




The Academy Color Encoding System – Overview

These days, a film or TV show is unlikely to use only one type of camera for every setup, and indeed different camera types may be used to cover different angles within the same setup. Between filming and final deliverables, the content will pass through a range of software applications, and be handled by a number of creative and technical teams, often in different locations. And it is likely that the deliverables will take multiple forms, each intended for viewing on a particular type of display in a particular environment. Now more than ever, it's vital to implement image processing pipelines that ensure color consistency at every stage.

The **Academy Color Encoding System** is designed to do exactly this. It defines a common working color space into which content from different sources can be transformed, and Output Transforms for all the common display standards, so that the captured images will be perceptually consistent no matter where or at what stage they are viewed. While big facilities with in-house color specialists may have their own custom workflows, an ACES-based pipeline simplifies the implementation of robust color management from filming, through editorial, VFX and DI, all the way to distribution. ACES includes clearly defined standards and best practices, that are implemented in a wide range of software tools, and can be added to many others through the use of plugins.

ACES Color Science

Modern digital cinema cameras — even many lower-end models — are able to record in raw or log modes to capture **scene-referred** images. This means that the captured images have a known (and invertible) mathematical relationship to the light in the photographed scene, as opposed to **display-referred** images, which are designed simply to “look good” on a particular monitor. The output of a sensor is always scene-referred, but traditional video cameras add a display transform to the image, converting it to display-referred before recording. ACES delays the application of the display transform (Output Transform or ODT in ACES terminology) to the end of the production pipeline, which allows not only greater image fidelity, but also more flexibility when it comes to VFX and color grading. Additionally, a variety of Output Transforms can be applied to an individual scene-referred image in order to target different displays, such as Rec. 709, ST 2084 (HDR) or DCI (cinema).



The scene-referred recordings from different cameras use different custom encodings that have been optimized for the capabilities of the particular camera. Log curves fit the dynamic range of the camera into a limited bit depth for recording, and color primaries are chosen to minimize wasted code values. These differences mean that unique display transforms are needed to map each camera's content to a given display. ACES circumvents this issue by defining a color space, ACES2065-1, that encompasses the entire range of visible colors, and an encoding that uses 16-bit floating point values. As a result, captured images from any camera can be transformed to the space using simple mathematical operations called Input Transforms (sometimes referred to as Input Device Transforms or IDTs). From ACES, the same Output Transforms can be used to render the images for different displays, regardless of the camera source. The same is possible for computer-generated imagery (CGI), meaning that all images can be handled in the same way.

Because the ACES2065-1 (AP0) gamut is larger than that of any current or future display, the final output is always a down-conversion, which preserves maximum image quality within the limitations of the output format. VFX and grading are also performed in the ACES2065-1 color space, or the ACEScg, ACEScc and ACEScct “working spaces” designed for those purposes. Transforming between the full ACES2065-1 color space and these working spaces is streamlined and lossless on systems that implement ACES. These working spaces are tailored to respond more intuitively to the math of grading and compositing operations, and they have primaries (AP1) that are more like the red, green and blue colors perceived by the human eye.

ACES-encoded images are viewed through an ACES Output Transform for the particular display in use, but this “rendering” is not baked in until final deliverables. An ACES master is preserved for archiving, and down the road, that master can be reprocessed through a different Output Transform, perhaps for HDR or some future display technology. In many cases, only minimal trim adjustments will be needed to take advantage of new display capabilities. The [ACES workflow diagram](#) shows an overview of this process.

Crafting a Look

ACES is designed to have a light touch on the image. Although a degree of subjectivity is involved in designing an Output Transform, the intent is to produce a neutral base result, giving each user the flexibility to create their preferred look. On set, LUTs can be used, either in-camera or in an external LUT box, that preview the base look in an ACES-based pipeline. These are available for [download](#) at ACESCentral.com. Even without a dedicated ACES LUT, if the default Rec. 709 output from most cameras is used for on-set monitoring, it is similar enough to the look of ACES-based content that surprises are unlikely when ACES is used in post-production (although testing is strongly recommended).

Before principal photography even begins, tools from ACES [Product Partners](#), such as [Prelight](#) from FilmLight and [LiveGrade](#) from Pomfort, enable DITs and cinematographers to work with the same tools that will be used in the final DI to craft a look within the ACES framework. The Look Transform (sometimes referred to as a Look Modification Transform or LMT) is the ACES equivalent of a “Show LUT” and can be exported as a LUT for use on set and as saved grades in a variety of formats, including ASC CDL. This allows the same image to be seen in video village on set, in dailies (whether on an iPad or projected), on the Rec. 709 monitor in editorial, on an sRGB or P3 calibrated monitor in VFX, and as a starting point for the colorist to finesse in the final DI.

Beyond the Image

Beyond color encodings, ACES defines a number of metadata standards to communicate additional information about the content. These include ACESclip, an XML sidecar file that travels with ACES files and carries information such as the IDT used and any LMTs applied. Together with metadata embedded in ACES-compliant EXR files (defined by SMPTE standard ST 2065-4) this allows compliant software to configure the viewing pipeline automatically, ensuring viewing consistency throughout. The functionality has already been implemented in some applications, and software developers are currently adding it to others. Meanwhile, the readily readable nature of XML means that if an ACESclip file is available, the correct settings can be configured manually. Even without ACESclip, only a small amount of information needs to be passed along to ensure that the appropriate Look Transform and any per-clip grade adjustments are applied when viewing ACES-based content. ACES already defines a comprehensive Common LUT Format (CLF) for applying Look Transforms. Work is in progress on an ACES archiving standard, which packages the finished work into an MXF file together with metadata about the display format and mastering environment. This standardization will enable the creation of self-contained masters defined by the open source ACES standards, greatly simplifying any future versioning or remastering.

The [ACES Central](#) forum contains a number of resources to help you get started with your first ACES project, including the [ACES Primer](#), which gives more detail than can be included in this short overview. Questions you have may well already have been asked and answered in the discussions there, and if not, you can start your own threads, to be answered by fellow ACES users and mentors. You can also find the technical [ACES Documentation](#) there.