

ALEXA Color Processing

Introduction.....	2
Color Processing Terminology.....	2
Image Files and Data.....	2
Grayscale Encoding.....	2
Color Encoding.....	3
HTML Demonstrations.....	3
ALEXA Output Overview.....	4
Logarithmic Encoded Data.....	4
Log C (film matrix off).....	4
Log C (film matrix on).....	6
Display Specific Encoded Data.....	7
Rec 709.....	7
DCI P3.....	9
Linear Encoded Data.....	9
Linear from ARRIRAW files.....	10
Linear from Log C files.....	10
ALEXA Workflow.....	11
Log C Images.....	11
Preview.....	11
Color Correction.....	12
VFX.....	13
LUT Size.....	13
Normalized and Photometric Video Conversion.....	14
QuickTime	14

Introduction

ALEXA is an extraordinary new 35 format film-style digital camera system designed for the motion picture and broadcast markets, consisting of three cameras and an extensive range of prime lenses, zoom lenses, accessories and recording solutions. ALEXA cameras offer exceptional image quality with the organic look and feel of film; their unequalled exposure latitude, high sensitivity and unique ARRI color processing provide sharp and natural images for HD and 2K with a cinematic image quality of breathtaking richness and detail.

This white paper gives an overview of the different output formats available from the ARRI ALEXA series of cameras from the point of view of the different color processing formats. Important terms are defined in the introduction, after which the different output formats are explained in detail. Differences between the color processing of Software Update Packet (SUP) 2.x and SUP 3.x are pointed out. This white paper is accompanied by a number of HTML files that allow an easy side-by-side comparison of different color processing formats.

Color Processing Terminology

Image Files and Data

Images can be stored in different file formats (QuickTime, DPX) or as binary data on tapes (HDCAM SR). Image files and tapes are just containers for an array of numbers, the values of the pixels. The meaning of those values needs to be known when the images are displayed or processed. Some image file formats contain additional information (**metadata**) that describes this meaning. In many cases, however, you need to know what's in the image container because there is no such data (as with tapes) or it's unreliable (as in DPX files).

The pixel values encode grayscale and color information. The information can describe the original scene or the displayed image, which is in general not the same.

Grayscale Encoding

There are three encodings of grayscale information commonly used in digital post production. The first two (linear encoding and logarithmic encoding) encode the grayscale characteristic of the original scene; the last (display-specific encoding) encodes the grayscale characteristic of the reproduced scene.

Linear Encoding

The image data can be a linear encoding of the relative brightness in the original scene. An object that is double as bright as another would be represented by pixel values double as high. It is the native response of digital image sensors.

Logarithmic Encoding

The pixel values could increase by a fixed amount with each increase of exposure measured in stops. This is a logarithmic encoding because stops are a logarithmic measure of scene brightness. It is obtained from scanned negative film or from the Log C signal of the ALEXA, for example.

Display Specific Encoding

The data can also be an encoding of the image that is designed for a specific type of display. Examples are images for television monitors (Rec 709) or images for digital cinema projectors (DCI P3). Since the two displays have different characteristics, the encoding needs to be adapted to the target display. Displays use neither a linear nor a logarithmic encoding. The relation between signal and brightness is a power function where the exponent is traditionally designated by the Greek letter **gamma**, γ .

Look up tables (LUT) are used to transform the grayscale characteristic of digital images.

Color Encoding

So far, we have been talking about the grayscale characteristic only. But the three values of a color pixel relate also to the amounts of three **primary colors** that are mixed to create that color. The primary colors define the outer boundaries of the **gamut** of the display. A display can not produce colors that are outside of this gamut.

Rec 709 and DCI P3

Two sets of primaries are important in digital postproduction. One is defined in ITU-R Recommendation BT.709, more commonly known as **Rec 709**. This is the international standard for HDTV displays. The other set is defined in the Digital Cinema System Specification and in SMPTE 431-2, we refer to it as **DCI P3**.

Film RGB

The meaning of the RGB values may be different when the image is a film scan. In this case the three values describe the density of the dyes in the negative film.

Camera RGB

When the image comes from a digital camera, the three values are the relative amounts of light seen through three color filters. It doesn't make sense to talk about the gamut of a camera since there are no colors a camera can't see.

The important question is if the image from the camera matches the **appearance** of the original scene. Because the conditions under which the image is viewed (e.g. dim room, low luminance) differ generally from the conditions of the scene (e.g. bright daylight) the reproduction can't be one-to-one with the original.

To transform the color information of digital images, one uses **matrices** or **three-dimensional look up tables (3DLUTs)**.

HTML Demonstrations

This document comes with several HTML pages that demonstrate different processing of ALEXA images. Those demonstrations are designed to be viewed on a Rec 709 (or sRGB) display.

[ALEXA Output Formats](#)

[Log C in SUP 2.x and SUP 3.x](#)

[Illumination Sources](#)

[Log C Film Style](#)

[Video in SUP 2.x and SUP 3.x](#)

[Size of 3DLUT](#)

[Normalized and Photometric Video Conversion](#)

The links are not working in some PDF viewers. In this case, open the file `html/index.html` in a browser and use the links on that page.

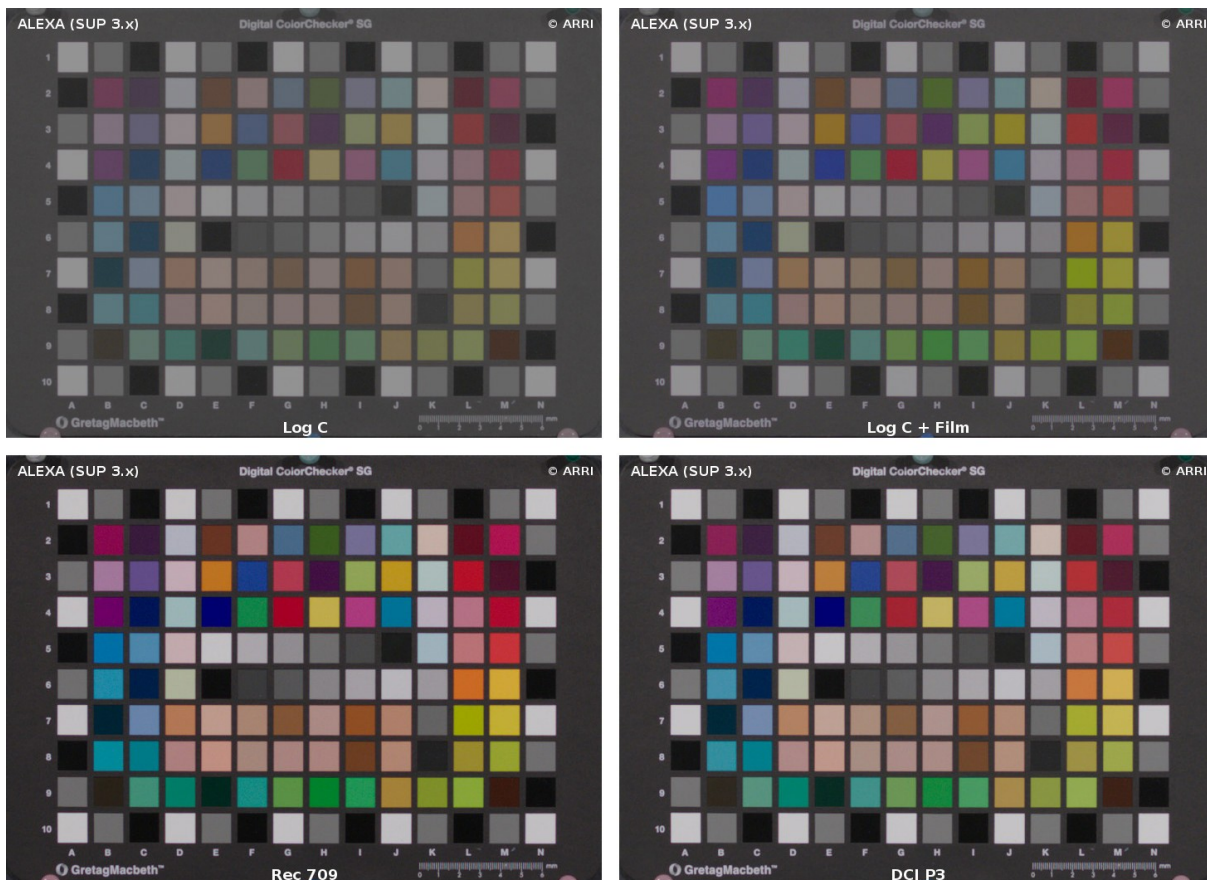
ALEXA Image Data

ALEXA Output Overview

The ALEXA cameras can output the following image data:

1. Linear encoding
 - ARRIRAW on T-Link
2. Logarithmic encoding
 - Log C (film matrix off) as SDI or as QuickTime
 - Log C (film matrix on) as SDI or as QuickTime
3. Display specific encoding
 - Rec 709 as SDI or as QuickTime
 - DCI P3 as SDI or as QuickTime

Open this [HTML](#) file in your browser to see a demonstration of the different ALEXA output formats.



Logarithmic Encoded Data

Log C (film matrix off)

The Log C curve was first introduced with the ARRIFLEX D-20 camera. It's an encoding with a grayscale characteristic similar to a scan from negative film. Because of the fundamental differences between digital cameras and negatives, the color characteristics remain different, though.

The Log C curve is depicted below. You see that the relation between exposure measured in stops and the signal is linear (straight) over a wide range. The slope of this part of the curve is called gamma. You see also the toe at the bottom of the curve. The toe occurs because the sensor can not see low light levels

with the same quantization as higher levels. The overall shape of the curve is similar to the exposure curves of film negatives.




Figure 1: Log C curve used in ALEXA SUP 3.x

Note also that the curve does not start at a signal level of zero. This is mostly done to be compatible with scanned negative where the clear base, which is black in the image, is traditionally represented by a code value around 95 in the 10 bit encoding (that is a signal level of approximately 10%). .

Differences between SUP 2.x and SUP 3.x

Prior to Software Update Packet (SUP) 3.0 the ALEXA camera used a similar curve as the D-20 with a lower and variable gamma. The gamma changes from 0.54 at EI 200 to 0.49 at EI 1600, which is within the gamma range of contemporary color negatives



The main improvements to the Log C curve in SUP 3.x are:

- Soft-shoulder for exposure index settings greater than EI 1600.
- Fixed black level of 95/1023.
- Larger linear part of 9-10 stops.
- Wide gamut RGB instead of camera raw RGB.

The last point needs some additional explanation. In SUP 2.x the Log C image was in native camera RGB. Starting with SUP 3.0 the Log C image is transformed in a wide gamut RGB. The gamut is larger than that of Rec 709 and DCI P3.



The changes to the Log C in SUP 3.0 have two visible results. The images have slightly more saturation and, for higher EI settings, the black is lower.

Open this [HTML](#) file in your browser to see a demonstration of the differences in Log C between SUP 2.x and SUP 3.0. Notice the difference in color saturation between the image from SUP 3.0 on the left side and the image from SUP 2.x on the right side.



Another advantage of the transform into the wide gamut RGB is that it allows a correction for different illumination sources. With the adjustment of the correlated color temperature, the camera does not only change the RGB gains to maintain a balanced image. It changes also the color matrix that is applied to the image. This adjustment was missing in the Log C images from SUP 2.x. Open this [HTML](#) file in your browser to see a demonstration of the test target photographed under two different illumination sources. In the HTML document you can see the differences in the Log C and in the Rec 709 image. Here only the latter is shown because the difference is better visible.

When you examine the images carefully, you'll notice that in the image made under the warmer tungsten light (left side) the red colors are more pronounced than in the image made under cooler HMI light (right side). That's the idea behind the correction: to make the images similar without removing the characteristic of the light used in a scene.





Log C (film matrix on)

Also a new feature in SUP 3.0 is a film style matrix that can be applied to the Log C output. The same transform is also available as a 3DLUT for post-processing of Log C footage.

The film style matrix makes the color characteristics of the Log C image similar to negative film scanned on an ARRISCAN (since scanners differ in their characteristics too, we can not say how the film style matrix relates to the output of other film scanners). The matrix is most usefully applied when the data is previewed or converted with a print film emulation (PFE). This is the common workflow in Digital Intermediate where the PFE is applied as a 3DLUT in the display path.

Open this [HTML](#) file in your browser to see a demonstration of the film style matrix. Below you see the ALEXA Log C image with film style matrix rendered through a PFE on the left side. On the right side you see a scan from a contemporary color negative film rendered through the same PFE. Only printer light adjustments (RGB offsets to the logarithmic images) were used to match the two the images. Both images were taken under tungsten light. Compare both images to each other and to the Rec 709 output of the ALEXA, which can be displayed on the right side too.



Differences between SUP 2.x and SUP 3.x

The film style matrix is introduced into the camera with SUP 3.0. The same transform is also available as a 3DLUT for post-processing of Log C footage. For software 2.0 only the 3DLUT is available. The film style look provided by the 3DLUT for SUP 2.x and SUP 3.x differ slightly.

Display Specific Encoded Data

Rec 709

The Rec 709 output of the camera is either used for on set preview or when the program is edited for television without extensive color correction. Those images are displayed without any further transformation. While this eases the post production it reduces the possibilities in color correction. The images have been tone-mapped and transformed into the target color space.

The tone-map curve is applied to the Log C data. This transform is also available as a LUT for post-processing of Log C footage. As you can see in the figure below the curve compresses the shadows and highlights of the image. While this provides a nice image on a video monitor, it also means that some information is no longer available in the image.





Figure 2: Tone-map curve (including video gamma correction)

The transform into the target color space brings also a loss of color information with it. Remember, that a display can not produce colors outside of its gamut. Those colors will be mapped or clipped to colors in the gamut.


Differences between SUP 2.x and SUP 3.x

The biggest improvement of the color processing in SUP 3.x is a better handling of those out-of-gamut colors. Their mapping into the target color space is more gentle and smoother. This can be seen in this [HTML](#) document. The images on the left side were generated by SUP 3.x. With increasing exposure the color patches in those images maintain more of their color as in the images on the right side, which were generated by SUP 2.x. The transition into over-exposure happens in a smoother way.

This difference can also be seen in the blurred highlights in the images below, where more color is retained in the image on left side.



The new processing reduces or avoids artifacts coming from out-of-gamut colors. The example below shows some very saturated green colors in foliage illuminated by strong lights. Again, the color is outside of the Rec 709 color space and clipped green spots appear in the image on the right side. The new processing in SUP 3.x avoids those artifacts.



On request of many users we use a softer tone map curve in SUP 3.x and reduced the overall saturation of the images.



DCI P3

The use of DCI P3 as a target color space for images is fairly new. This color space is greater than Rec 709 and is supported by all digital cinema projectors. An increasing number of LCD displays support it as well.

All that has been said about Rec 709 equally applies to DCI P3. It eliminates the needs for a LUT in the display path but it reduces the possibilities in color correction. The images are transformed by the same tone map curve as explained in the discussion of Rec 709. It's just the gamma compensation and the color encoding that is adjusted.


Because the DCI P3 color space has more saturated primary colors than Rec 709, images encoded for DCI P3 look desaturated when displayed on a video monitor.

Differences between SUP 2.x and SUP 3.x

The DCI P3 output option is introduced into the camera with SUP 3.0. We offer 3DLUTs for conversion of Log C footage recorded with SUP 2.x. The 3DLUTs for processing of Log C footage recorded with SUP 3.x are based on the same enhanced processing that is applied in the camera.

Linear Encoded Data

Linear data is mostly used in VFX processing because it is the natural encoding for computer generated elements. The ALEXA camera has a dynamic range of 14 stops which is a linear range of more than 15,000:1. In most cases floating point numbers are used to store this range (the OpenEXR image format, for example, is based on 16 bit floating point numbers).



Linear from ARRIRAW files

The most direct way to linear files is to record ARRIRAW with a T-Link certified recorder and to process the data with the ARRIRAW Converter (ARC). Several third party software vendors also support ARRIRAW in their systems.

Linear from Log C files

Another way to obtain linear data is to undo the Log C curve of images recorded on tape or in QuickTime ProRes 4444 files. ARRI provides LUTs for this purpose.

Differences between SUP 2.x and SUP 3.x

Data obtained from linearizing SUP 2.x images are in native camera RGB. Data obtained from linearizing SUP 3.x images are in wide gamut RGB.

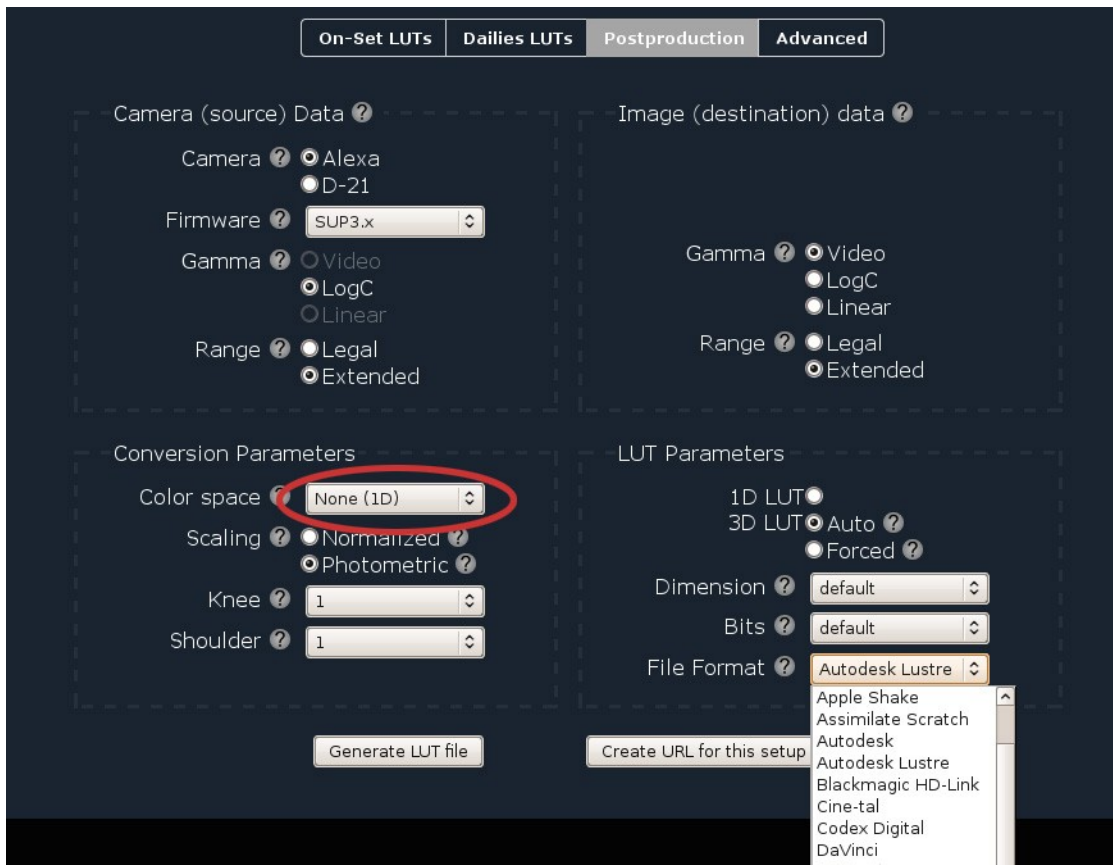
ALEXA Workflow

Log C Images

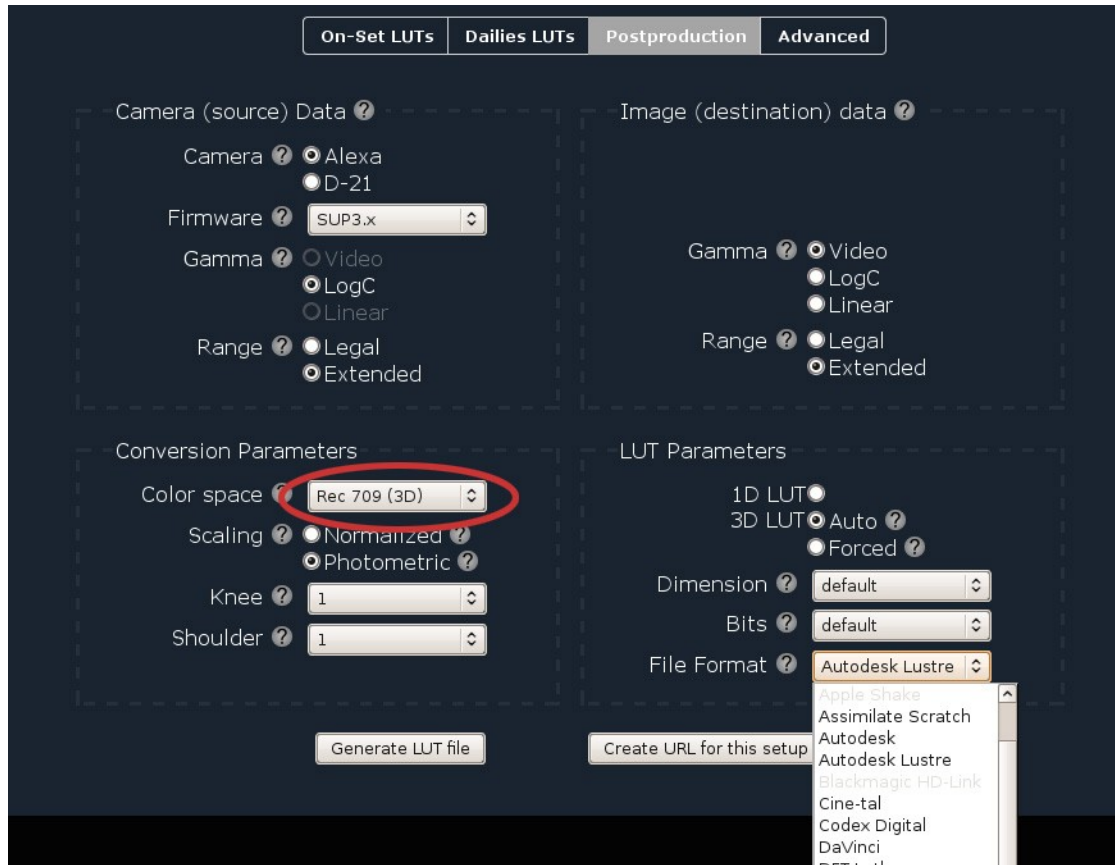
Preview

One can display Log C images on a monitor without any LUT. But you will see a very flat image with desaturated colors. Remember, a Log C image is an encoding of the scene and not an image rendered for display.

Depending on the possibilities of your equipment, you want to apply a LUT or a 3DLUT to the images. The one-dimensional LUT does the tone-mapping described above. The resulting image will at least have a grayscale characteristic more suited for display. The transform into the target color space, however, needs to be done with a 3DLUT. This type of LUT contains both, the grayscale and the color transformation. Both types of lookup tables can be downloaded from www.arridigital.com. The web application can generate one-dimensional and three-dimensional lookup tables in a variety of formats. In the screenshot below *None* is selected as color space and the resulting LUT will include the tone-map curve only.



A 3DLUT is generated when a color space transformation is selected like in the example below. Note that some systems do not support 3DLUTs and are disabled in the format menu.



There is a certain degree of freedom when a Log C image is rendered for display. The situation is similar to choosing a print film. You may prefer a more contrasty or a softer look. You may prefer fully saturated colors or the bleak look of a bleach bypass process. Therefore, using a 3DLUT with a custom look is a perfect alternative to using the LUTs provided by ARRI.

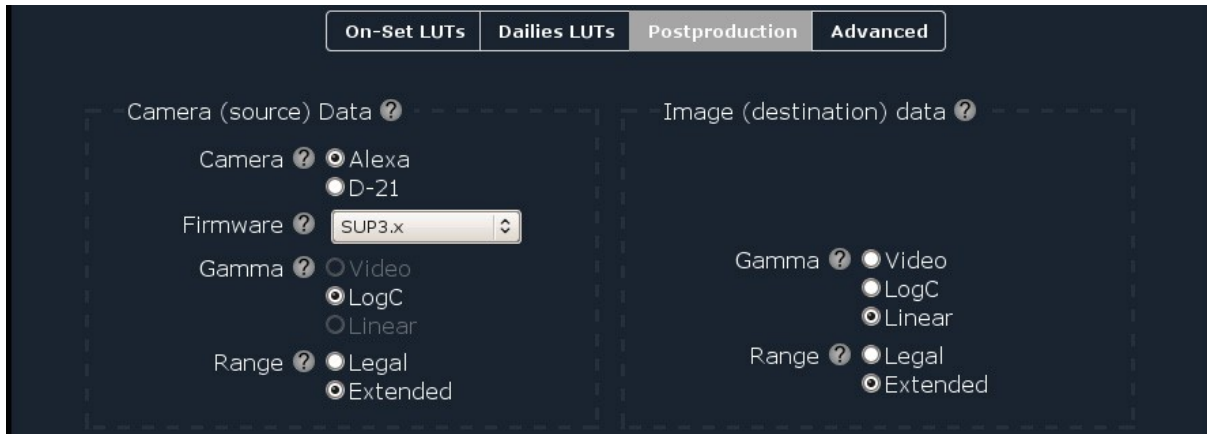
Color Correction

The same transform as used for preview is applied in the display path of the color correction system. All color correction is applied to the Log C data as done in a Digital Intermediate production based on film. For the DCDM (Digital Cinema Distribution Master), the images are rendered through the same 3DLUT into display RGB and further converted into the CIE XYZ color space. For recording to film, a color management system needs to be used that converts the display RGB values to recorder code values needed to produce a print matching the digital images. The ARRILASER comes with a set of conversion tables from standard color spaces like Rec 709 and DCI P3.

Alternatively, a 3DLUT emulating the recording/printing process (PFE) may be used in the display path. This eliminates the need to convert the data for the film recorder. You should be aware, though, that motion picture print film has been designed for color negatives. And color negatives see colors in a different way than a digital camera. To start with, color negatives are not balanced in the sense that the three layers do not have the same gamma (remember, it's the slope of the linear part of the curve relating density to exposure). In most negatives, the layer sensitive to red light has the lowest gamma. Print film compensates for this characteristic by having a higher gamma in the layer building the cyan dye. The camera RGB signals, however, are balanced. Neutral objects will be represented by equal RGB values across the range from black to white. As a result, the colorist may have to compensate for a shift in color balance while changing the overall exposure level. We noticed also that in ALEXA images, when viewed through a PFE, green colors like those of foliage turn yellow.

VFX

For compositing with computer generated elements images are normally converted to a linear encoding. The web application at www.arrydigital.com offers LUTs for this purpose as well.



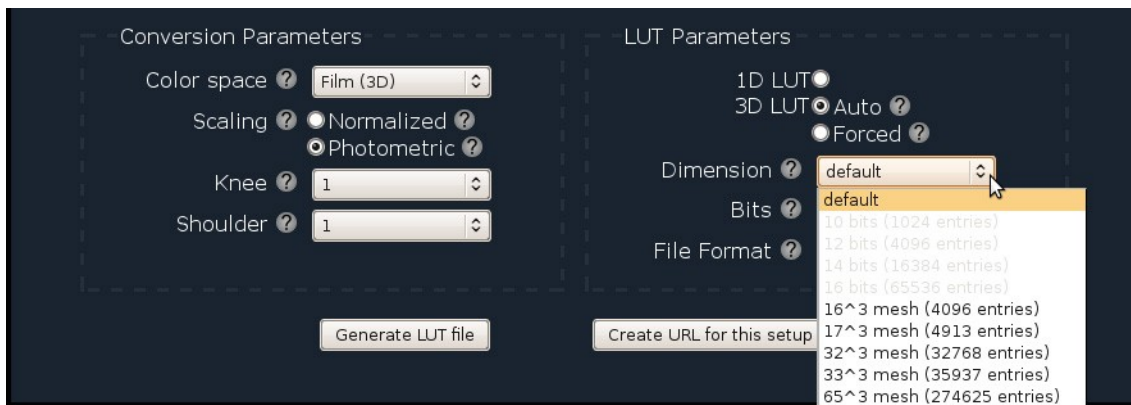
The output of this transform will be linear sensor code values. A conversion to relative linear exposure can be done as following, where we assume that the exposure index was EI 800 and that mid gray shall be represented by a value of 0.18.

1. Convert to linear sensor data using the inverse LogC LUT for EI 800.
2. Subtract the black offset of 256/65535.
3. Multiply data by 36 (0.18 / 0.005).

We offer more details about the Log C to our clients and affiliated post production facilities. Please contact ARRI-CINE - Digital Workflow at digitalworkflow@arri.de.

LUT Size

The bit size of a one-dimensional LUT rules how many distinct input values can be encoded. The preset for all LUT formats uses a sensible default for those formats that support multiple sizes. For special needs, you can change this size. Some formats are only available in a fixed size and they will ignore this setting. A three-dimensional LUT uses less input values in each channel because the total size of the table raises with the cube of the number of input values. Systems that apply a 3DLUT interpolate between the values in the table.



As a consequence of this interpolation, an image may change slightly if the number of mesh points in the 3DLUT is low, say 16^3 or 17^3 .

This [HTML](#) document demonstrates the influence of the size of a 3DLUT. On the right side you see a video image (Rec 709) processed by the ALEXA camera. You can compare this image with several versions of the same image on the left side that have been generated by converting a Log C image with a

3DLUT in different sizes. You will not be able to see any difference between the original and the results of lookup tables with 65^3 or 33^3 mesh points. The image generated by the smallest LUT, however, differs slightly from the original.

The influence of the LUT size is much stronger with Log C images from ALEXA SUP 2.x and the corresponding conversion lookup tables. We strongly recommend to use a 3DLUT with 32^3 or more mesh points for Log C images from ALEXA SUP 2.x.


Normalized and Photometric Video Conversion

The sensor of the ALEXA camera has a dynamic range of 14 stops. The exposure setting determines how many stops will be over and under the mid gray. With an exposure of EI 400, for example, the ALEXA has headroom for $6 \frac{1}{2}$ stops. Exposing at EI 800 increases the headroom by one stop (while you lose one stop in the shadows).

This difference is seen in the Log C output of the ALEXA camera, where the code values represent scene exposure. At EI 400, the maximum code value is 917/1023 and at EI 800, the maximum is 976.

When the Log C image is converted to a video image, the resulting code values are scaled to a fixed range. We call this conversion normalized because the maximum exposure value is always represented by the maximum video code value (1023 in 10 bit).

When the lookup tables are generated one can decide to omit this scaling. We call this photometric conversion.



One advantage of the ALEXA SUP 3.0 is that the difference between normalized and photometric conversion is reduced compared to SUP 2.0. This is shown in this [HTML](#) page. On the left side you see the video image generated by the processing in SUP 3.0 and on the right side the image generated by SUP 2.0. At the bottom of the page you can switch between normalized and photometric conversion. When you do this, you will notice a visible shift in brightness in the images on the right side. Because of the fixed black point of the new Log C curve in SUP 3.0, this shift does not occur in the images on the left side.

QuickTime

QuickTime is a container for image sequences. The ALEXA camera can generate QuickTime movies using the ProRes codec in the following variants.

- ProRes 422 (Proxy)
- ProRes 422 (LT)
- ProRes 422
- ProRes 422 (HQ)
- ProRes 4444

It is possible, and many people have already successfully done so, to write Log C images into the QuickTime.

No matter which type of image is written into the QuickTime, the data is always in legal range. This means that the minimum value is 64 and the maximum value is 940 (those values are used for the 10 bit variants of ProRes, the values are scaled accordingly for the 12 bit ProRes 4444). All known systems, however, will automatically rescale the data from QuickTime to the more customary value range in computer graphics, which goes from zero to the maximum value allowed by the number of bits used in the system (e.g. 255, 1023, or 4095). Therefore, LUTs generated at www.arridigital.com for processing or display of ALEXA QuickTime movies should always be generated for extended range input.