photo by Sara Terry

HDR Explained

by Josh Limor

High dynamic range (HDR) has generated a tremendous amount of media coverage and analysis over the past year, with the more breathless headlines focused on the implications of having multiple formats that may confuse the creative community.

To be clear, studios and content distributors are excited about the commercial prospects associated with HDR and are asking the creative community – including cinematographers – to incorporate this rapidly rising image technology into their projects. This raises interesting questions about whether or not there are truly creative benefits to be gleaned by artists from HDR. This is an especially cogent question when you consider that the hype around technologies like 3D and 4K seemed to add significant complexity to the creative process without delivering a commensurate improvement in the audience experience.

My view is that HDR does offer creative opportunities that will be well received by consumers interested in cinematic, episodic and live programming. I work with artists every day who are using the expanded spectrum of light and colors that HDR provides, helping to accentuate story elements and heighten the overall experience for the audience.

At its core, HDR is not a complicated proposition. It simply offers content creators access to a much wider palette with which to work, and introduces new opportunities to inject nuances associated with contrast management that were not previously available. While this is exciting, this does not materially alter the tools and skills that cinematographers and colorists have been using for years now.

So what exactly does the creative community need to know to begin getting the most creativity out of HDR?

From a cinematographer's perspective, HDR represents an opportunity to revisit some important concepts that we learned from capturing content on film. Film acquisition is predicated on capturing the widest latitude and highest dynamic range of light on the negative in order to ensure that the post-production team – including colorists – has access to the fullest range of information. In digital acquisition, the tendency for many content creators drifts toward trusting what is seen on the reference monitor as the final image. For years now, digital cameras have been able to acquire much more dynamic range than what we have been able to see on available monitors – which we might retroactively refer to as today as *standard dynamic range* monitors. The emergence of HDR monitors that are not only available in production environments but are starting to penetrate the home allows creatives to take advantage of color and luminance information so that new dimensions can be added to the story.

The risk of working with more dynamic range is that it can allow more noise to be apparent in the image. As a result, cinematographers should ensure that steps are taken to protect the shadows from any undesirable, encroaching noise. Similarly, protecting details in the highlights – or brighter areas of the image – ensures that colorists can now create more depth in the picture they are painting.

One of the biggest advantages of HDR is that artists are able to work with this depth to develop the illusion that people are no longer looking at a flat screen, but are gazing through a window. That said, even though a wider palette of light and color is available with HDR, not every scene and shot demands it. It is always important to remain true to the story and use the additional range only where it makes sense. While a specular highlight off a metal watch might make sense in an advertisement, a bright practical light source may be distracting in a drama.

Once the content has been created, there has been a lot of discussion about how to get HDR images into people's homes. The studios, technology companies, content distributors and consumer electronics manufacturers have come together to develop open specifications around HDR. The market, however, is still wrestling with the fact that multiple formats for content currently exist.

Here's a quick breakdown of the variation in formats:

SDR

The rise of HDR has created a category that requires a name for the existing way images are displayed, standard dynamic range, which is also known as BT.1886, Gamma, Rec. 709. (Gamma is a relative transfer function that stretches or compresses code values evenly across the display from its peak luminance to its minimum luminance.) This mouthful of jargon is what has allowed consumers to experience interoperable high-definition images using the pre-HDR range of luminance that the content ecosystem has had at its disposal. As we move beyond SDR, the entire industry has agreed that the color space will expand from Rec. 709 to BT.2020, which provides a wider color gamut. Where things remain unresolved is regarding what the best replacement should be for BT.1886 Gamma – which defines the luminance curve. As we discuss HDR, some of the key differences between the formats revolve around the characteristics of their luminance curves.

PQ (Perceptual Quantizer)

This curve is the foundation upon which HDR10 and Dolby Vision are built. It is what the technical community refers to as an "absolute transfer function" for luminance. Each code value corresponds to a specific measurement of luminance (typical luminance is measured in nits, or cd/m2) from the screen. PQ spaces the 1024 code values, in a 10-bit signal, up to 10,000 nits. As a frame of reference, today's consumer SDR TVs display peak luminance of 100–350 nits, while OLED HDR TVs display peak luminance of 540 nits and higher; LCD HDR TVs support peak luminance of more than 1,000 nits. Typically, PQ is used for images that want to maintain as much of the original light characteristic that can be handled by a given display. So if you have an image that was graded to 1,000 nits presented on a consumer display that supports 600 nits, then working with PQ may maintain the image accurately to 540 nits and then roll off (or compress) the additional highlight detail into the remaining 60 nits. While PQ was designed as a perceptual distribution of light, it was not optimized for compressed home delivery. This means it can potentially limit the image quality by the time it reaches the home because most current encoding technologies – which compress images – were designed before PQ existed.

HDR10 and Dolby

The only difference between HDR10 and Dolby Vision is that HDR10 relies solely on static metadata describing the entire piece of content and hardware inside the TV to map and roll off the highlights. By contrast, Dolby's HDR technology relies on a combination of metadata from the content-creation process and hardware inside a Dolby Vision TV. Dolby's metadata is created with a specific additional trim pass on top of the HDR10 version and requires Dolby's proprietary hardware.

HLG (Hybrid Log-Gamma)

This curve, which has been supported by the BBC and NHK, is a "relative transfer function" that is designed for HDR content. This ostensibly allows content to be sent without the need for metadata. (This point is disputed by some critics.) HLG is designed to allow for a greater use of code values with the assumption that most images may only be graded up to 1,000 or 2,000 nits. This assumption allows code values to be better distributed throughout the bits available in the signal.

S-Log3

This is a proprietary curve designed by Sony for acquisition in professional environments – not for consumer distribution and display. It maximizes the dynamic range potential of their camera sensors while limiting the noise.

The key thing to keep in mind is that there have been specific market conditions that have led to the development of each of these formats. In other words, there is a constituency behind all of these approaches to HDR. The challenge is to figure out how to manage this complexity for content creators and consumers. It turns out that the math exists to create algorithms that can harmonize and normalize this array of formats by creating a common baseline from which all content can be recreated. As a result, it is possible to build a universal distribution platform that can carry any piece of content (regardless of format, including SDR) from any source to any display.

An example of such a platform is Advanced HDR by my firm, Technicolor, which was designed with compression quality in mind to allow artists to create any format of content to tell their story. The platform ensures that HDR content can travel over any distribution vehicle, and then be consumed by audiences over any device or display. ICG

Josh Limor is Vice President, Technology and Ecosystem Development, Technicolor, in charge of developing the portfolio of next-generation immersive video solutions. He has spent over a decade in business and product development for the video industry, and has a track record for defining market opportunities and creating successful product portfolios. Prior to Technicolor, Limor worked at Panavision, where he designed, developed and implemented a range of digital workflow solutions that helped to transition the market from film, to digital, to higher dynamic range.