colorfront

Using the Human Perceptual Model for Multiple Display Mastering

Table of Contents

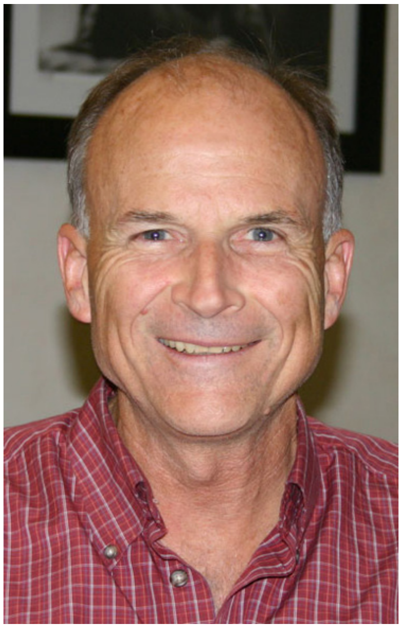

William Feightner, Colorfront CTO

Feightner began his career as technical director of Compact Video, before moving on to similar roles at Laser Edit and Composite Image Systems (CIS). He was co-founder and, most recently, CTO and executive vice-president of technology at EFILM Digital Laboratories (part of Deluxe Entertainment Services Group Inc.) At Laser Edit, Feightner created a live, real-time, multilayer VFX compositing system, and continued this pioneering approach at CIS, where he helped to develop the 2K pin-registered telecine system that revolutionized the process of interactive image compositing for feature films.

At EFILM his innovations included: new software for digital laboratory calibration; image processing and image management software; end-to-end, multi-site, collaborative workflow procedures and software; the world's first fully-digitally timed DI pipeline on We Were Soldiers (2002); and the world's first 4K DI finish on Spiderman 2 (2004). He was also responsible for the workflow on Extremely Loud & Incredibly Close (2011), the first US feature to shoot and post using the ARRIRAW format. During his time at EFILM, Feightner worked closely with Colorfront, harnessing the company's technology on many projects.

Feightner has earned multiple awards. In 2007, he received a technical achievement award from the Academy of Motion Picture Arts and Sciences (AMPAS) for the creation of a Digital Color Separation process for the archival of motion pictures. He also received an Emmy Award for special effects work on the TV series Moonlighting, plus a Monitor Award for the opening segment of The Magical World of Disney. In 2013, SMPTE awarded Feightner the 2013 Technicolor/Herbert T. Kalmus Medal, which recognizes outstanding contributions to the highest standards of quality and innovation in motion picture post-production and distribution services.

Introduction

Maintaining a consistent, creative look across today's rapidly growing field of new display technologies for television, cinema, and portable devices is an immsense challenge. This is especially true when considering the wide disparity in possible brightness, contrast, and viewing environments. The traditional show LUT based approach of a dedicated finish for each possible combination is not practical. Even if that approach were taken, maintaining creative consistency would be, at best, extremely difficult.

By utilizing the knowledge of how the human visual system works, the perceptual model, processing source content to fit a given display's brightness and surround can be applied while maintaining the original creative intent with little to no trimming. This white paper will explain how the Colorfront Engine's Perceptual Processor is currently addressing these real-world use cases.

The Colorfront Engine is a state-of-the-art parametric color processing pipeline mapping various input formats, including camera original (scene-referred) and graded (display-referred) images, to a wide range of SDR and HDR output formats at user definable brightness levels and gamuts while maintaining the creative intent. It is based on the extensive knowledge of how our eyes see light, and uses an internal processing color space where the perceived color and tonal relationships are preserved.

The Deliveries Quandary

When considering current display technologies, we see a great variation in dynamic ranges, color gamut, and the recommended (or realistic) display environment in which they are watched. It is a mistake to think about these displays as either HDR or SDR in a binary way. In reality, there is a range of different dynamic ranges between the two. There is no magic number defining where HDR begins and SDR ends.

As an example, a 48-nit cinema version is lower in peak brightness compared to the 100-nit broadcast deliverable. Does that mean 48-nit cinema is a sub-SDR, if we call 100-nit SDR? Regardless of labeling them both SDR, they are simply displays with different capabilities rendered in different viewing environments. Enhanced dynamic range cinema venues may be only slightly brighter than 100-nit SDR and yet deliver a high dynamic range experience, because it is not just about the peak luminance, it is about the difference between peak luminance and the black level. By achieving very deep black level in a very dark cinema environment, the viewers can have a wonderful experience with the enhanced dynamic range.

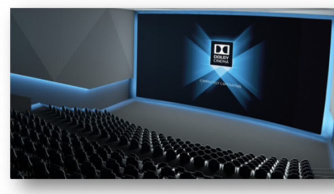

That said, how do we master for all of these versions? Do we have a colorist producing every single one of these deliveries? To avoid having to grade for each display, a managed color workflow is required that translates between these different nit levels and color spaces while maintaining a perceptual match. In other words, when the viewers look at the content, they need to see the same thing.

Preserving The Creative Intent

What do we mean about preserving the creative intent? The creation of images starts with capturing some source footage. The source footage then goes through various transforms and is finally put up on what we call a **reference master display**. This is a very important concept because this is the hero image. This image is exactly what we want all deliveries to look like, no matter whether they are viewed in a brighter or darker environment, or whether the peak brightness matches the current cinema level, or if it is a much brighter HDR experience at 1000 nits or above. The looks need to be maintained across all of these without having to adjust each one of them. Looking at a mastering display and creating a reference image can be done in multiple stages of the workflow, from starting on-set to all of the way through post-production. In our industry, we want the "look" to be accurately maintained and to travel with the project from end to end.

The reference master display is something that has been quantified; we know what the viewing environment is, what the dynamic range is, and what the color primaries are. We want the image to maintain the exact integrity of what the artist saw on the reference display in a particular ambient surround. The reference image is a display-referred image that represents the light being emitted by the screen. It is quantifiable as the viewing environment has been quantified, so that it is what we refer to as our hero image. It is exactly what we want to be conveyed, all of the way down the line.

The number of ways the images will be consumed is going to keep increasing. As we go off to handhelds and tablets, we have all sorts of different conditions for viewing, and we want the viewers to have the same creative experience with each of those viewing options.

Current Industry Practices

In an ideal situation, a reference master display should be a superset of all of the deliverables. In other words, if our maximum deliverable is a 1000 nits, we should be able to visualize what all of the other possible deliverables would be. Oftentimes, however, this is not possible. Instead, we may be limited to the existing viewing technology and may not be able to exceed it. We need systems that can translate from our particular arrangement to a brighter or darker system.

Having a reference master with very high peak brightness has its own challenges, as the mapping to a lower dynamic range display with possibly darker ambient surround viewing condition is not trivial. Without the right tools, it is very easy to diverge from the original creative intent and to produce an unfortunately altered look for the various deliverables. Below we discuss some typical use cases and current approaches that attempt to deal with these problems.

HDR To SDR: Dynamic Range Of Reference Master Larger Than Deliverable

A typical use case scenario is when the reference master's dynamic range is large and graded on a brighter display, while we need a deliverable with a lower dynamic range. For example, we have a master graded on a bright 1000-nit HDR display and we need to deliver it to the cinema with 48 nits.

There are various ways to scale down the image and turn the brightness down, but they all come with the loss of the full range image. While clipping of the highlights produces poor image quality at the clipping points, soft clipping also results in the loss of important tonal and color relationships. Using a complex tone curve could be a better solution, however color hue, saturation, and colorfulness would not match in this case either. Although using these techniques could work fairly well with black and white images, this would not be valid approach because, in most present-day cases, we work with color images.

A technique that is often used is grading down a higher range image by using color correction. A skilled colorist may then have to spend a lot of time scaling down the image to fit it into a lower deliverable format for every single scene, and then even have to use shapes and secondaries to produce a good looking picture. This process comes with wasting a lot of creative time instead of spending the colorist's artistic resources on improving the image and implementing a creative look on the piece.

Consistency must be maintained across the scenes and the different deliverables as well, which is something that also consumes a significant amount of time. With the proper transform, however, preserving the perceived colors and tonal relationships can work very well and this method would guarantee a match.

SDR To HDR: Dynamic Range Of Deliverable Larger Than Reference Master

The opposite scenario is where we master a deliverable at a lower nit level. For example, when we have a 100-nit reference master and need to go up to HDR. One could manually boost the brightness range, but that would only result in a bright SDR image that would not look good.

A very typical use case is retargeting 48-nit cinema masters to be shown on a new emissive cinema wall. As it is very time consuming to remaster movies when using current industry practices, most films currently playing on these new screens, unfortunately, do not take full advantage of the capabilities of the higher dynamic range of these displays.

To manually grade the image brighter in order to get a pleasing picture can be very time consuming and laborious. Another technique is to only expand the highlights, however, the relationship between the colors and gray scales would still not match and the creative intent would be lost. We want to be able to maintain all important color relationships and have an image that has the sense of the original one. When we do this in the perceptual processing environment, all characteristics fall into place and work very well.

Transmission Color Spaces

Once we come up with the master look and have it rendered to the desired nit level, the delivery is fed into transmission systems, in other words, handshake spaces. For HDR, the main transmission systems are HLG and HDR10. When properly implemented, HLG and HDR10 should reproduce the same colors and look exactly the same. Transmission systems and handshake spaces should not alter the look in any way.

In the situation that we need to master a delivery with a wider color palette or we have colors going out beyond what the display can do, it is very important being able to do proper out-of-gamut remapping. If we only clipped the colors and hues shifted at the color boundary, the image would be bright and oversaturated. Therefore, it is very important that we do a proper, perceptually processed, out-of-gamut remapping, in other words, constraining colors back into a narrower color space.

Previewing

Soft proofing or soft previewing is essential. We need to be able to preview the final grading on a higher dynamic range display in addition to all of the deliverables of different nit levels. Nonetheless, the monitor settings should not be changed: we should be able to just flip a button in the system to preview all of the deliverables and make certain our gradings are holding up. We should have the ability to soft proof all of the various dynamic ranges and color constraints (for example, a rec.709 at 100 nits or 300 nits for a cinema wall) on the reference master display.

Creative Trimming

As discussed above, using the typical creative tools, such as lift, gain, and gamma to create various SDR and HDR versions from a graded reference master is not a good practice. These tools are not designed to maintain color perception and thus require advanced skills and a significant amount of time for desirable results.

There is no question that even when using the best color processing there are some rare cases where manual adjustments (i.e. “trims”) are needed to create a proper final image. When the target dynamic range is significantly different from the reference mastering display’s properties, the tonal relationships in the perceptually remapped image may not hold anymore, aesthetically. When expanding from an SDR source to HDR, a bright highlight in the background may draw the viewers’ attention from the hero, and similarly an SDR image may look too contrasty in a dark surround environment when derived from an HDR source.

In order to fix these issues on a scene-by-scene basis, the proper tools are required. Normal color correction tools are designed for creative painting, not for matching to different display nit levels and viewing surrounds. The appropriate tools for such a job could be:

- *Dynamic range control* to drive the dynamic compression or expansion within the output nit range
- *Ambient surround compensation* from dark to bright and vica versa
- *Colorfulness adjustment* to control the perceived look of the remapped colors

Current Industry Output Transforms

In our industry, current output transforms have not been designed for a perceptual match between the different deliveries. Each transform has its own look, which requires a separate grade for each delivery. Even 6 Using the Human Perceptual Model for Multiple Display Mastering with devoting a lot of time to each delivery’s ODT, it is very difficult to maintain the creative intent across all of the different deliveries.

One use case, for example, is when the image is rendered with Rec.709, which is one of the standard ODTs. In this scenario, in post-production the image is rendered into cinema DCI P3 while keeping the brightness the same. Without perceptually matching the images, there will be a color temperature and saturation difference, which the colorist would then have to adjust.

Another use case is when the master is rendered at 100 nits for dim surround and we want to deliver it at 48 nits for dark surround by using one of the existing standard output transforms. The delivery will be fine, but it will not truly match. However, if we apply a perceptual space and do the same render, the 48-nit delivery will get more brightness and be a much closer match to the original master.

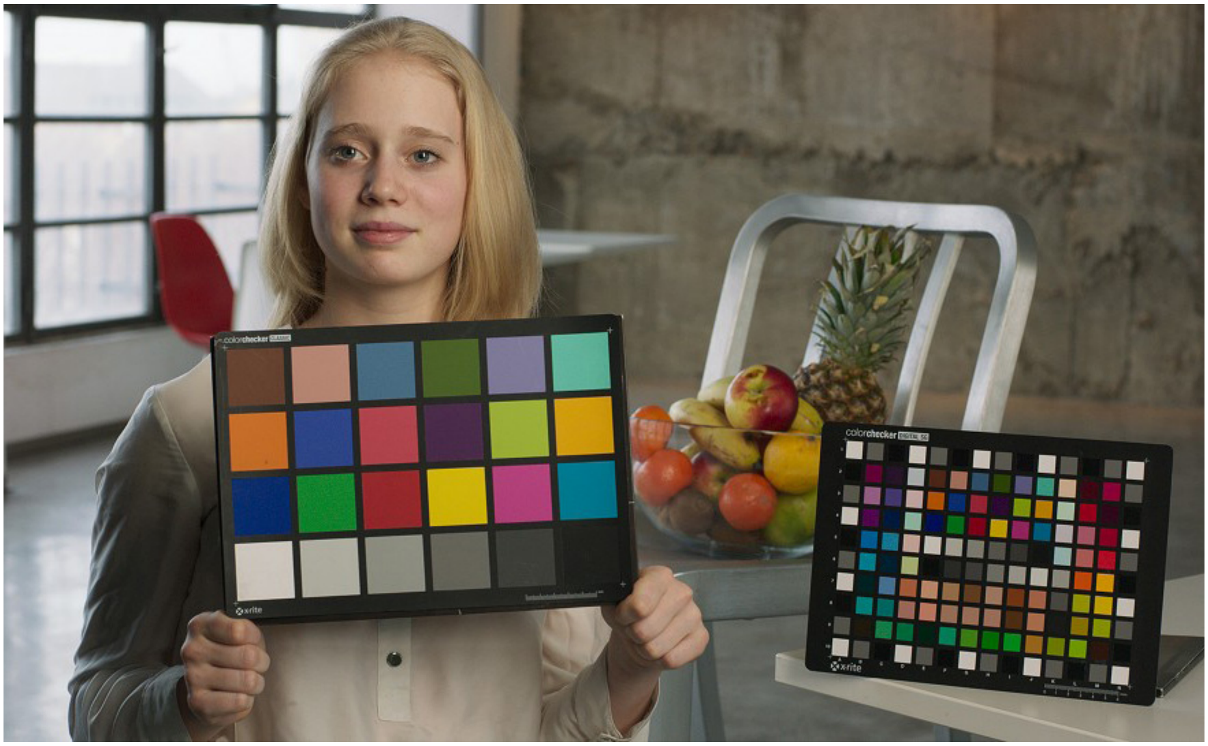

Mismatching 709 and DCI Cinema ODTs

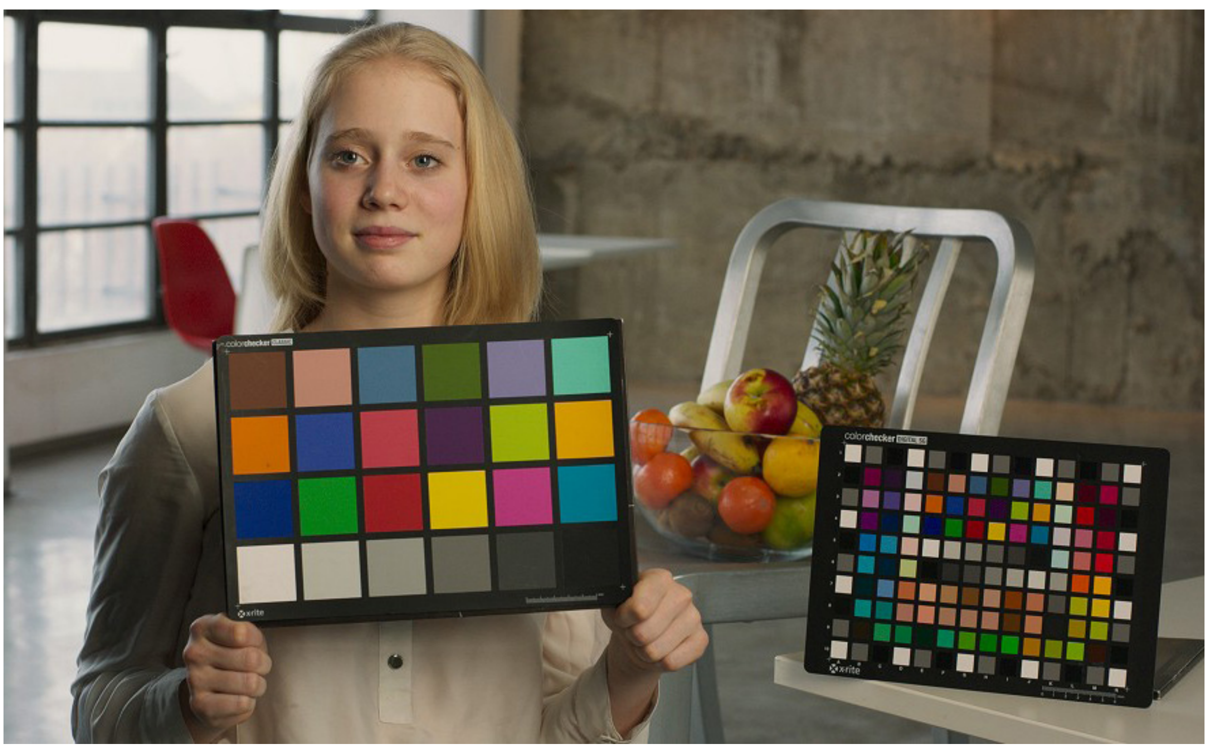

The image on the top, generated with the Rec.709 ODT, does not match the image below, generated with the DCI ODT and converted back to Rec.709 to be comparable. The second image has an overall warmer color tone, which is very visible on the skin tones.

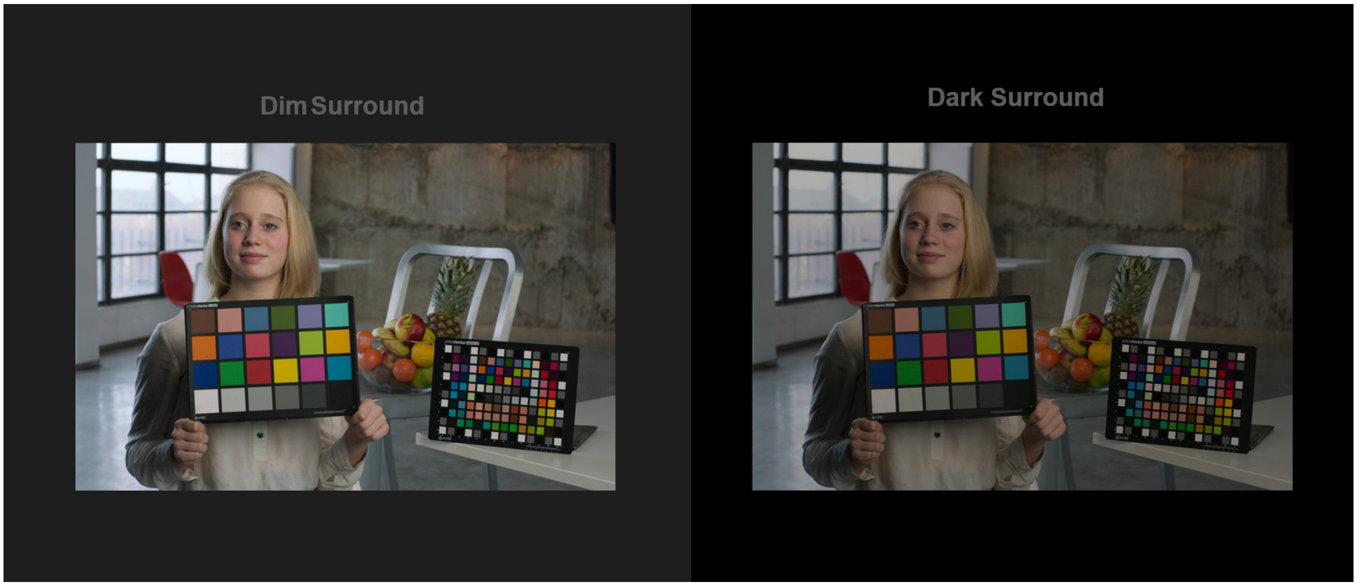

Cinema ODT: 100-nit master in dim surround and 48-nit deliverable in dark surround

When remapping a Rec.709 100-nit image to a darker, 48-nit version without considering how the human visual system works, we get a result that looks slightly different than the original. Skin tones on the left and on the right do not look the same. Please note that these differences are much more pronounced in a proper viewing environment and may be difficult to distinguish in the small images embedded here.

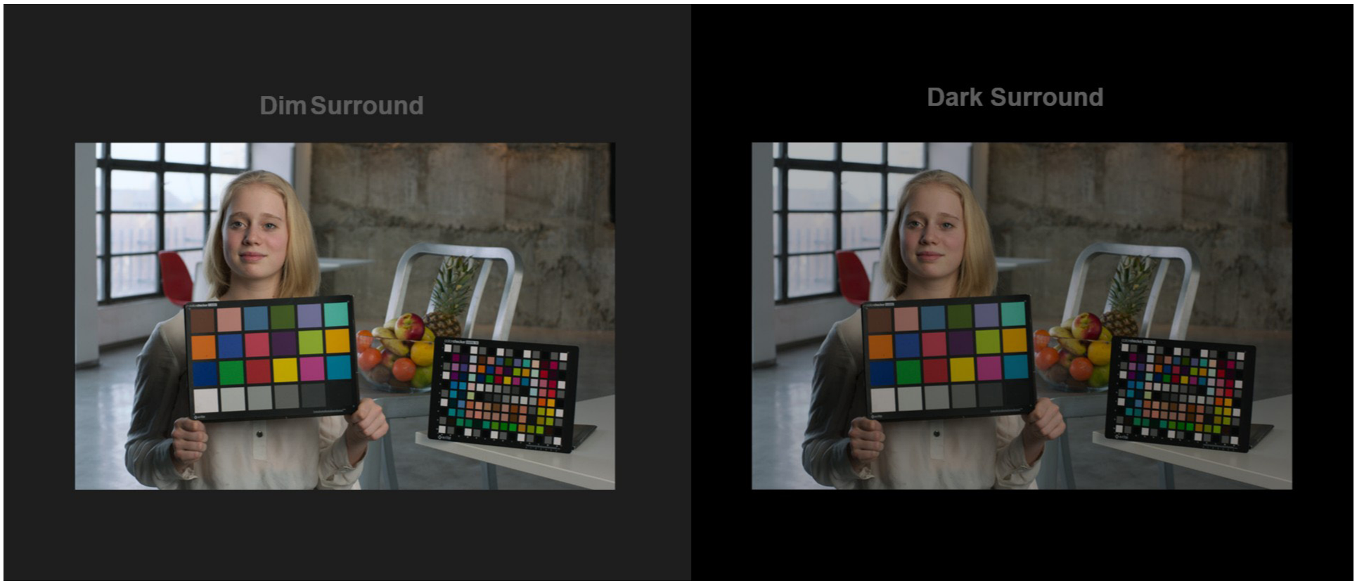

Cinema ODT: 100-nit master and perceptually matched 48-nit deliverable

When using a perceptual transform to make a darker version of the image, skin tones track much more accurately. This is even more important when converting a master to a significantly brighter target.

Perceptual Processing In The Colorfront Engine

Delivering to different target displays of specific brightness, contrast, and color representation capabilities requires a brightness adjustment of the master. If this brightness change is applied in an arbitrary processing color space, the percieved colors will be distorted. However, if these adjustments are applied in a carefully constructed perceptual processing color space, the color perception is maintained, thus the creative intent is preserved.

One practical implementation of such a processing system is built into The Colorfront Engine. This custom processing color space has been designed on the basis of the extensive knowledge of human color perception available in scientific literature, as well as from the many years of practical experiences in high-end feature film look design and color management. Below are listed the particular effects regarding the human color visual perception to be considered when building such a system.

Visual Perception Phenomena

Hue Appearance

Bezold-Brücke hue shift

The hue of monochromatic light changes with luminance.

Abney effect

The hue of monochromatic light changes with the addition of white light.

Contrast Appearance

Stevens effect

Contrast increases with luminance

Bartleson-Breneman effect

Display image contrast increases with the luminance of surround lighting.

Left and right squares have the same pixel values

In the first picture, the left and right squares have the same pixel level although they look different. In the second picture, the right squares are perceptually corrected to look the same as the left squares, but their pixel values now differ.

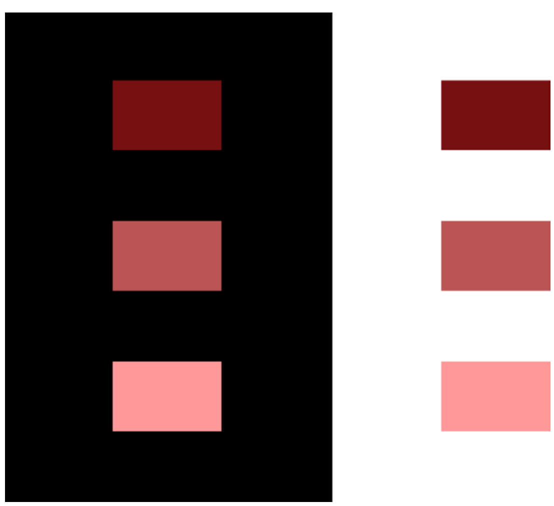

Right squares are corrected to appear the same

Colorfulness Appearance

Hunt effect

Colorfulness increases with luminance. Similarly, objects at dusk appear less vivid than on a bright summer afternoon. This property of color perception cannot be ignored when mapping an image to a brightness level not matching the reference mastering display.

Hunt effect - colorfulness not compensated

Hunt effect - colorfulness compensated

The Hunt effect is when a picture's increase in brightness results in it becoming more colorful. In this case, images have to be compensated to look as they do on a darker reference monitor. Once properly compensated, the colorfulness of the face matches.

Brightness Appearance

Helmholtz-Kohlrausch effect

Brightness increases with saturation.

Chromatic Adaptation

Human color perception adapts to the white point of the illuminating light source when viewing a reflective object.

Perceptual Space

When we put images into a perceptual space, the content is fit into the given display's brightness and gamut capabilities, and they all correspond to each other. Because this is not a linear space, on a scope you might see that a bright image changes to a certain color after its hue shifts. This preceived change is due to the perceptual processing model, which is the way the eye works. The perceptual transform space takes into account the non-linearities of the eye so that the viewing results match.

The ambient surround that we are mastering for also influences the rendering of the image. The final image on the given display needs to match the reference master. The perceptual transform will faithfully reproduce images on different displays in various viewing conditions, while the look of the image is maintained across different display brightness levels and color primaries, and will remain consistent with the varied levels of surround brightness.

Single Master Process

In a single master process, the image comes through an IDT (Input Device Transform) to our common grading space. The key to this system is the output transform that processes our reference look to any other different output levels required. We can also tune the brightness and surround parameters to compensate for dark or bright environment. An inverse transform is also available here in a legacy delivery so that we can invert already pre-existing Rec.709 or other HDR materials back into the grading space, as these transforms are all fully invertible.

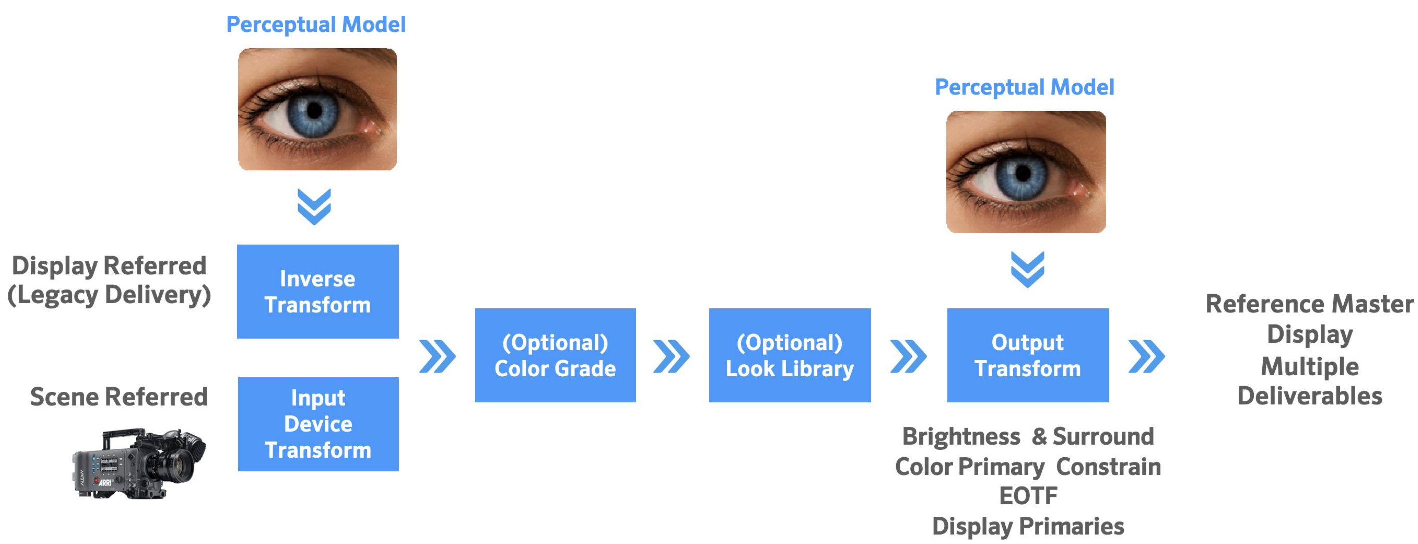

Processing blocks in the Colorfront Engine pipeline

The Colorfront Engine In Practice

Live Broadcast And On-Set Look Management With FS-HDR

FS-HDR, a 1RU, rackmount, universal converter/frame synchronizer, is designed specifically to meet the HighvDynamic Range (HDR) and Wide Color Gamut (WCG) needs of broadcast, OTT, production, post, and livevevent AV environments where real time, low latency processing and color fidelity is required for 4K/UltraHDvand 2K/HD workflows. Developed in partnership with Colorfront, the FS-HDR's HDR/WCG functionality is powered by the Colorfront Engine™ proprietary video processing algorithms.

AJA FS-HDR unit implementing Colorfront Engine to be used on-set or in live broadcast

HDR Dailies In Express Dailies And On-Set Dailies

Express Dailies provides the option for dailies workflows to utilize the Colorfront Engine project type. This allows dailies operators to use ACES compliant IDTs to map all input sourced into the internal processing color space, where dedicated tools are offered to adjust exposure, color temperature, tint and physically correct saturation. From the project, SDR 709, HLG or HDR10 deliverables can be rendered simultaneously utilizing the perceptual processing engine, thus making tandem SDR and HDR dailies effortless.

On-Set Dailies, with the customizable node-based pipeline, offers even more flexibility to design color pipelines delivering any kind of output in a parallel fashion. By using the "Color Processing Wizard" in the node page, the user can create the Colorfront Engine pipeline for any camera format, combined with a color correction tool of preference.

Simultaneous SDR and HDR dailies deliverables in Express Dailies using a Colorfront Engine project

Mastering And Re-Mastering In Transkoder

A wide range of analysis and processing tools are available in Colorfront Transkoder to support single master workflows. The HDR/SDR Remap tool, in conjuction with the floating point precision color space conversion tool, can be used to remap any kind of display or scene-referred source to any target display. Dedicated controls are provided to set the output brightness, ambient surround compensation, HDR expansion amount and colorfullness. The remap can also be broken out into an input transformation (reverse tonemap) and an output transform, so custom grading tools can be added in the middle. This way a pipeline can be built for remastering workflows, where a colorist can utilize creative grading tools to fine-tune the few shots where the automatic mapping does not produce the desired aesthetic result.

The integrated second head analyzer can be configured to work in any kind of color space so the mastering operator can inspect brightness and gamut levels in the input or the output. In a time when new projectors and immersive displays are announced more and more frequently, it is very important to be able to master in accordance with the capabilities of any specific display in today's modern mastering workflows.

Perceptual color volume remapping tool in Transkoder

Colorfront Engine SDK

The modular color processing functions underneath all of the above implementations are available in a software development-kit (SDK) for partners to implement. By applying a relevant subset of functions, various solutions can be designed for specific color processing problems.

Input Formats

The Colorfront Engine SDK supports camera original log files of all major manufacturers, as well as the latest mezzanine, broadcast and cinema formats. The input images are mapped, and, if applicable, reverse tonemapped by the system into the common internal processing color space.

Perceptual Processing

The internal perceptual processing engine offers physically accurate adjustments of exposure, saturation, color temperature and tint. High-quality interal processing of floating points is non-destructive and handles extreme colors near or above the gamut boundaries.

Look Module

Scene-referred inputs are mapped to the final output by applying a user configurable look. A library of looks includes motion picture film looks, high quality broadcast and the latest Hollywood cinematic looks. The Colorfront master look has been specifically designed to produce a modern, colorful output with perfect skin tones, neutral grey axis, and proper handling of out-of-gamut colors. For broadcast applications such as for sporting events, a less dramatic version is available as well. These looks can also be mixed with each other to produce the ideal output.

Output

When generating the final output, both the container color space and the display capabilities can be defined, such as variable nit levels to adapt to any peak brightness and flexible gamut constraints (Rec.2020, P3 or Rec.709). Compensation for ambient surround brightness is also available. The color and relative contrast perception tracks perfectly across different outputs.

The Colorfront Engine SDK is available as a simple-to-integrate lightweight C++ SDK on Windows, OSX, and Linux. Depending on the intended use, the host application may employ a limited set of parameters or the full range of controls. In order to minimize processing cycles, the highly optimized 32-bit floating point processing engine can be easily multi-threaded.

Conclusion

There are many requirements and challenges of applying a color processing pipeline when mastering for multiple displays and various surround brightness levels. The following questions should be carefully considered when choosing a color processing pipeline for a mastering/delivery workflow:

- Does it support the common master workflow?
- Does it handle both SDR to HDR, and HDR to SDR?
- Does it support camera original and graded sources?
- Is it based on LUTs created with creative grading tools?
- Do they break with images pushing the color boundaries?
- Does it support various input and output nit levels?
- Does it support different output color spaces with gamut constraints?
- Does it support various ambient surround conditions?
- Will SDR look the same as HDR? Is the look of the image maintained?

Adopting a human perceptual model to ensure that the creative intent is maintained is an effective and the only adequate method that can handle the wide range of gamuts and brightness levels while preserving a consistent look across all of the abundant number of display targets in today's growing technological world.

The Colorfront Engine is an implementation of a perceptual color processing pipeline that meets the requirements listed above. This technology is currently available in different software and hardware products such as AJA's FS-HDR and HDR Image Analyzer, as well as in all Colorfront dailies and mastering products. Colorfront's state-of-the-art color processing Engine brings true plug-and-play simplicity, as well as a convenient and supreme solution to today's complex multi-source, multi-deliverable production needs.