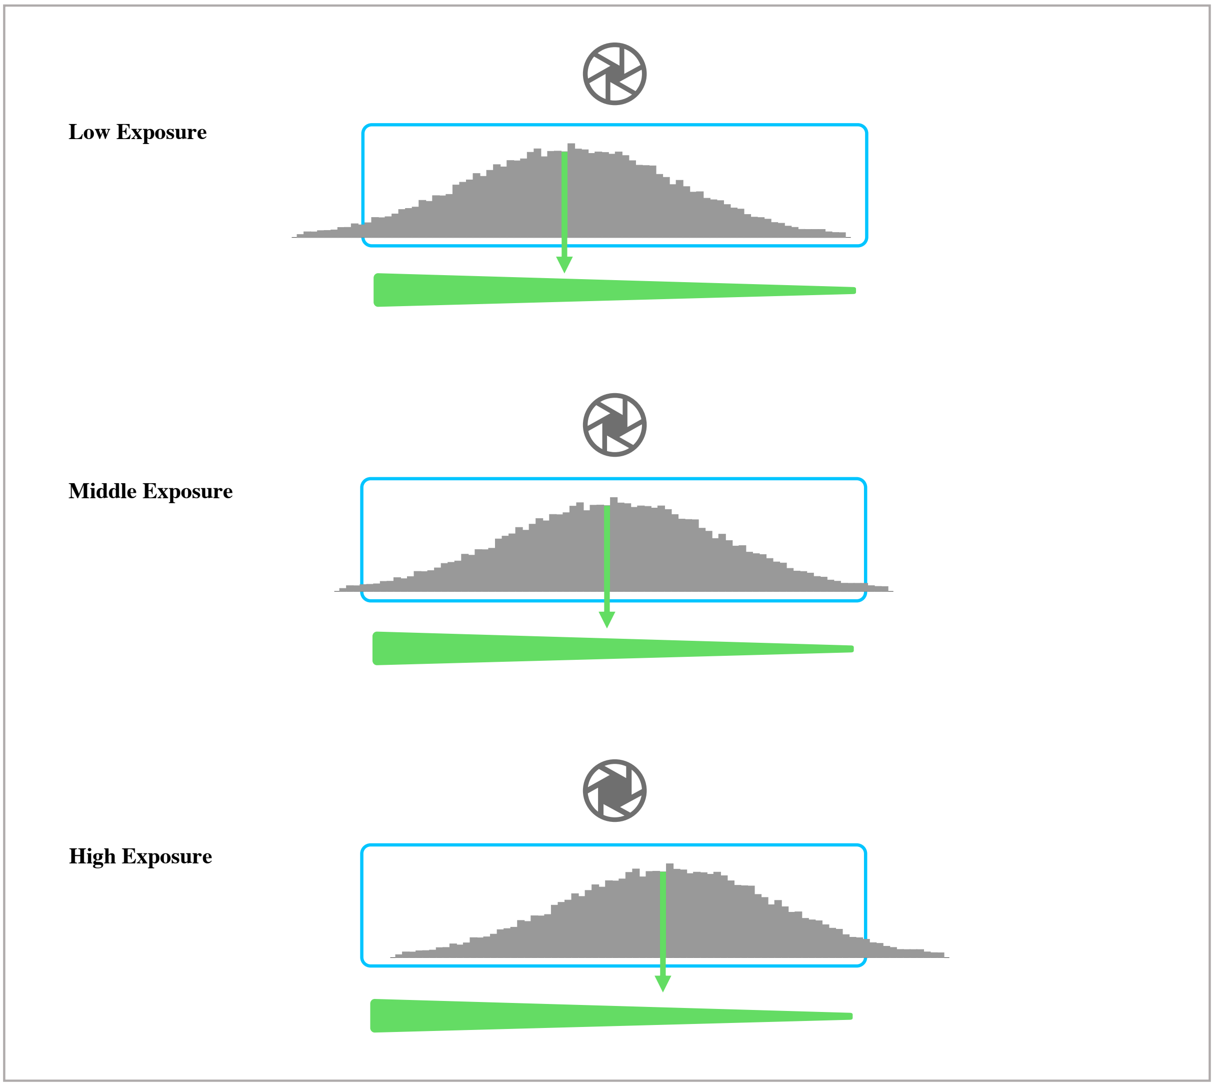

Figure 9: Exposure latitude

When the contrast in the scene is larger than the dynamic range of the camera, the cinematographer needs to decide if more of the shadows or more of the highlights should be captured. It is the change to exposure that is important, not the adjustment of the exposure index. When ARRIRAW is recorded, the exposure index can be set later without any loss of image quality.

When the exposure is reduced by one stop, one more stop of highlights can be recorded. On the other end, one stop of tonal values is pushed below the sensitivity threshold. When the exposure is increased by one stop, one more stop of tonal range will be above the sensitivity threshold, but the sensor goes into saturation earlier. The following table provides the exposure latitude below and above 18% gray for ALEXA cameras.

	ALEV-3-based cameras			ALEV-4-based cameras		
ISO	Below 18%	Above 18%	Total	Below 18%	Above 18%	Total
160	9,2	5,3	14,5	10,0	7,0	17,0
200	8,9	5,6	14,5	9,7	7,3	17,0
250	8,5	6,0	14,5	9,3	7,7	17,0
320	8,2	6,3	14,5	9,0	8,0	17,0
400	7,9	6,6	14,5	8,7	8,3	17,0
500	7,5	7,0	14,5	8,3	8,7	17,0
640	7,2	7,3	14,5	8,0	9,0	17,0
800	6,9	7,6	14,5	7,7	9,3	17,0
1000	6,5	8,0	14,5	7,3	9,7	17,0
1280	6,2	8,3	14,5	7,0	10,0	17,0
1600	5,9	8,6	14,5	6,7	10,4	17,0
2000	5,5	9,0	14,5	6,3	10,7	17,0
2560	5,2	9,3	14,5	6,0	11,0	17,0
3200	4,9	9,6	14,5	5,7	11,4	17,0
4800	NA	NA		5,1	11,4	16,4
6400	NA	NA		4,7	11,4	16,0

Table 4: Exposure latitude of ALEXA cameras

The table provides the exposure latitude below and above 18% gray in fractional stops. The sum of the latitude below and above equals the dynamic range of the camera. In the ALEXA 35 an encoding limit is reached for EI greater than 3200. This is explained in the next chapter.

Log C Encoding

The Log C encoding maintains a fixed relation between signal values and the relative exposure in the scene. For example, in LogC4 a signal value of 18% represents an exposure of 2 stops below mid gray (which is at a signal level of 28%). A signal value of 34% represents 1 stop above mid gray. These relations stay constant regardless of the exposure index. A table with relevant LogC4 values is given in the appendix.

In LogC3 the relation between signal values and relative exposure stops changes somewhat with the exposure index. The appendix includes a table for an EI of 800.

Because the number of stops above mid gray that can be captured increases with decreasing exposure, the maximum of the LogC4 function will change with exposure index. This can also be seen in Figure 9.

Note: In the ALEXA 35 an encoding limit in LogC4 is reached when the EI is above 3200. The maximum value of the LogC4 function becomes greater than 1, which can’t be encoded in a 12- or 16-bit integer number. This means that no tonal value higher than 11 1/3 stops above mid gray can be encoded. ARRI regarded this limitation as acceptable because the advantage of being able to select a high exposure index outweighs the possible disadvantage of clipping of highlights more than 11 stops above mid gray.

Contact

In case you have questions or comments, please contact: digitalworkflow@arri.de.

Appendix

The appendix provides some detailed information about the LogC3 and LogC4 encoding function.

The LogC4 function is a successor to the LogC3 function, which has been used in all ALEXA cameras since 2011. The LogC4 function is used in the ALEXA 35 camera.

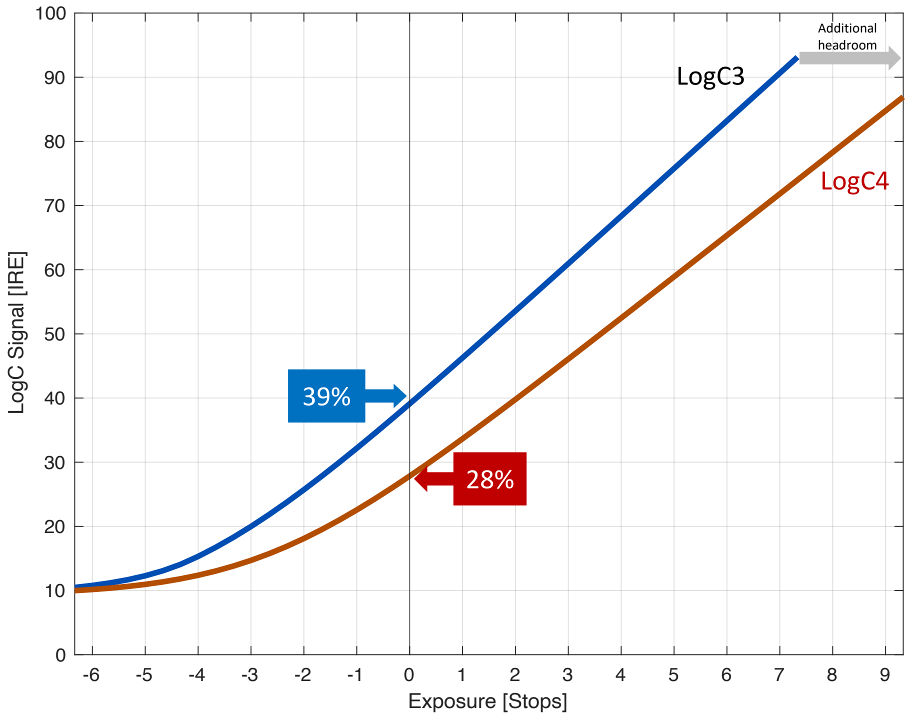

Figure 10: LogC3 and LogC4 encoding curves

The LogC4 curve has a lower slope. Displayed on a monitor without lookup table, the LogC4 image is darker than a LogC3 image. It maps the scene values to lower signal values and by doing so, it provides additional headroom for the dynamic range of the ALEXA 35 camera.

Displayed on a monitor without lookup table, the LogC4 image is darker than a LogC3 image.

Image 10: Images in LogC3 and LogC4

The image on the left side was captured with an ALEXA Mini LF in LogC3. The image on the right side was captured with an ALEXA 35 in LogC4. Both images are shown without lookup table as they would appear on a Rec. 709 monitor.

ARRI provides lookup tables for the conversion of images from LogC3 and LogC4, respectively, to monitor color spaces like Rec. 709. These lookup tables are adapted to the different signal values in the two versions of the Log C function. For example, the conversion lookup table for LogC3 will map a

LogC3 signal of 39% to a video signal of 40%. The conversion lookup table for LogC4 will map a LogC4 signal of 28% to the same video signal of 40%. As a result, the video images originating from the two Log C encodings will look very similar.

Image 11: Images in Rec. 709

Stops	LogC3 IRE	LogC4 IRE	Stops	LogC3 IRE	LogC4 IRE
-6 1/3	10.48%	9.99%	2 2/3	58.49%	43.97%
-6	10.79%	10.16%	3	60.95%	46.08%
-5 2/3	11.18%	10.38%	3 1/3	63.41%	48.20%
-5 1/3	11.68%	10.64%	3 2/3	65.88%	50.33%
-5	12.30%	10.96%	4	68.35%	52.47%
-4 2/3	13.09%	11.35%	4 1/3	70.83%	54.61%
-4 1/3	14.07%	11.82%	4 2/3	73.30%	56.75%
-4	15.31%	12.38%	5	75.78%	58.90%
-3 2/3	16.74%	13.04%	5 1/3	78.25%	61.05%
-3 1/3	18.29%	13.81%	5 2/3	80.73%	63.21%
-3	19.98%	14.70%	6	83.21%	65.36%
-2 2/3	21.79%	15.71%	6 1/3	85.69%	67.51%
-2 1/3	23.70%	16.85%	6 2/3	88.17%	69.67%
-2	25.71%	18.11%	7	90.65%	71.83%
-1 2/3	27.80%	19.49%	7 1/3	93.13%	73.98%
-1 1/3	29.96%	20.97%	7 2/3	95.61%	76.14%
-1	32.18%	22.56%	8	NA	78.30%
- 2/3	34.45%	24.25%	8 1/3	NA	80.46%
- 1/3	36.76%	26.01%	8 2/3	NA	82.62%
0	39.10%	27.84%	9	NA	84.78%
1/3	41.47%	29.73%	9 1/3	NA	86.94%
2/3	43.86%	31.67%	9 2/3	NA	89.10%
1	46.27%	33.65%	10	NA	91.26%
1 1/3	48.70%	35.67%	10 1/3	NA	93.41%
1 2/3	51.13%	37.72%	10 2/3	NA	95.57%
2	53.58%	39.78%	11	NA	97.73%
2 1/3	56.03%	41.87%	11 1/3	NA	99.89%

Table 5: Tonal values and the resulting signal values in LogC3 and LogC4

Glossary

brightness
the subjective impression of →luminance

candela per square meter (cd/m^2)
the SI unit for →luminance

clipping
increasing the exposure beyond the →signal saturation will not yield a higher signal value, the signal is clipped

dark noise
arises because of random events on the atomic scale in the light detector

Note: This is the reason why the digitized signal from a sensor doesn’t start with the value zero. In an ARRI camera for example, a zero exposure (no light) is represented by a digital value of 256. This number is somewhat arbitrarily chosen but it provides more than enough range to encode the random fluctuations of the dark noise. Hence, when one examines an ARRIRAW frame recorded without light (lens cap closed), one will get an average value of 256 but the individual pixel values fluctuate around this value. Values below 256 represent negative values. There is a linear relationship between positive (>256) values and exposure. Conceptional difficulties arise when the sensor signal is converted to scene values since there can’t be less than zero light.

exposure index (EI)
numerical value that is inversely proportional to the exposure provided to an image sensor to obtain an image

Note: Setting the exposure index selects which sensor signal becomes mid gray in the Log C image and consequently in the monitor image. This is explained in more detail in the section on page 13.

footlamberts (fl)
a unit for →luminance, 1 fl = 3.426 cd/m^2

HDR
high dynamic range refers to images displayed on monitors that exhibit a significantly larger contrast than →SDR monitors

Log C
a family of transfer curves used in ARRI cameras

Note: Cameras based on the ALEV 3 sensor (ALEXA Classic, ALEXA XT, ALEXA SXT, ALEXA Mini, AMIRA, ALEXA LF, ALEXA Mini LF, ALEXA 65) use the LogC3 curve. The ALEXA 35 uses the LogC4 curve. In this document we use Log C (without version number) in a context where the version doesn’t matter.

luminance
is the photometric measure of the →brightness of surfaces reflecting or emitting light

middle gray

a tone that is perceptually about halfway between black and white

Note: 18% has been universally adopted as the average scene reflectance. A distinction must be made to the signal value that represents this tonal value in a video image. It isn't standardized and values between 40% and 50% are mentioned in various sources. It depends very much on the type of scene and the creative intent. The value chosen by ARRI produces an image at the lower end of the range, with more headroom for highlights.

The signal value for mid gray in the LogC3 curve is 39%. The signal value for mid gray in the LogC4 curve is 28%.

mean value

the arithmetic mean of a set of *n* values x_i is $\bar{x} = \frac{\sum_i x_i}{n}$, the sum of the values divided by n

nits

a term used in the US for →candela per square meter (cd/m^2)

photon

represents the minimum amount of energy of electromagnetic radiation

Note: This amount of energy is very small and, conversely, the number of photons emitted from a light source is extremely large. A light bulb, for example, emits in the magnitude of 10^{20} photons per second. Nevertheless, the number of photons arriving at a single photosite during the exposure of one frame is a few thousands (based on the average scene exposure at ISO 800).

photosites

the single light detectors of an image sensor

Note: These are not image pixels, because a pixel has red, green, and blue components. Each of the photosites is covered with a red, green, or blue filter. The process of reconstructing the RGB information for each pixel from that mosaic of colors is called demosaicing. The number of photosites on a sensor and the number of pixels in a final image may be different numbers.

color saturation

describes the purity of a color

SDR

standard dynamic range refers to images that are displayed on conventional video or desktop monitors

Note: Conventional LCD-based monitors can display a contrast ratio of up to1000:1. While this is equivalent of 10 stops, it doesn't mean that one can only reproduce a scene contrast of 10 stops. Photographic images are no linear representation of the world. Non-linear and local tone mapping is used to render a larger scene contrast in the image.

signal saturation

maximum detectable exposure, beyond this point the signal is clipped (→clipping)

shot noise
a form of noise observed in light detectors that arises because of the quantum nature of light

SNR
abbreviation for signal to noise ratio, the →mean value divided by the →standard deviation

standard deviation
the standard deviation of a set of *n* values x_i is $\sqrt{\frac{\sum_i (\bar{x}-x_i)^2}{n}}$, the square root of the average quadratic distance to the →mean value; it's a measure of the amount of variation of a set of values

stop
a measure for the relative exposure in a scene, equivalent to opening or closing the aperture by one f-stop, doubling or halving of the light entering the lens