Proposal for the Academy Linear Grade (ALG)

Author: Daniele Siragusano
Email: daniele@filmlight.ltd.uk
Date: 02.03.2020

Current Situation for Colour Decision Interoperability

The current way of exchanging colour decisions (in 2020) is via the ASC CDL (Colour Decision List). The CDL was invented at a time where many different workflows were used. Display-referred and scene-referred workflows both needed a standard grading operator, but using the CDL in modern workflows has some significant drawbacks.

Drawbacks of the ASC CDL

The ASC CDL is not suitable for directly modifying ACES Linear encoded files (both AP0 or AP1) - mainly because of the order of the gain and offset operations and the choice of saturation coefficients.

In ASC CDL the multiplication happens before addition, but in linear-light workflows we need the opposite order - more below.

Also, the ASC CDL is explicitly *working colour space agnostic*. Its design fundamentally requires a perceptual encoded (log or power function) tone curve. Additionally the saturation matrix (using *ITU. BT. 709* coefficients) asks for *working primaries* that are similarly aligned to the Rec.709 colour gamut.

To apply ASC CDL values, a colour space conversion from the linear encoding space to a perceptual working space (typically log encoding) is needed.

Precision

The conversion from linear to log is not free. Precision errors can occur when converting from linear to log and back again. These precision errors can come from a variety of sources:

- Lookup tables (if used to perform the colour space conversion) can elicit interpolation errors.
- Representing logarithmically encoded images in half-float can also cause rounding errors.
- In VFX workflows, a CDL and its inverse are typically applied twice (at the beginning and at the end of compositing). This duplicates the errors explained above.

Complexity

Especially in VFX workflows (which utilise linear-light working spaces), the additional conversion for the CDL application is a potential source for error. Further communication of the log-like working colour space is required. The Academy tried to standardise the log working colour spaces; with considerable effort, it has already come up with two log encodings (ACEScc and ACEScct), but other log encodings are popular too. Then, of course, all of those encoding OETFs can be mixed with a variety of virtual primaries.

> The ASC CDL format has no provision for communicating the working space. So the CDL numbers alone are highly ambiguous.

Unsuitable Operations

Typical production grades can be split into three components:

- Per-shot camera matching
- Look modification (LMT)
- Display rendering (RRT + ODT)

The ASC CDL is not suitable for the second and third components. These are typically accomplished by more complex colour transforms.
Therefore the primary use-case in which one would like to communicate CDL on a per-shot base is *per-shot camera matching*.
The operations include:

- Flare control
- IDT homogenisation for explicit gear used in production
- IDT correction for unusual light conditions
- Exposure control
- White balance
- Slight saturation adjustments

Also, the *per-shot camera matching* should not destroy the scene-linearity.

All of the colour operations above are very hard (or impossible) to achieve with the ASC CDL, and the ASC CDL does not necessarily maintain scene-linearity. In contrast, these operations are trivial to achieve in linear-light colour manipulation.

Requirements for a Modern Per-shot Grading Tool

Let us examine the required per-shot colour operations in greater detail:

Flare Control

Flare has different causes but can be quickly and sufficiently modelled with an addition in linear light.
Pre-flaring images in VFX is also an everyday routine to get rid of negative colours in the noise floor.

Therefore, the ALG needs an addition as the first operation.

IDT Homogenisation for Explicit Gear Used in Production

IDTs provided by camera manufacturers are not ideal because no camera nowadays meets the Luther condition. Depending on the choice of optimisation parameter, different vendors might map the same scene spectra to different XYZ coordinates.
A specific production could introduce a slight correction to the camera manufacturer-provided IDTs to improve the match for the important scene spectra of the show, in the form of a second 3x3 matrix applied to the linear-light data.

IDT Correction for Unusual Light Conditions

Similar to the subsection above, IDTs are calculated for stereotypical light sources. But a production team could calculate a correction for the actual light used on set.
A specific production could introduce a slight correction to the light source component of the provided IDTs in the form of a third 3x3 matrix applied to the linear-light data.

Exposure Control

Typically the exposure of a camera is defined by the contrast ratio of the scene and the available light. So 18% grey can land on different linear light code values (expose to the right).
A specific production could introduce an exposure correction to the linear-light image data in the form of a fourth 3x3 matrix - using the diagonal entries.

White Balance

Due to the use of filters (like ND), linear-light camera data always needs to be white-balanced. This is typically achieved with a colour space conversion into an LMS-ish space (3x3 matrix), then adjusting the three channels via gain factors and then applying a colour space conversion back (3x3 matrix).
The three steps can be concatenated to a single 3x3 matrix.

Slight Saturation Adjustments

Sometimes the saturation of a shot needs to be adjusted to compensate for perceptual factors like the *Hunt effect*, to match the appearance of the shot on a different viewing condition.
There are many (free and proprietary) ways to model saturation, and most can be expressed with a 3x3 matrix.

Academy Linear Grade (ALG)

All of the above 3x3 matrices can be concatenated to a single 3x3 matrix (under normal conditions). So an ideal modern per-shot matching operator can be expressed with an addition in linear light followed by a single 3x3 matrix.

This would give us 10 numbers (one for flare and nine for the 3x3 matrix).

The Academy would not necessarily specify how those values are derived. Each post-production company and grading manufacturer could incorporate their own models and tools to calculate the 10 numbers. Also, the user interface to the 10 numbers could be different from tool to tool. We think this answers exactly the need for both a standard for interoperability and room for differentiation between competitors across all fields.[1]

> It is essential that grades can travel from application to application in an unambiguous way.

Also, the ALG could be applied directly to the linear-light encoded image data, making additional colour space conversions obsolete.

ALG does not destroy scene linearity, so it is ideal for VFX pre-grades.

ALG can only be used in colour-managed workflows like ACES. It is not possible to successfully grade with ALG in display-referred or telecine-style grading styles, so a production company could spot incorrect workflows early on (if the DIT or colourist complains that they cannot grade with ALG - it means that their setup is not right).

Limiting the ALG to essential but powerful operations also makes a clear distinction between per-shot grading and look development. Every higher-order modification like non-linear contrast or colour changes need to go into an LMT in the form of a CLF.[2]
It also helps to clarify the line between a DIT (using ALG per-shot) and a colourist crafting the LMT:

- ALG influences the exposure of a shot.
- LMT influences the lighting, photography and production design of a show.

The Academy should specify the linear-light gamut that the ALG should be applied in, to avoid ambiguity. It does not really restrict the actual working space on how the 10 numbers are derived (except for flare).

1. It should be a general Academy strategy to specify abstraction layers instead of explicit manifestation of processes. ↩
2. The author strongly suggests that CLF should also be defined in the same linear-light encoding space, simplifying the application of ALG and LMT. ↩