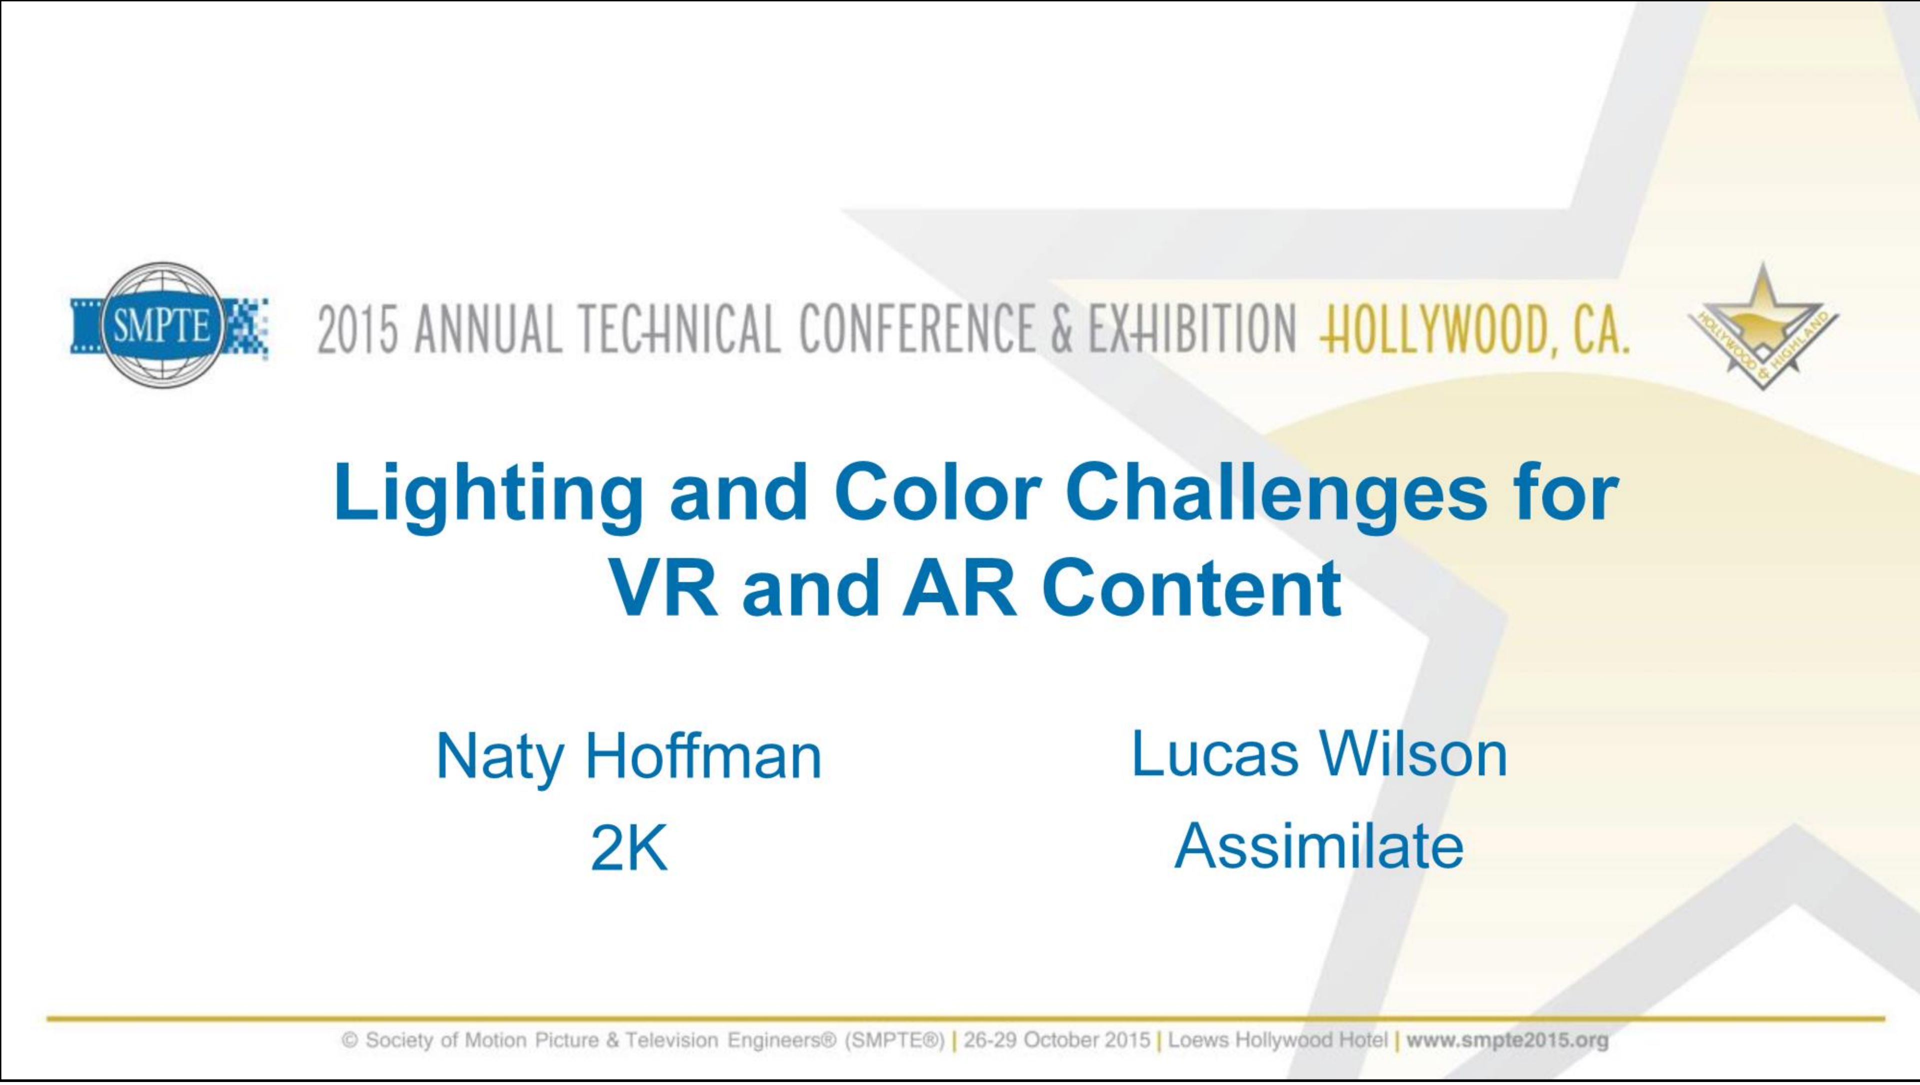

Hi, and thanks for coming to our talk this afternoon. I'll be presenting the first two parts of the talk: general background and videogame-specific considerations. Lucas will present the third part, which will cover considerations for live-action VR and AR.

Background: Color and Lighting

Naty Hoffman, 2K

We'll start with some fundamentals regarding color and lighting.

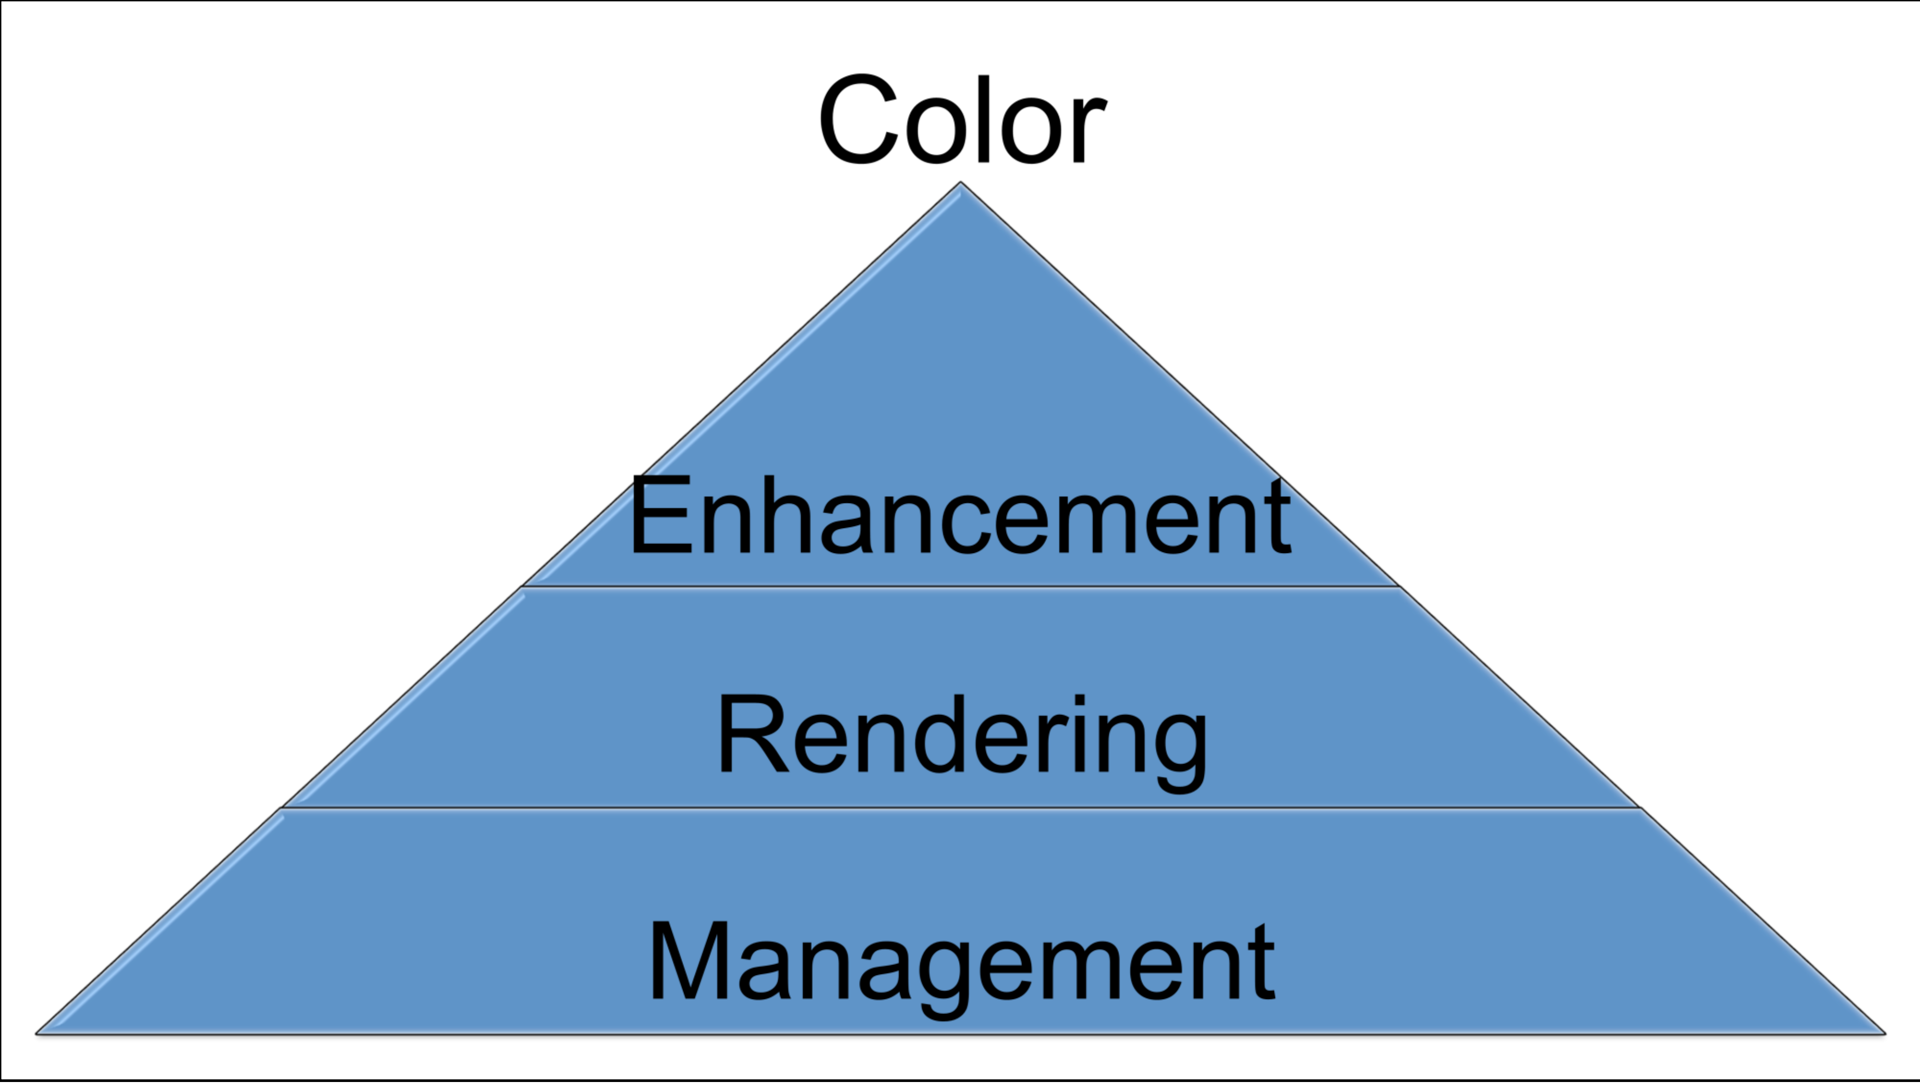

The different aspects of color form a conceptual pyramid, with each layer relying and building on the layer below it. Color management is at the base: precise definition of color spaces, viewing conditions, image states, reference standards etc. On top of that, we have color rendering (not to be confused with computer graphics rendering)—the "scene to screen" transform which is essential for correct color reproduction. On top of that we have color enhancement, which covers creative modifications such as color grading.

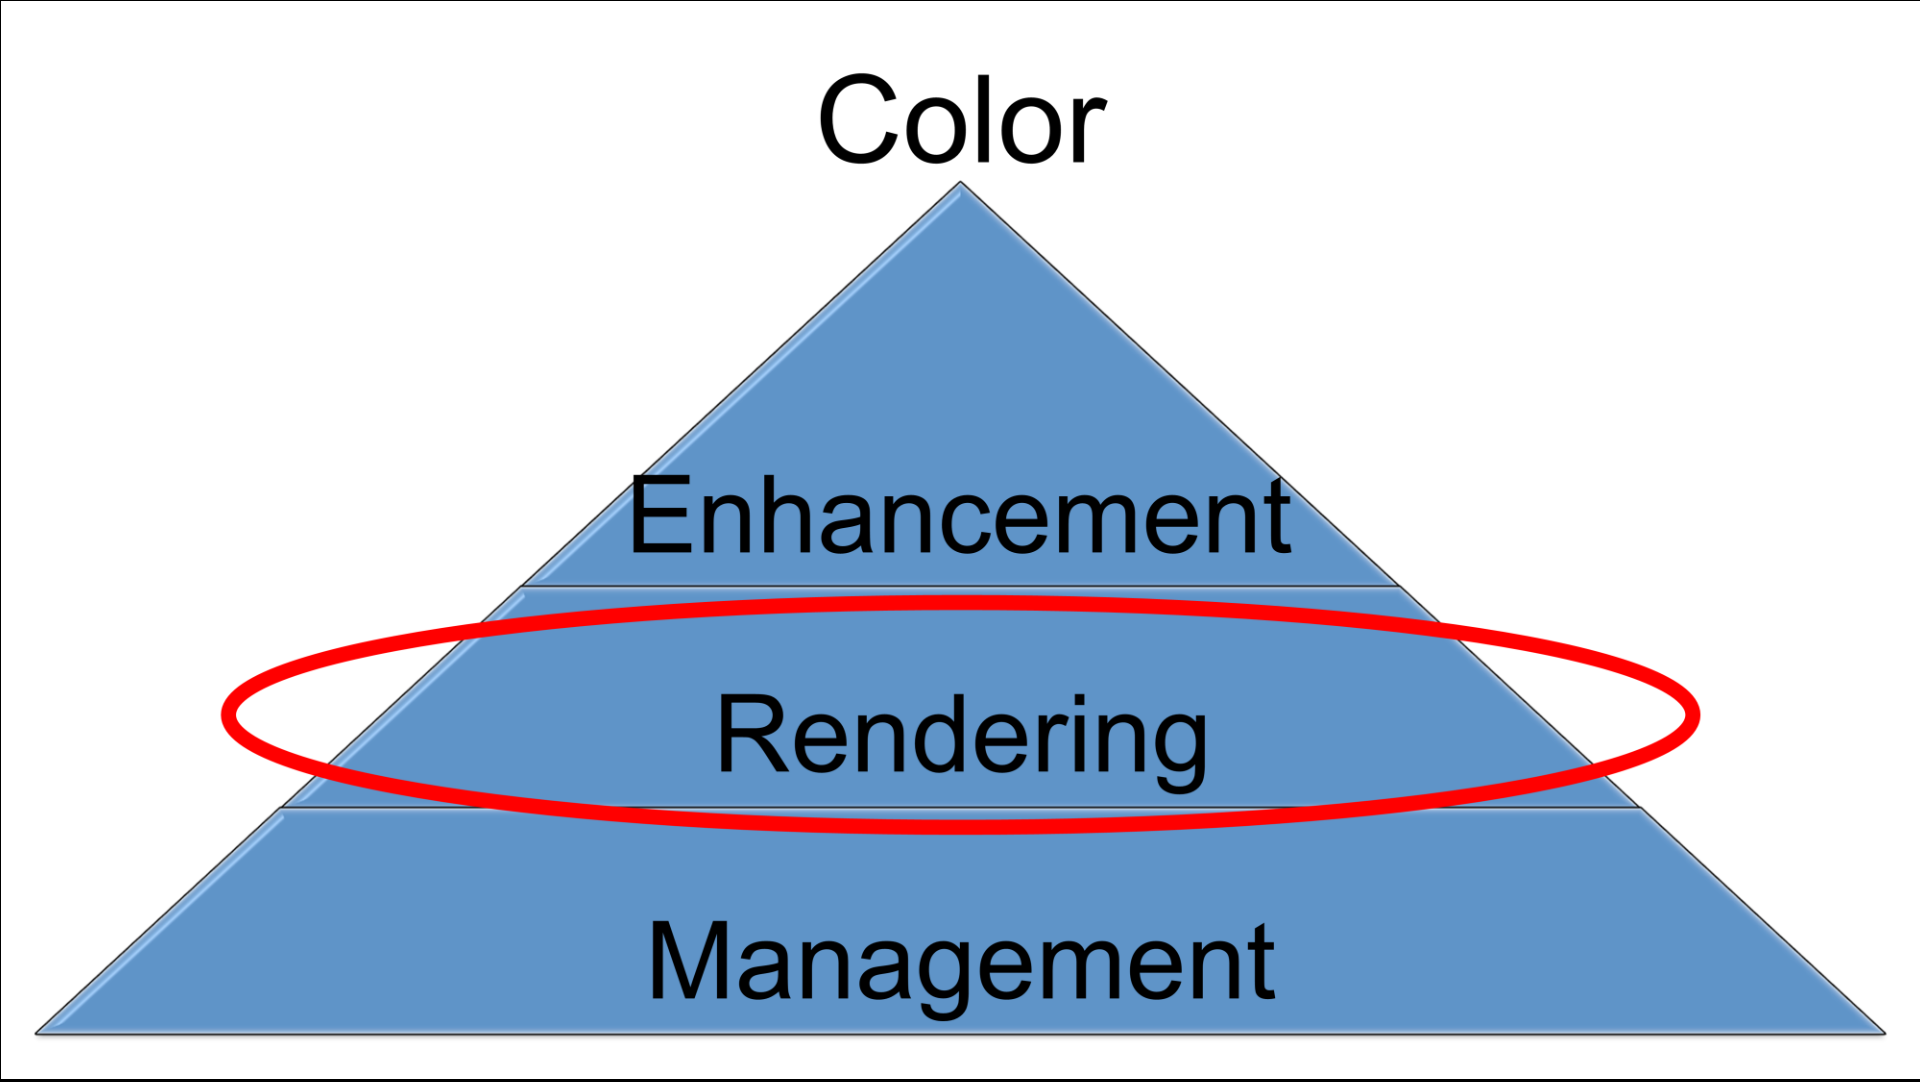

I’ll discuss color rendering in a bit more depth, since of the three, this is the layer that is likely to change most when going from traditional imaging to VR and AR.

The Image Reproduction Problem

Color rendering is a solution to the image reproduction problem, the core challenge of imaging science.

One way to state this problem is to start with the appearance of an actual scene (for example this sunlit meadow). We want to reproduce the appearance of this scene onto a display device, which can be a…

…computer monitor in an office,…

…a television set in a living room,…

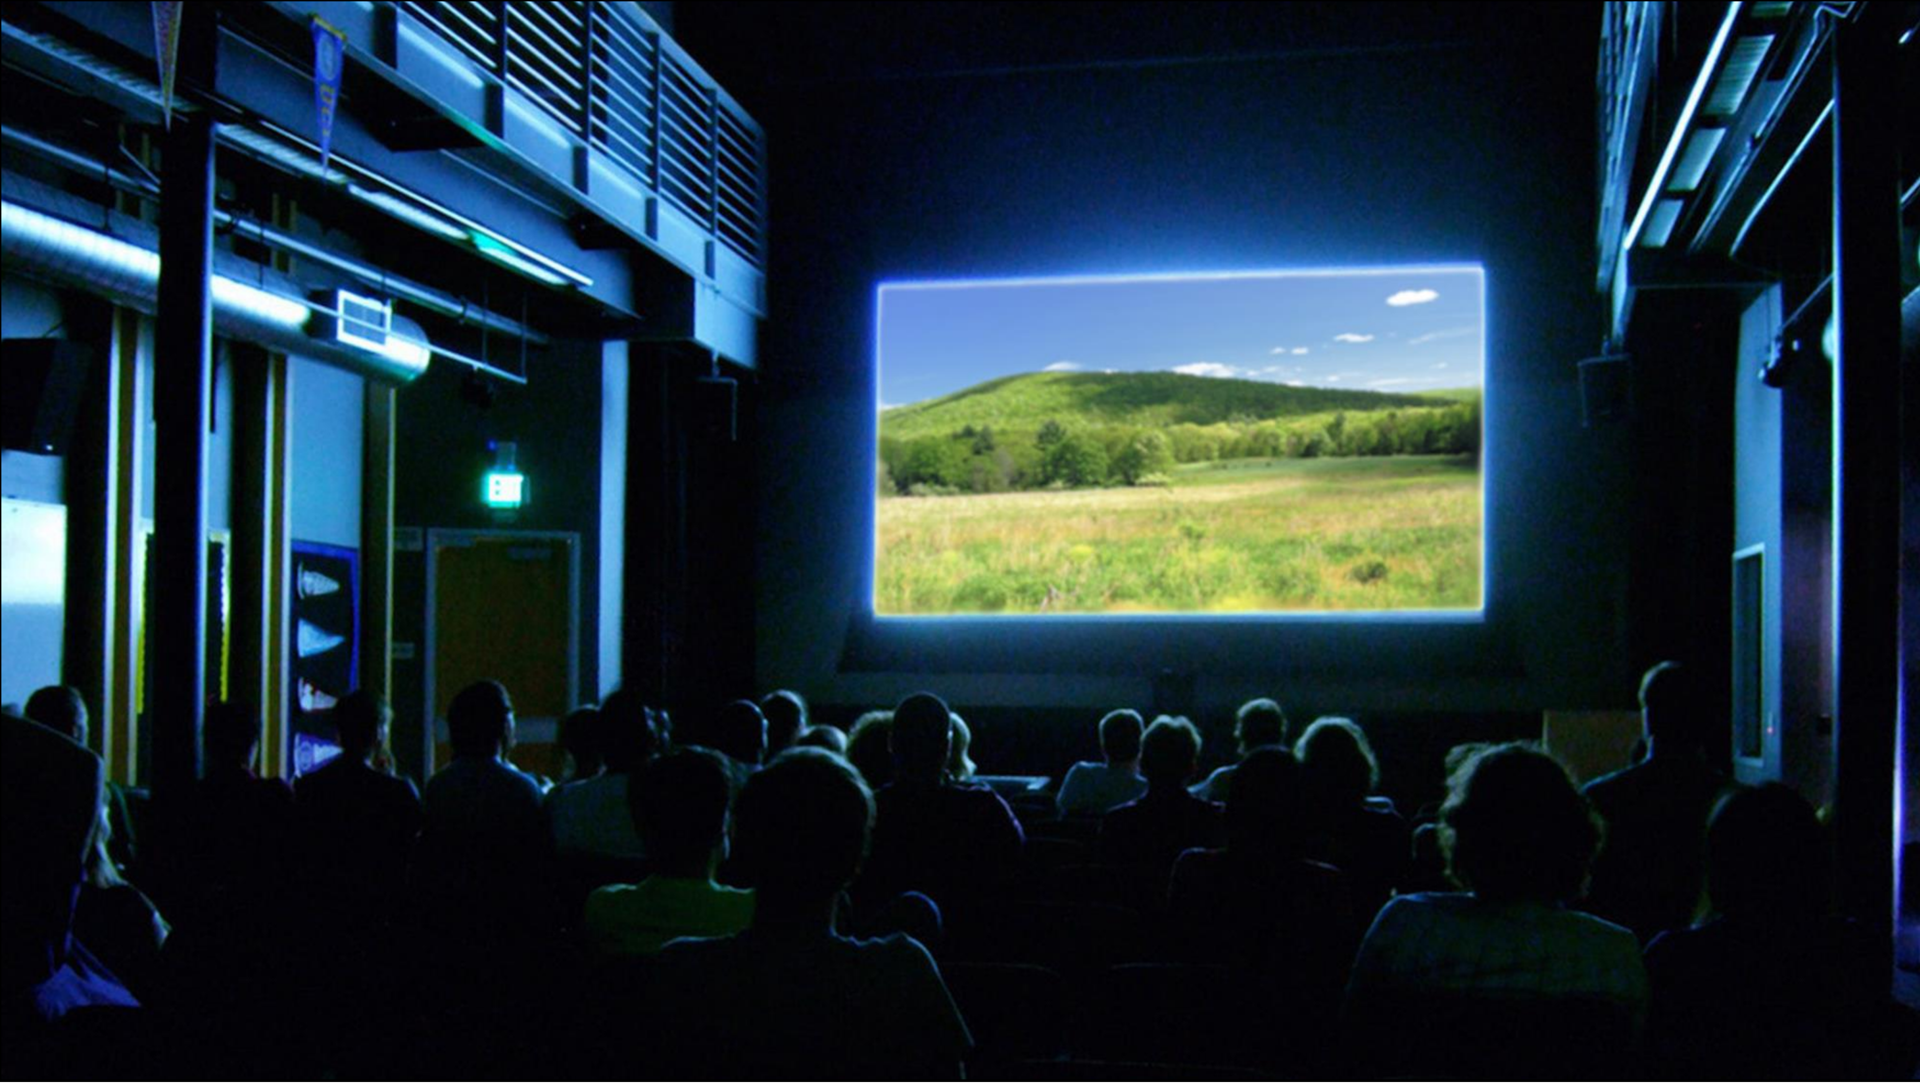

…a projection screen in a movie theater,…

(Image “Watching a blank screen” by Flickr user Kenneth Lu, used under CC-BY 2.0 / Cropped from original, “Hillside” image pasted into frame)

…a framed photographic print on a desk, etc.

Of course, the image reproduction problem doesn't make sense without the context of a well-defined…

…viewer and viewing environment. The problem of image reproduction is defined as trying to reproduce the same perceptual experience that a viewer would have had if they had been present in the original scene…

(image of man with briefcase from Microsoft Office Clip Art, used according to Microsoft Office EULA)

…onto a viewer in a different environment, watching a reproduction of the scene on some display device. This perceptual impression is affected by the visual adaptation of the viewer to the lighting in the room, the environment surrounding the image, etc.

In other words, the image reproduction problem is: “how do we create an image on the display such that the viewer gets the same (or as close as possible) perceptual impression as they would from the original scene?”. If you think about it, it’s surprising that this is possible to any degree, considering that the absolute luminance and dynamic range of the original scene is typically several orders of magnitude over anything the display device can produce. Fortunately, certain qualities of the human visual system cause our perceptual impression to be relatively independent of the absolute brightness of an image.

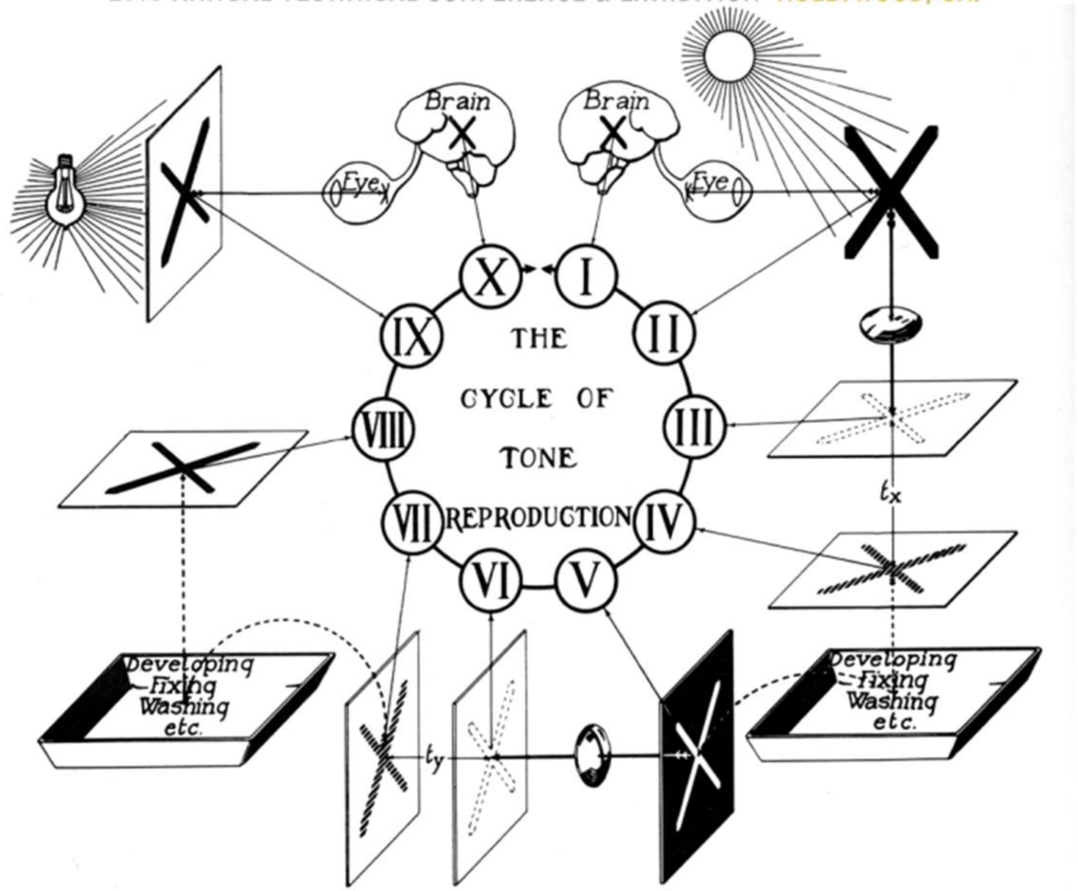

Cycle of operations in tone reproduction.

The image reproduction problem is about as old as SMPTE, which is celebrating its centennial next year. It was first formalized by L. A. Jones (Kodak Research Labs) in 1920; this is an illustration Jones made at the time—quite similar in principle to the previous slide.

(Diagram from “From Dry Plates to Ektachrome Film” used for scholarly commentary under “fair use”)

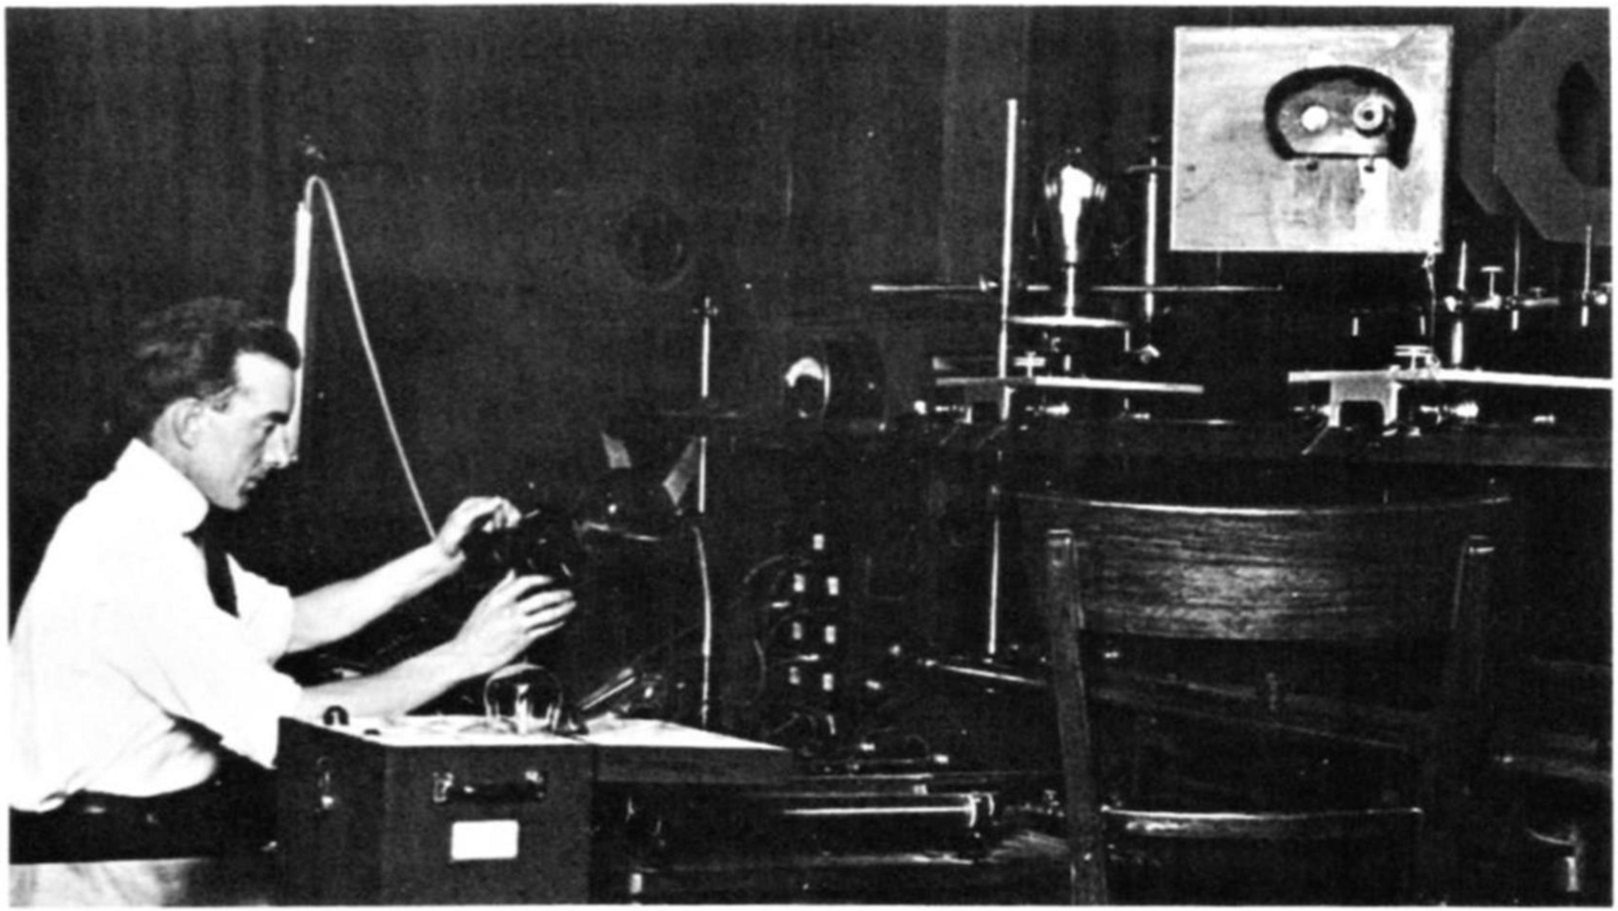

A photometer being operated by L. A. Jones.

L. A. Jones was one of the first hires at Kodak Research Labs when it opened in 1913; four years later he was promoted to chief physicist, a position which he held for the next four decades (he also served as president of SMPE—the precursor of SMPTE—from 1924 to 1925). Over his career, Jones focused on finding, refining and improving a solution to the image reproduction problem; this solution is still in wide use today.

(Image from "From Dry Plates to Ektachrome Film" used for scholarly commentary under "fair use")

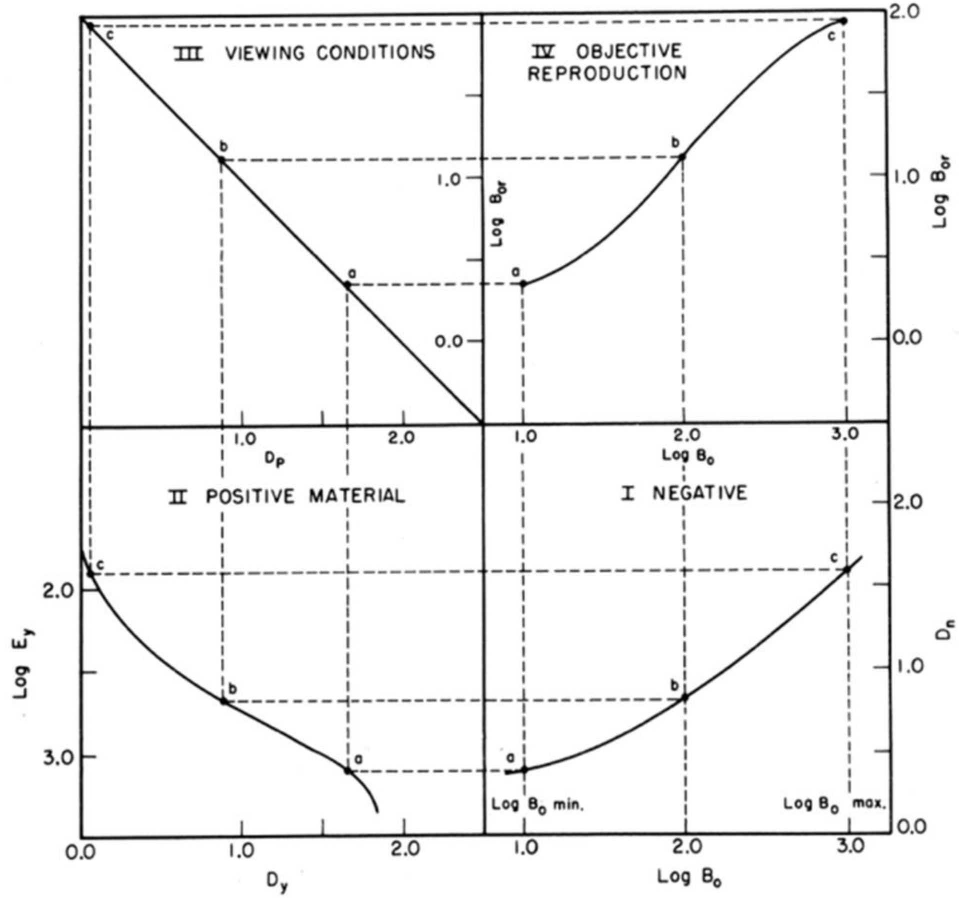

Graphic solution of the tone reproduction problem.

Jones' solution (shown here in a type of diagram he invented, which is unsurprisingly called a Jones diagram) was the photographic tone reproduction operator. This operator drove the design of film emulsions, and later the behavior of electronic imaging systems. The innovation of this operator is in its psychophysical properties – chemistry was merely its first implementation.

(Image from "From Dry Plates to Ektachrome Film" used for scholarly commentary under "fair use")

All the displays this transform has been used for over the years share some basic similarities. Whether it's television,…

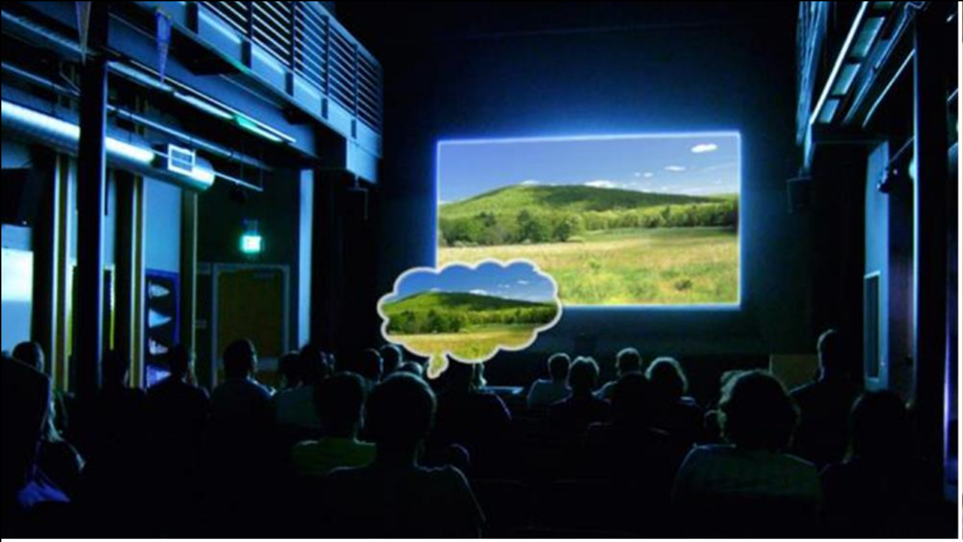


…movies,…

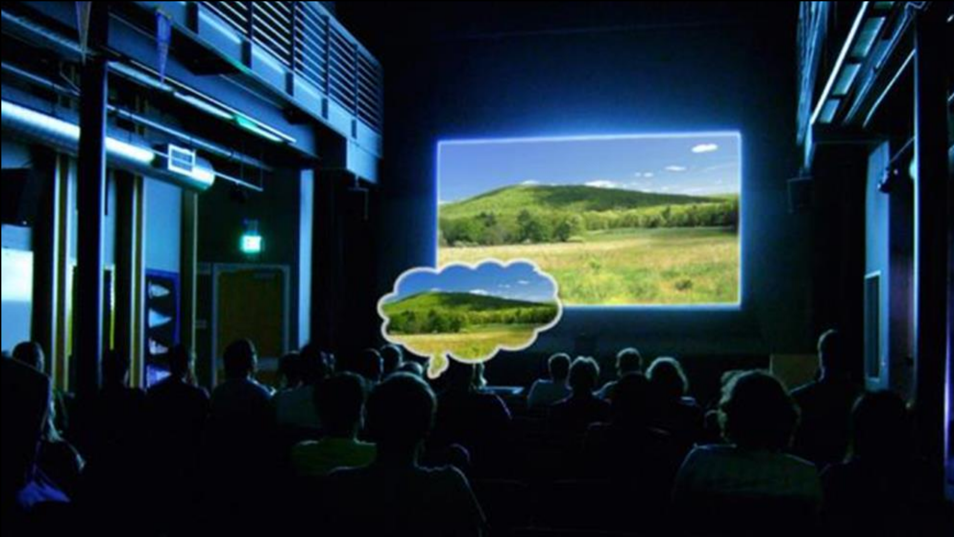

…photographic prints,…

…even paintings (which since DaVinci’s time have been a pretty good, albeit slow image reproduction technology), they all are…

(“Portrait of a man standing besides a table” by Arthur Devis, 1745, public domain / Cropped from original), (“State Drawing Room in Boston Manor House” by Flickr user Maxwell Hamilton, , used under CC-BY 2.0 / Cropped from original, “Hillside” image pasted into frame)

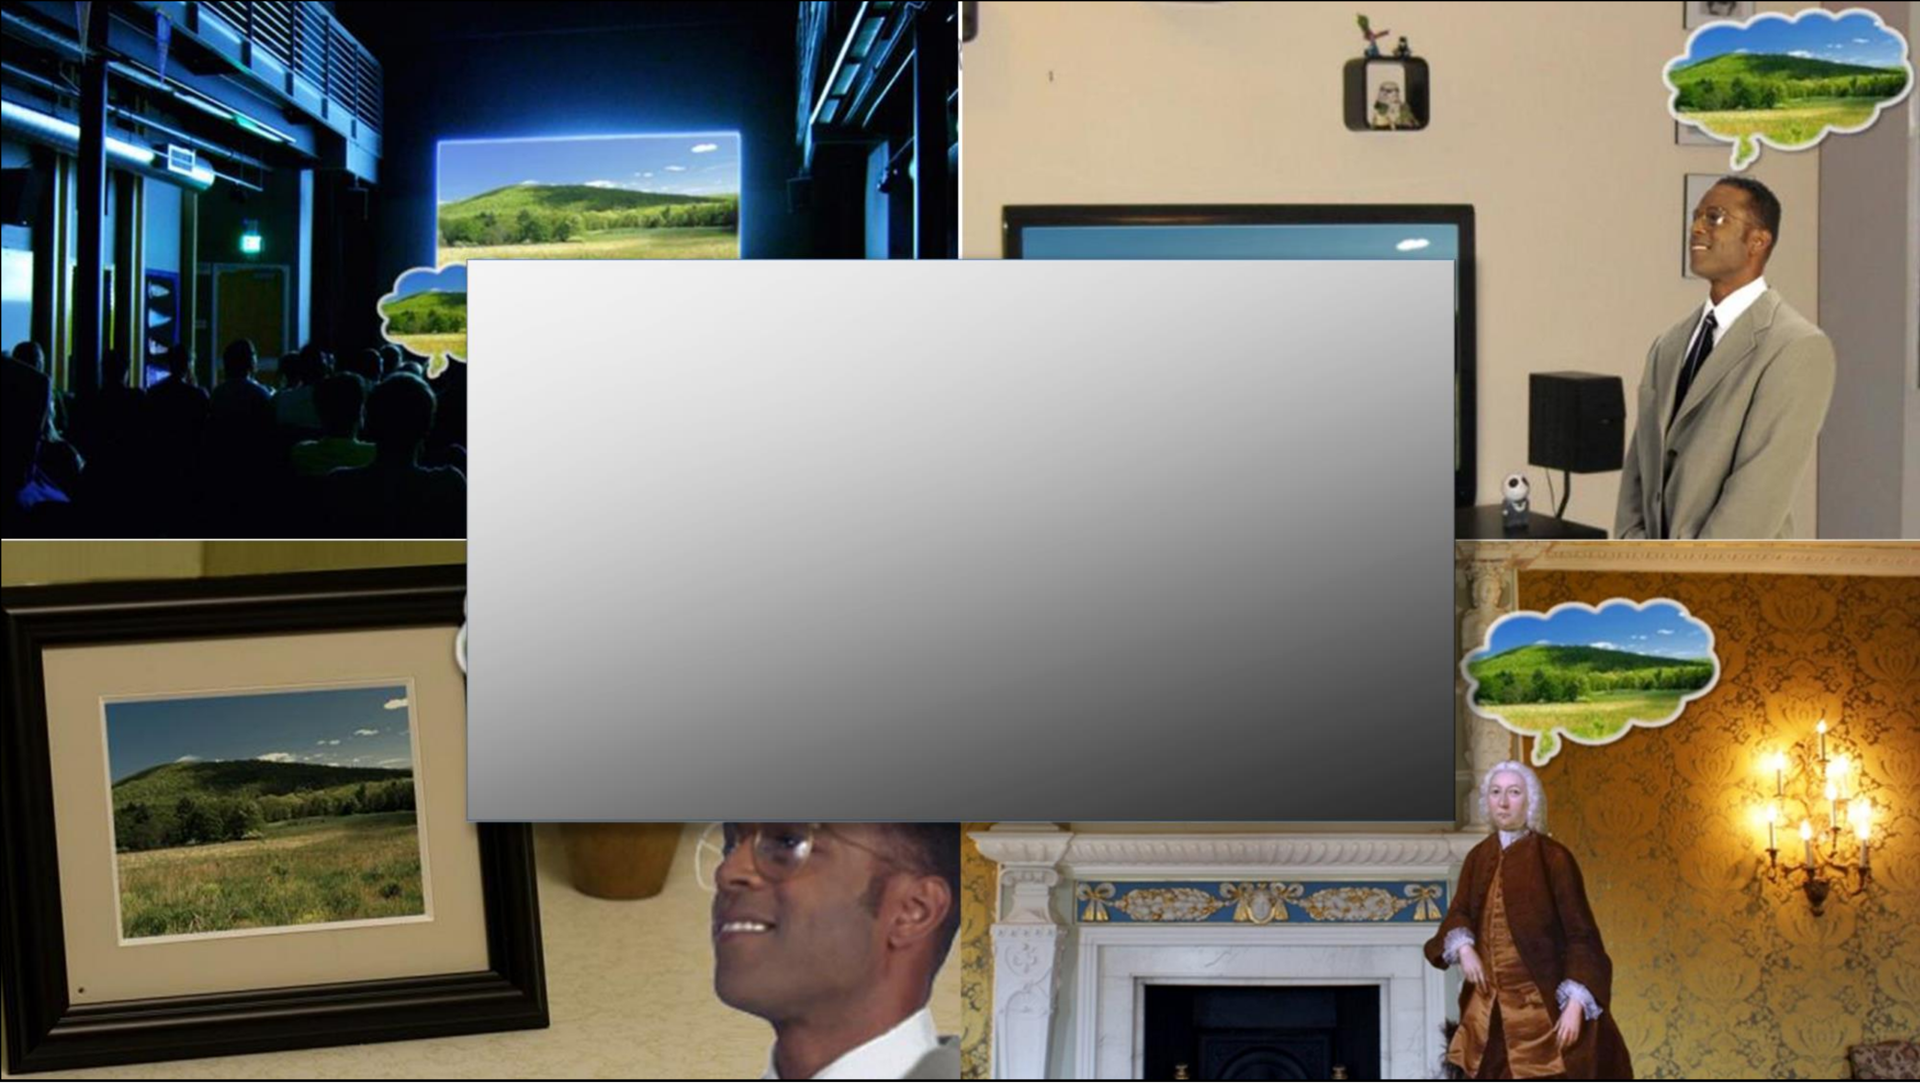

…rectangles taking up a not-very-large percentage of the total field of view, and have some sort…

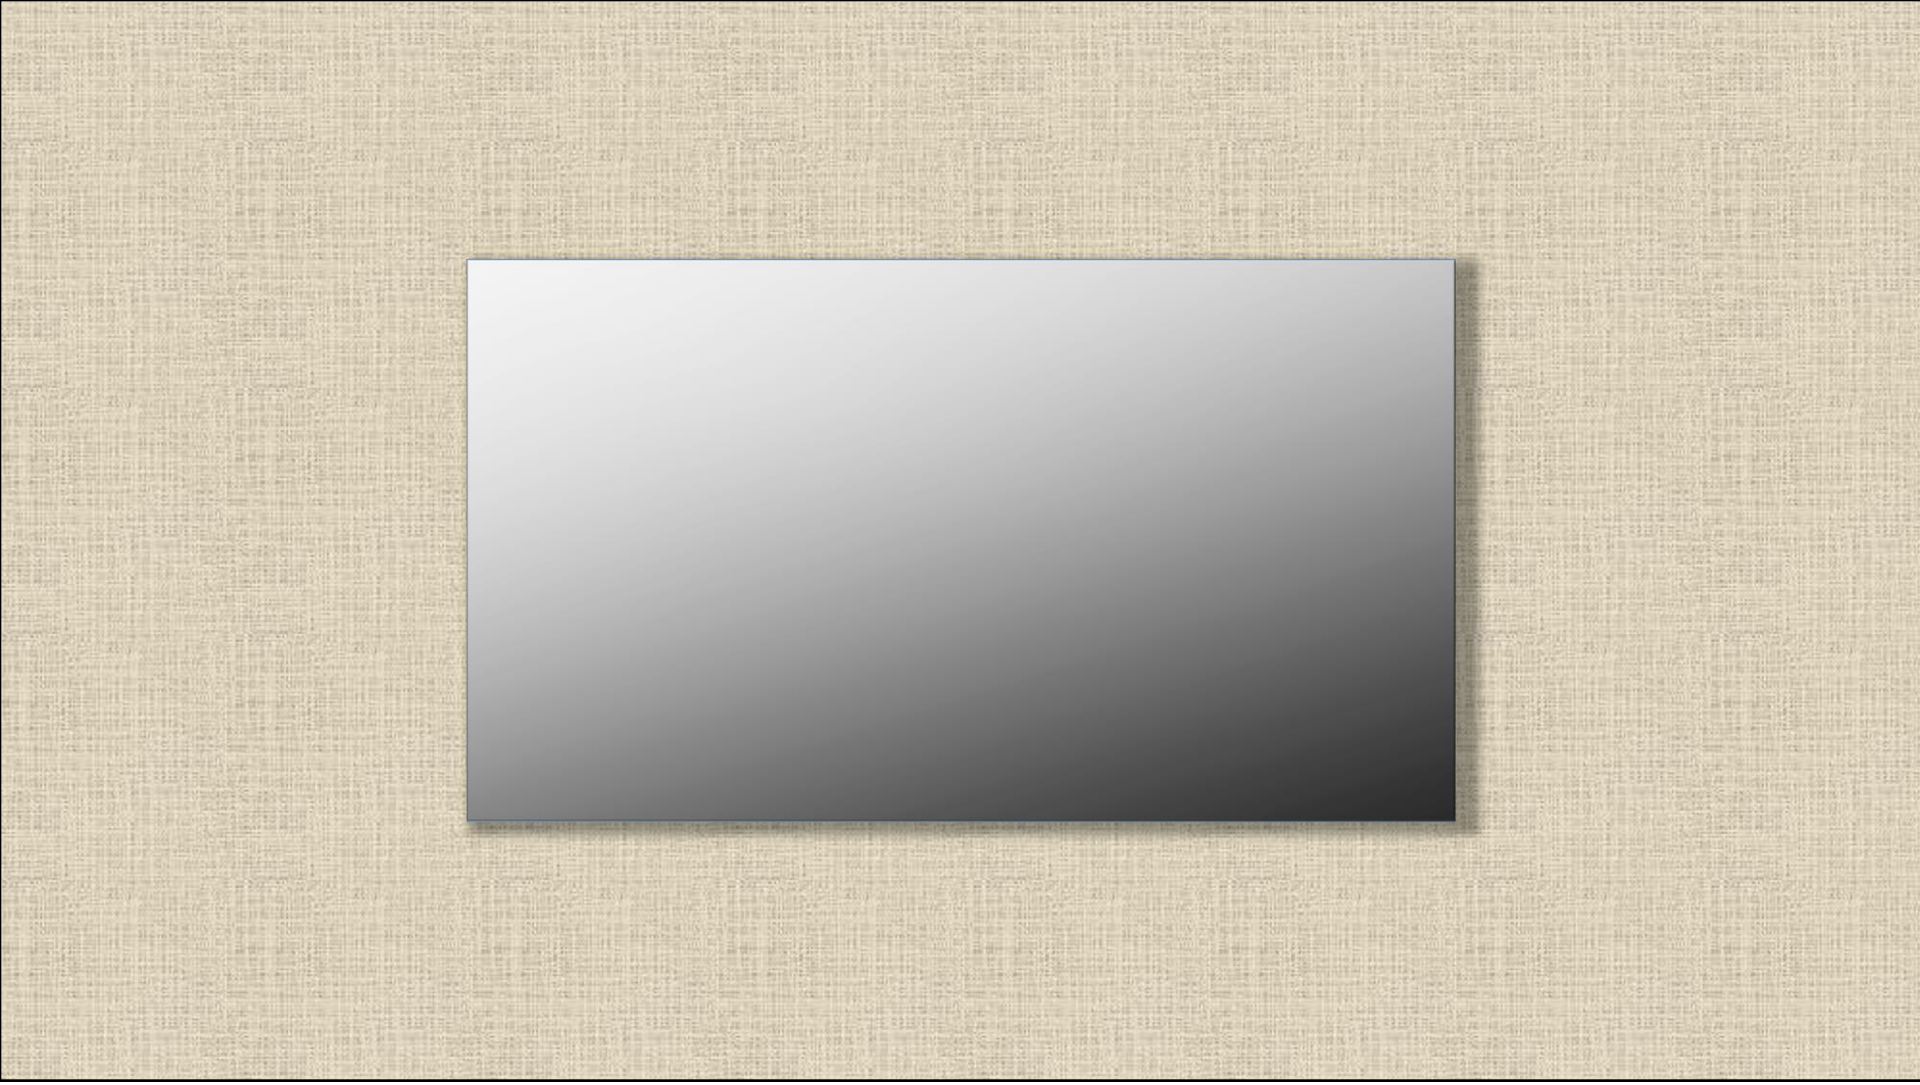

…of surround illumination. So we have a century of imaging science and several centuries of practice doing image reproduction in this situation.

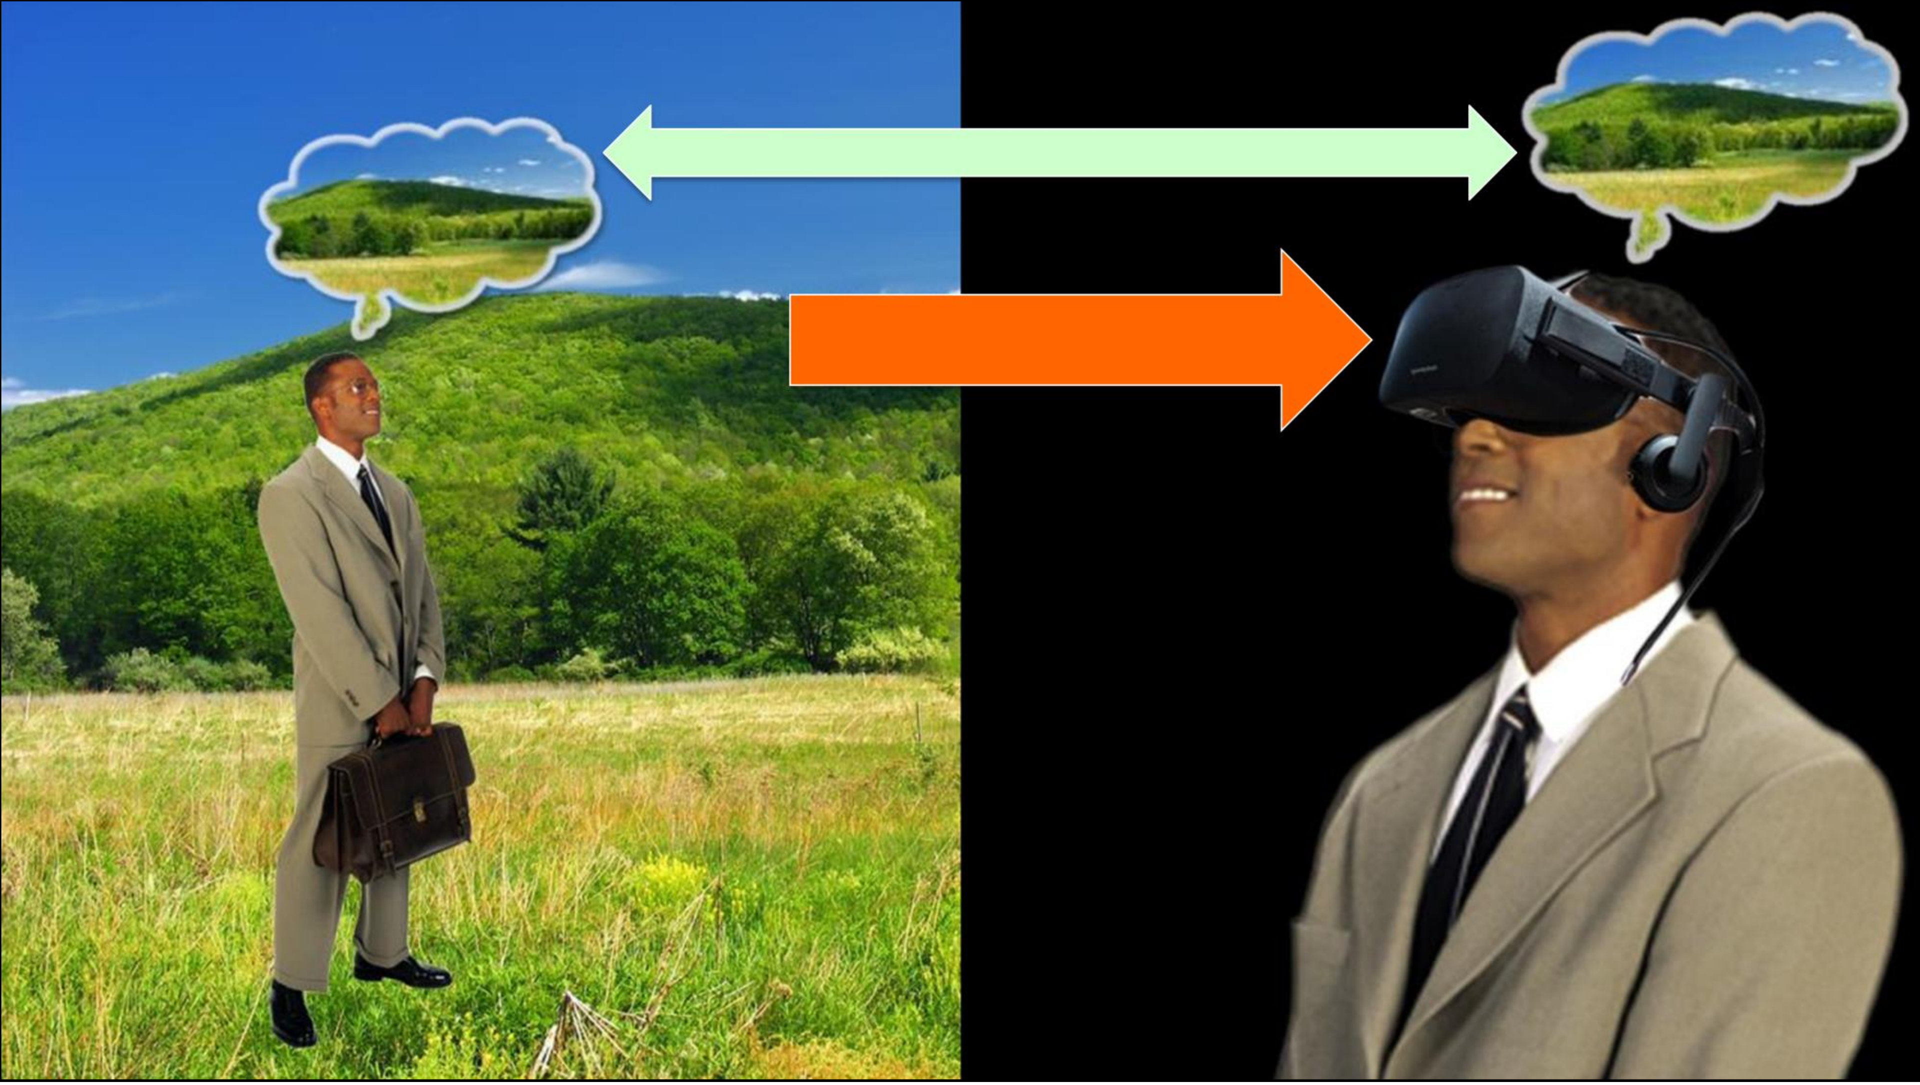

And now we have VR. Which in some ways is very different than all the previous displays. No more narrow view rectangle with a surround – now we have a full virtual environment surrounding the viewer. In VR, the image IS the surround.

(Image of Oculus Rift used for scholarly commentary under "fair use")

The VR Viewer is MORE Sensitive to…

- Quantization (banding, greyscale color shifts, etc. are more visible)

In many ways, VR is the most critical viewing environment. Quantization is much more noticeable; the JND (just noticeable difference) is much smaller in VR. Viewer is very sensitive to small level variations - visible banding. Certain types of errors like small hue shifts over the neutral scale are also very noticeable.

The VR Viewer is LESS Sensitive to…

- Global chromaticity or brightness changes

Visual adaptation is almost absolute in VR – there's no surround to inhibit it. One of the consequences is that sensitivity to global chromaticity and brightness changes is very low - even the sensitivity to such differences between the two eyes is not high (gamma differences are noticeable since they cause visible spatial detail to be different in shadows).

Color Management in VR

- Colors can't change after VR processing
  - Due to chromatic aberration correction, etc.
  - OS-side color management is tricky, at least until the VR layer and the OS are better integrated.
  - Application-side color management doesn't cause problems, but not all applications will do it.

Pixel colors can't change after the VR processing, since VR requires multiple low-level adjustments such as chromatic aberration correction. This means that OS-side color management is tricky (eventually VR processing will likely be integrated with the OS in some way, making this problem go away). Any application-side color management is fine since it occurs before passing the image to the VR system, though that's not a solution for everyone.

Another CRT-like "standard"?

- Low switching times are pushing VR IHVs to OLED
  - The brands have large (& similar) dynamic range, color gamut.
- Physical form factor and optics drive additional similarities
  - Flare is relatively high (hurts simultaneous contrast)
  - Bright displays are easier to make in small form factors
- New CRT-like "unofficial standard"?
  - Wide gamut, high sequential contrast, mediocre simultaneous contrast, high luminance

But constraints such as switching time, form factor and optics are driving certain similarities (OLED display tech, relatively high viewing flare). So we have a group of VR headsets with similar wide color gamuts and high sequential contrast from OLED, low simultaneous contrast due to flare caused by the VR optics, and very high absolute luminance (small OLED displays can be made very bright). With OS color management not working that well (at least at first), we may very well see content being tuned to an unofficial standard, just like TV used to have an unofficial standard based on the properties of CRT technology.

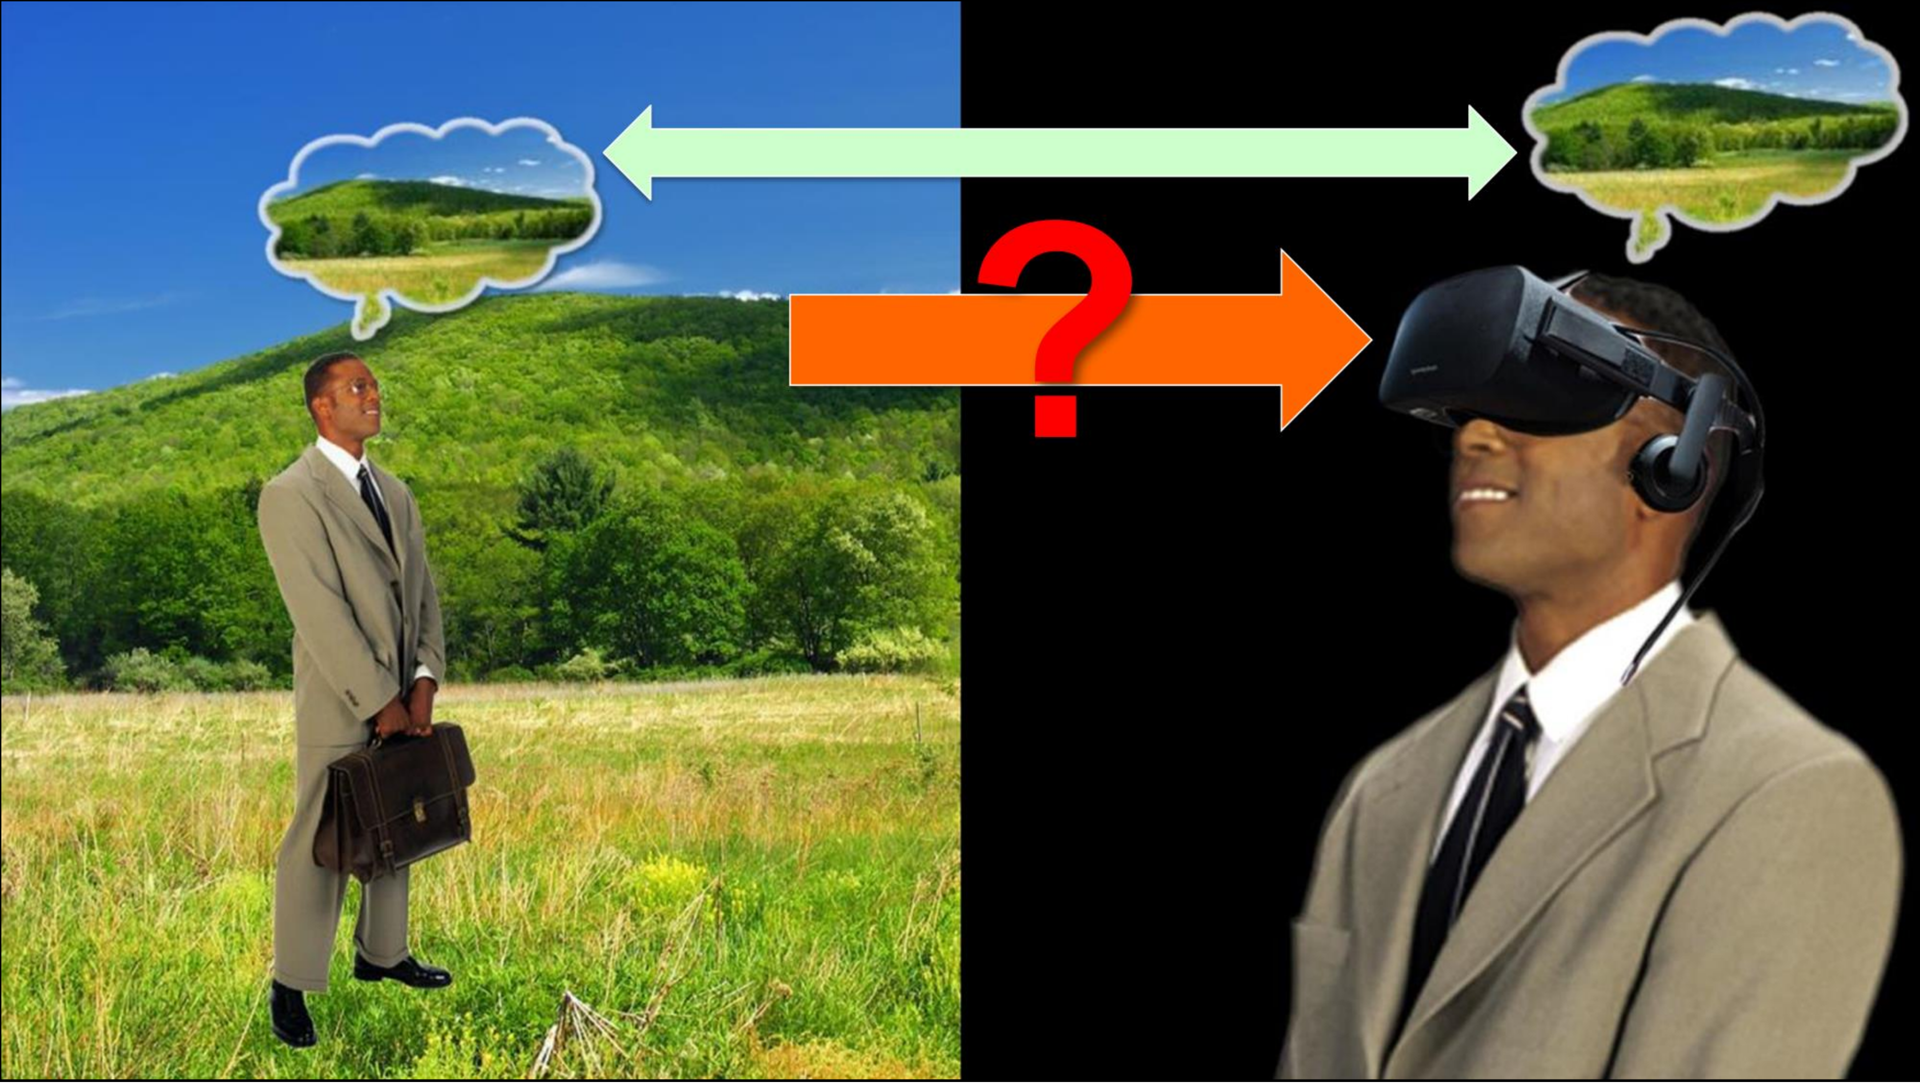

Going back to the image reproduction problem – what's the optimal transform for VR? The answer is that nobody knows for sure, but it's very likely to be different from the optimal curve for 2D displays.

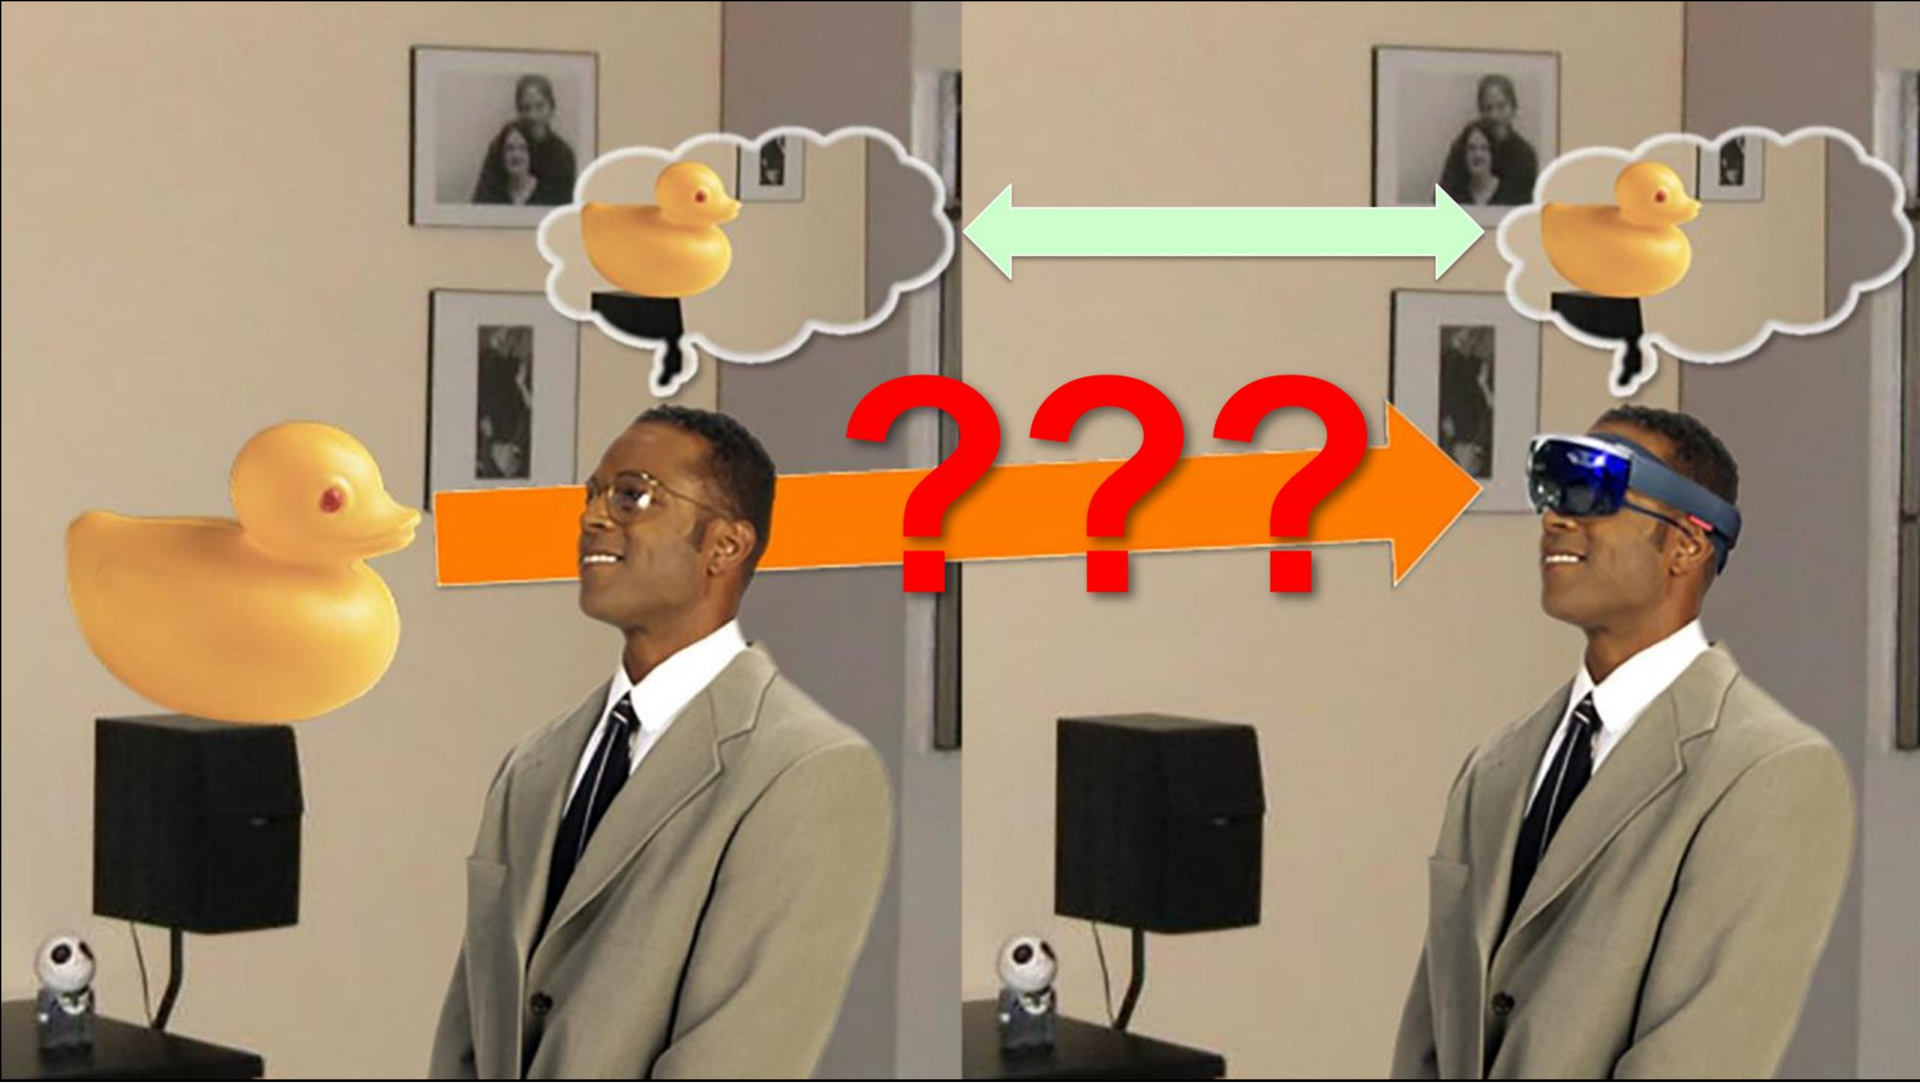

AR is the opposite—instead of no visible surround, surround is visible around, between and even (since most AR displays are additive) through the virtual objects. To appear opaque, virtual objects must be brighter than what's under them (challenging if they are supposed to look like physical objects in the scene). AR displays can't use warping optics, so field of view tends to be smaller & angular resolution higher than in VR headsets. Overall, AR image reproduction is a much harder problem than VR with many more unknowns.

(Image of Microsoft Hololens used for scholarly commentary under "fair use")

Background: Lighting

I'll give a bit of background on lighting as well.

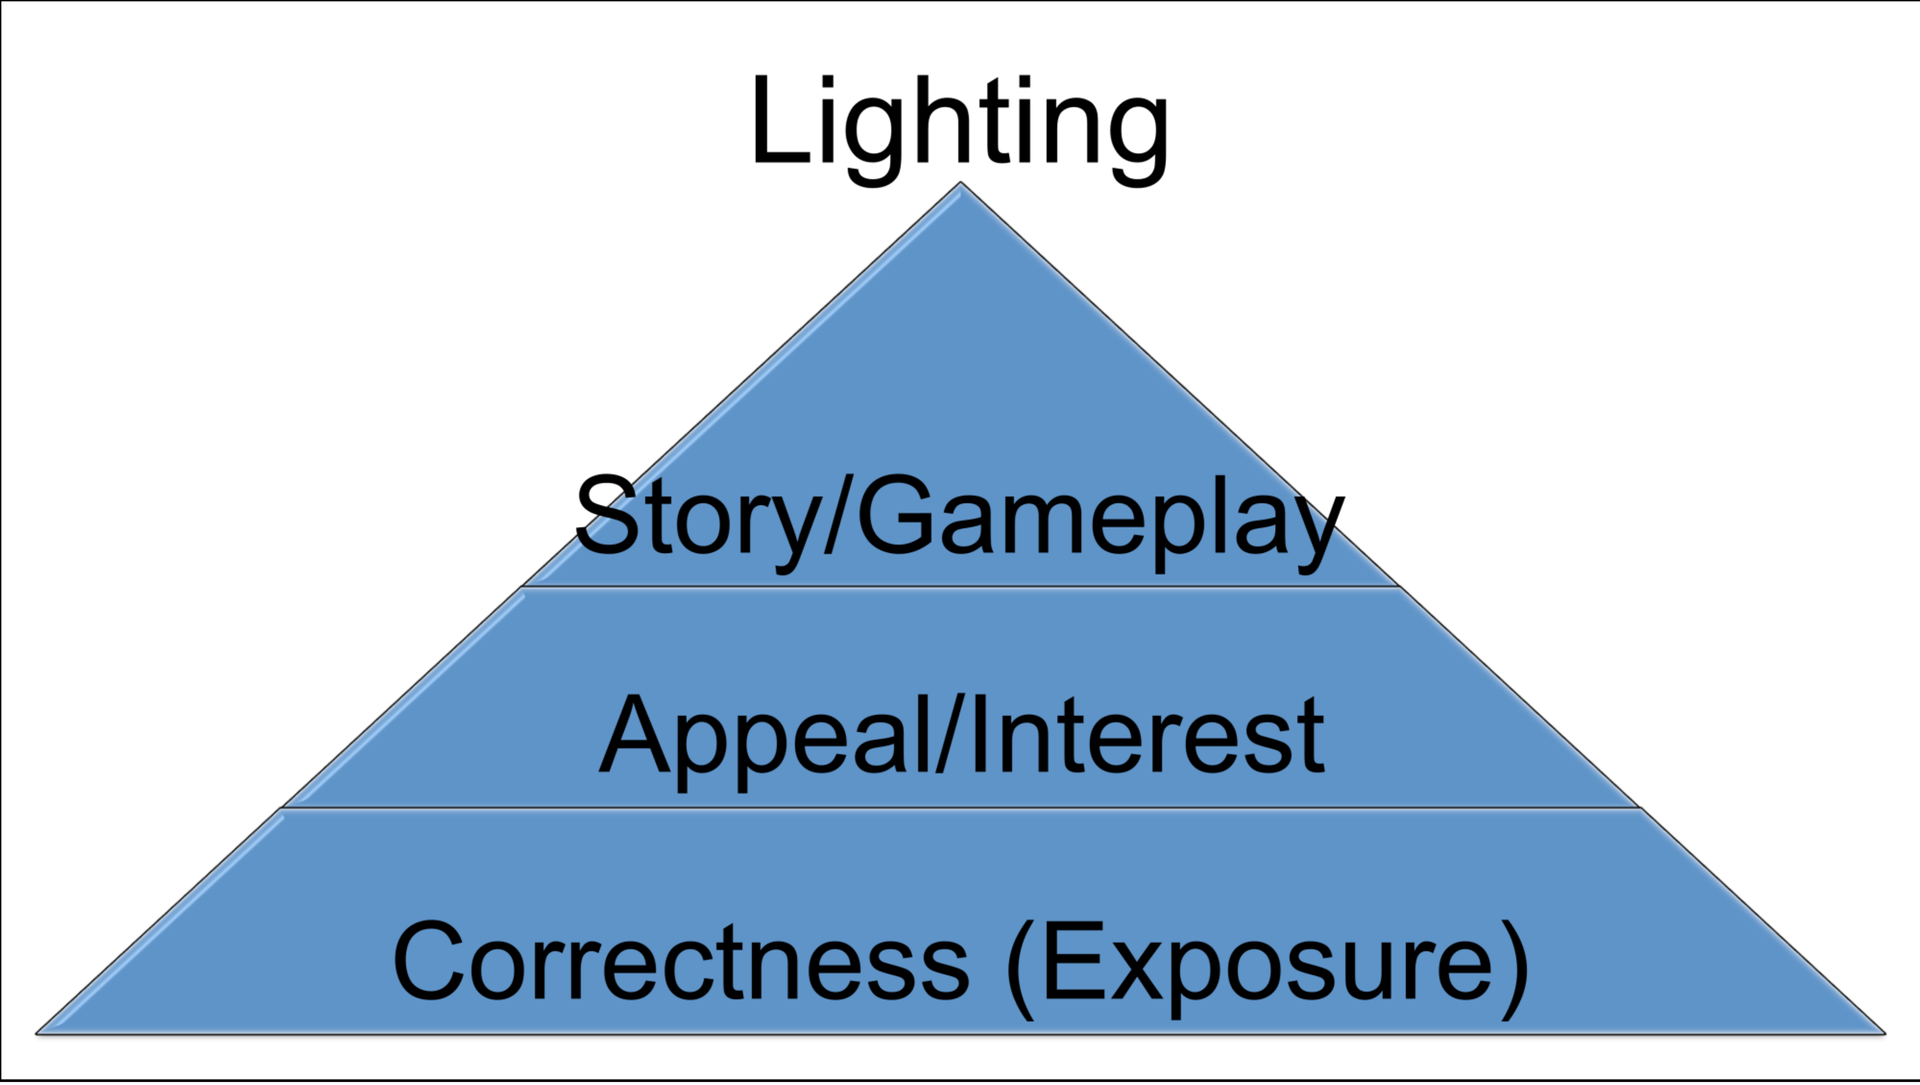

Like color, we can think of different aspects of lighting as forming a conceptual pyramid. The lighting and camera exposure control where the scene falls on the tone reproduction curve; getting this right is the foundation on which the other layers rest. The next layer is knowing how to deploy lighting to create an appealing, engaging or interesting image. And finally the lighting needs to serve the story, and in the case of games it needs to serve the gameplay: to guide the player in a certain direction, to ensure enemies are visible, etc.

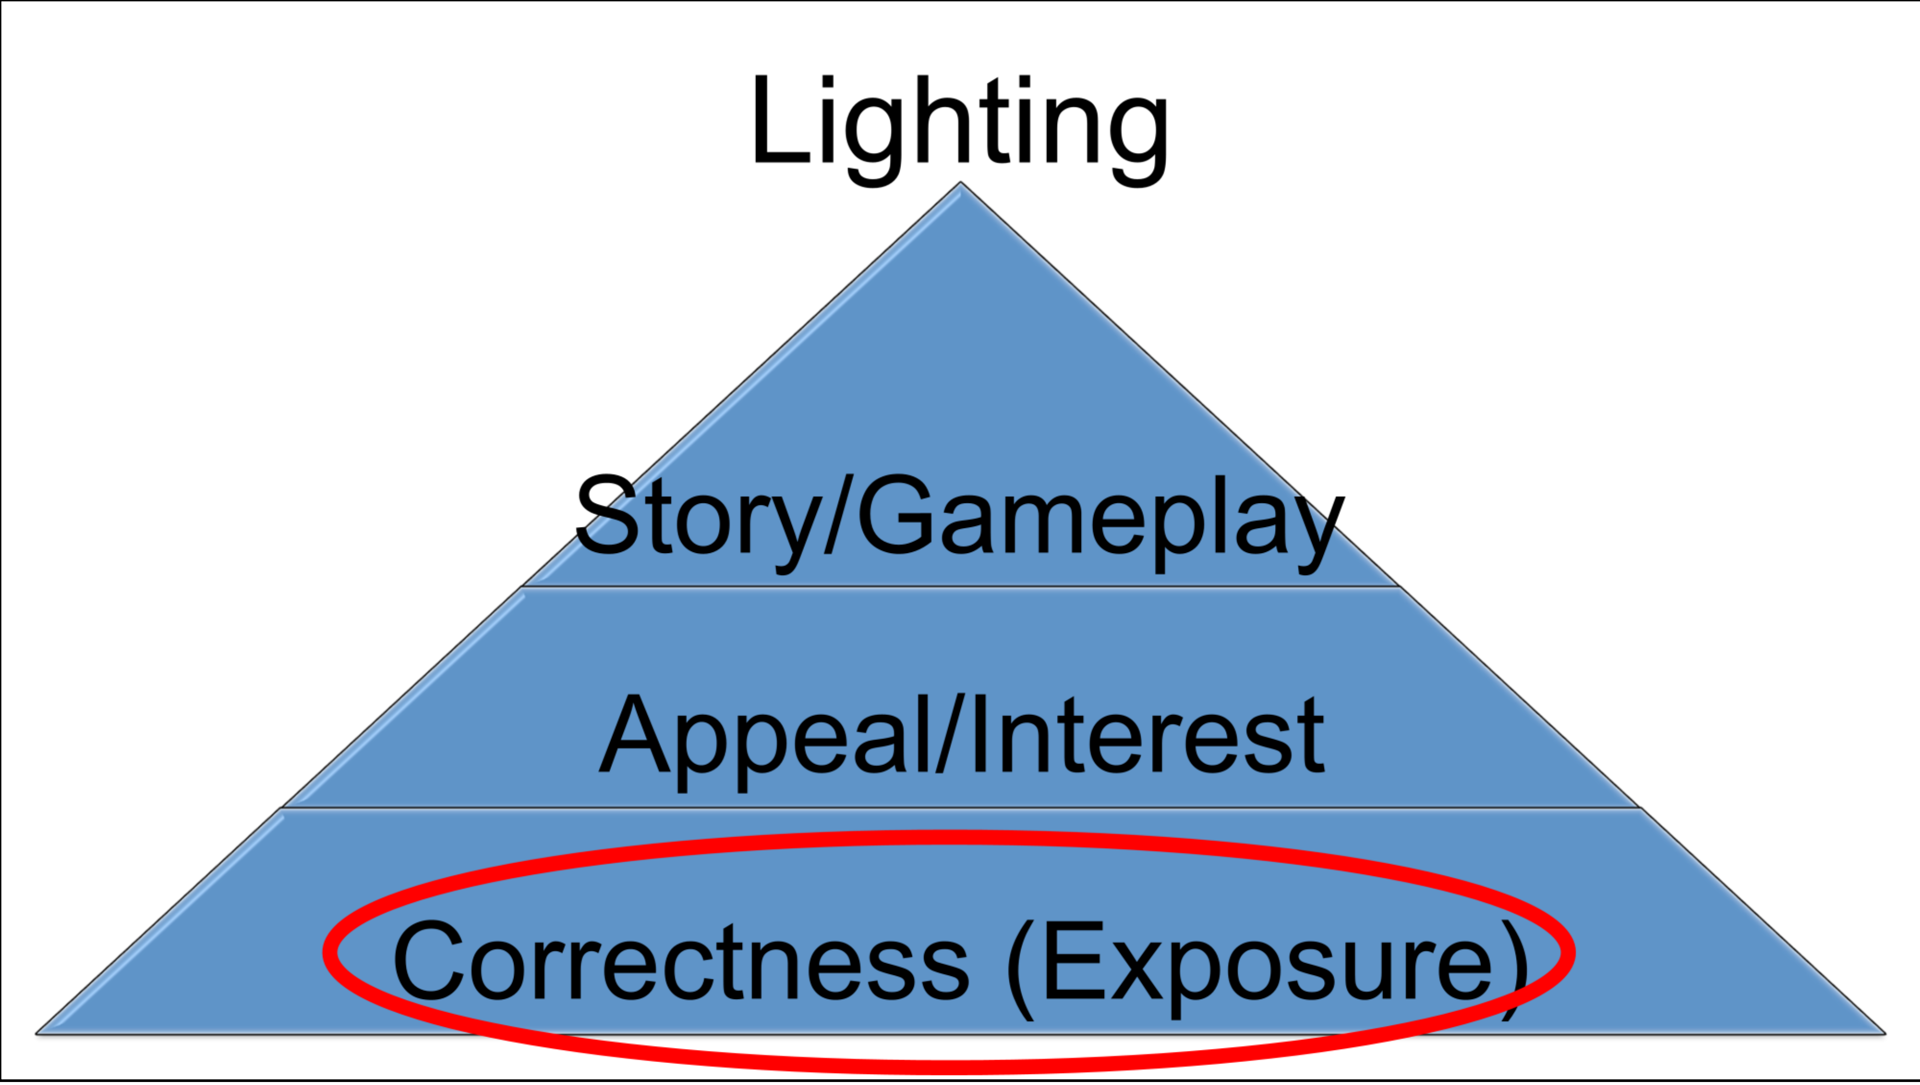

Here we'll only discuss the lowest level, which ties into tone reproduction.

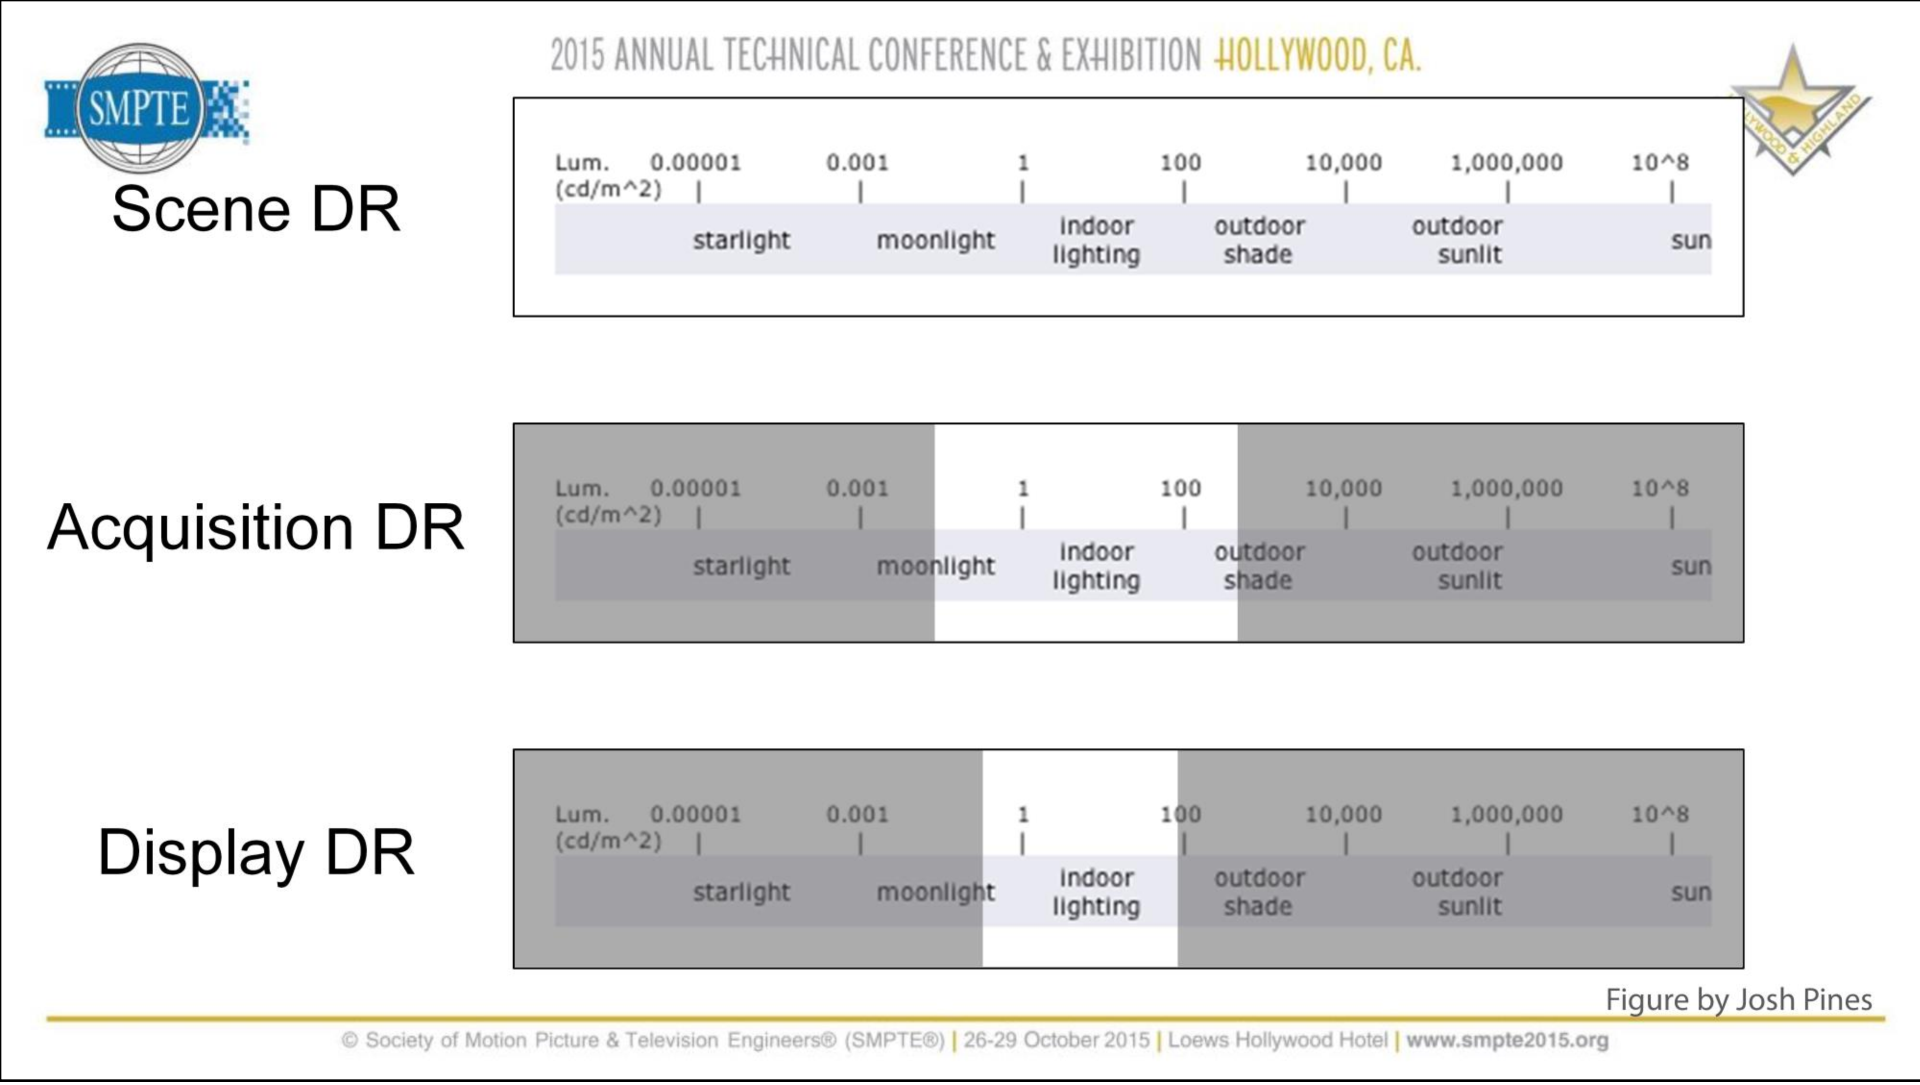

This slide (from a presentation by Josh Pines and used with his kind permission), shows how the scene dynamic range relates to the camera’s acquisition dynamic range, and to the dynamic range of the display. Since not all of the scene’s dynamic range can be acquired, the first step is to ensure that the important content is. That can be done by…

Scene DR

Acquisition DR

Display DR

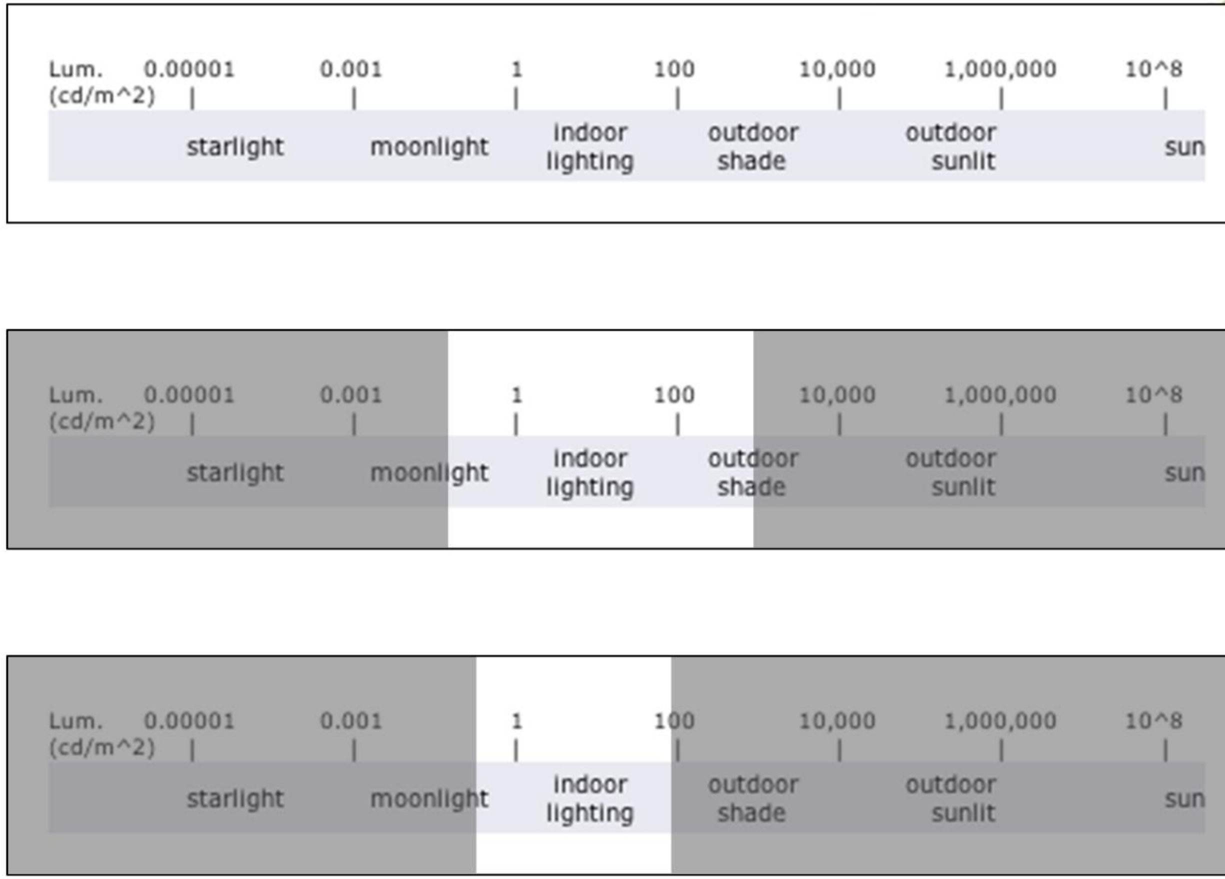

Figure by Josh Pines

…adjusting exposure…

<NOTE TO PRESENTER: flip back and forth over this slide and the next two a few times before moving on>

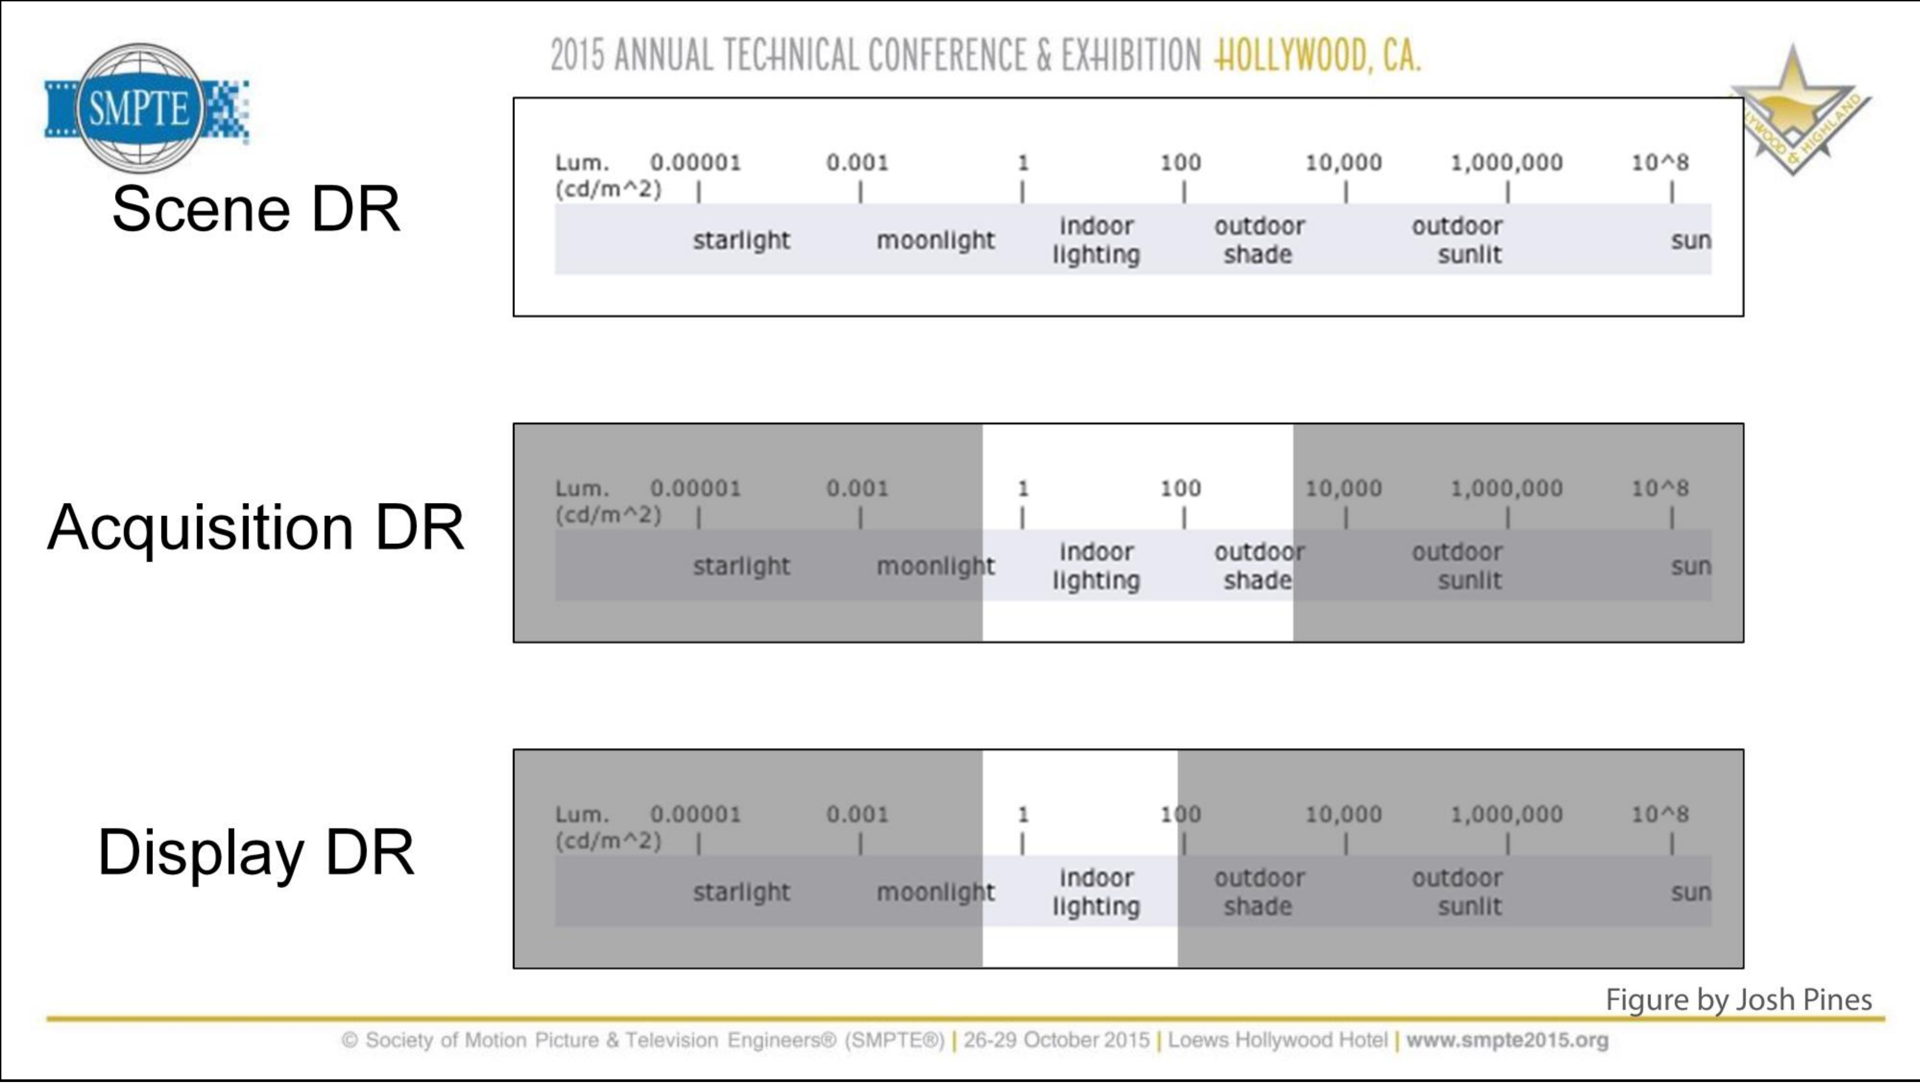
2015 ANNUAL TECHNICAL CONFERENCE & EXHIBITION HOLLYWOOD, CA.
SMPTE
Scene DR
Acquisition DR
Display DR
Lum. (cd/m^2)
0.00001
0.001
1
100
10,000
1,000,000
10^8
starlight
moonlight
indoor lighting
outdoor shade
outdoor sunlit
sun
Figure by Josh Pines
© Society of Motion Picture & Television Engineers® (SMPTE®) | 26-29 October 2015 | Loews Hollywood Hotel | www.smpte2015.org

…

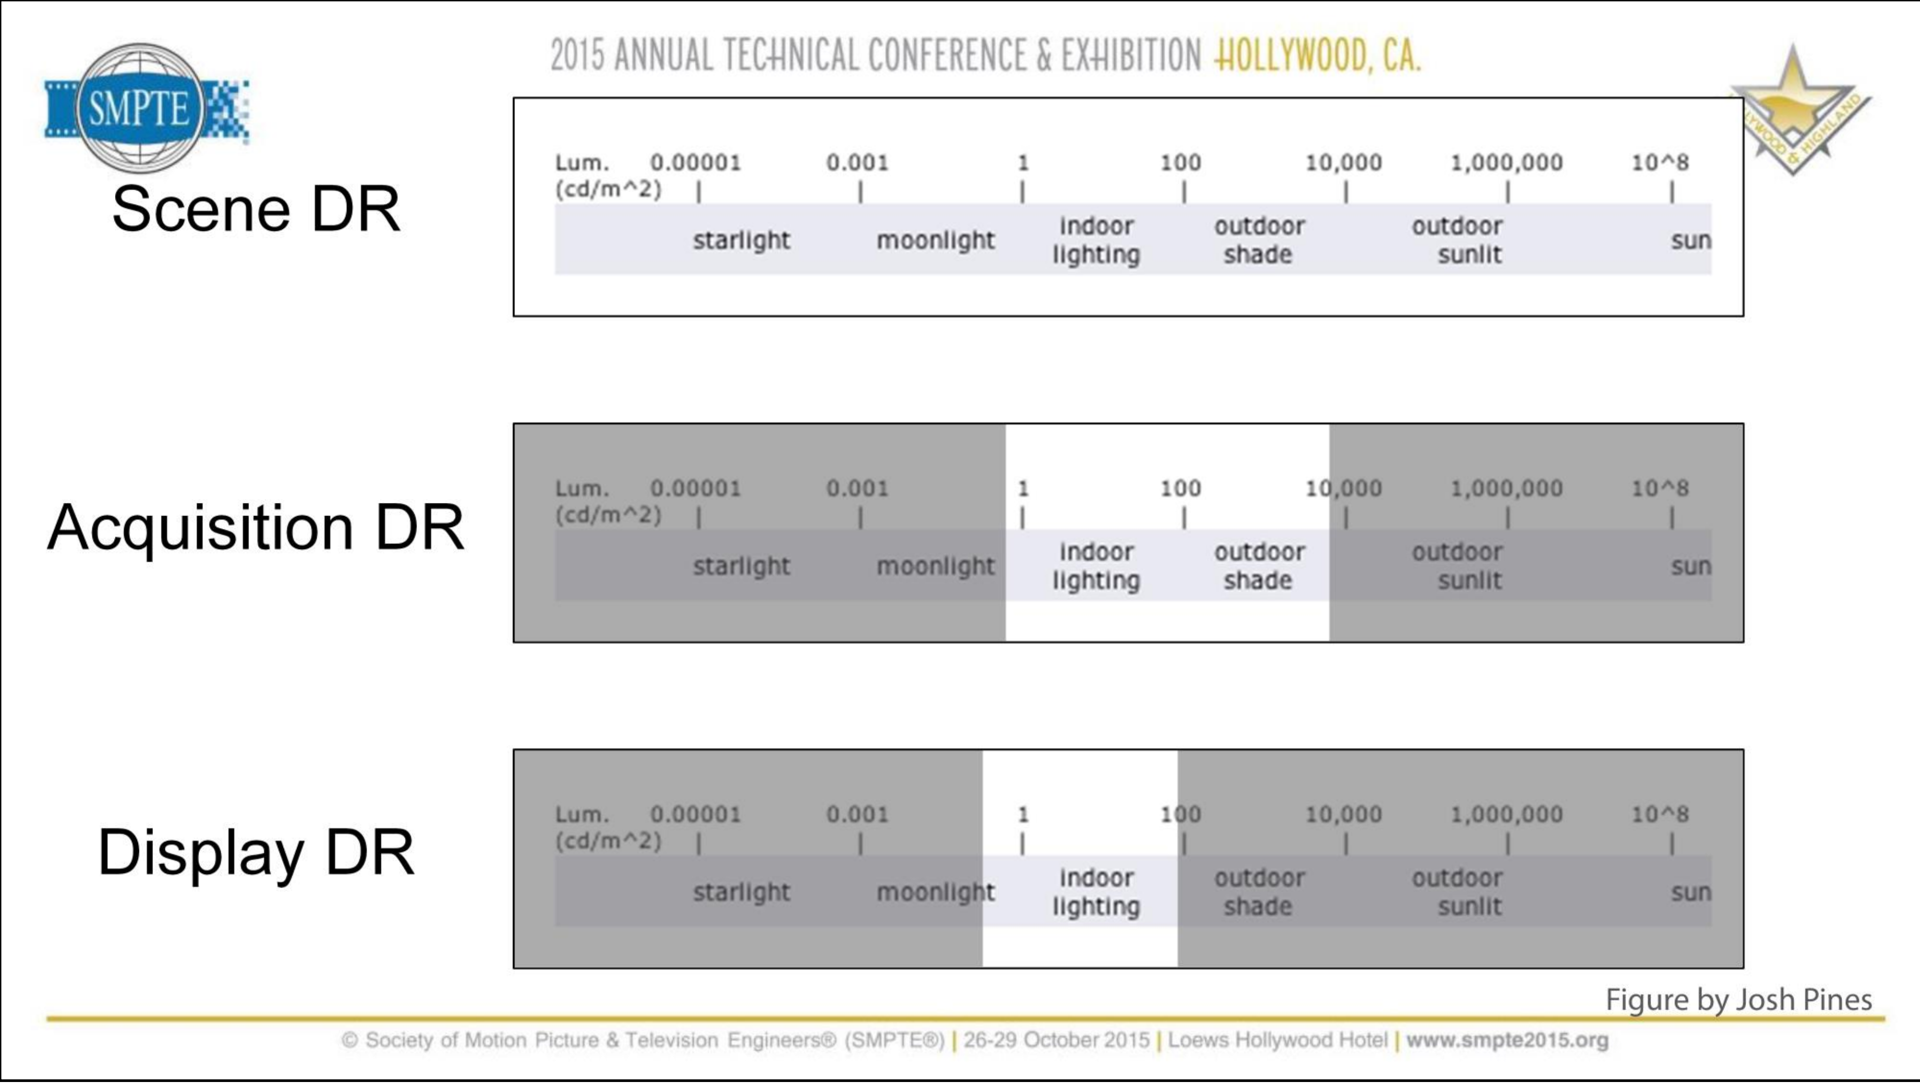

…or by adjusting aspects of the lighting (especially if the dynamic range needs to be reduced).

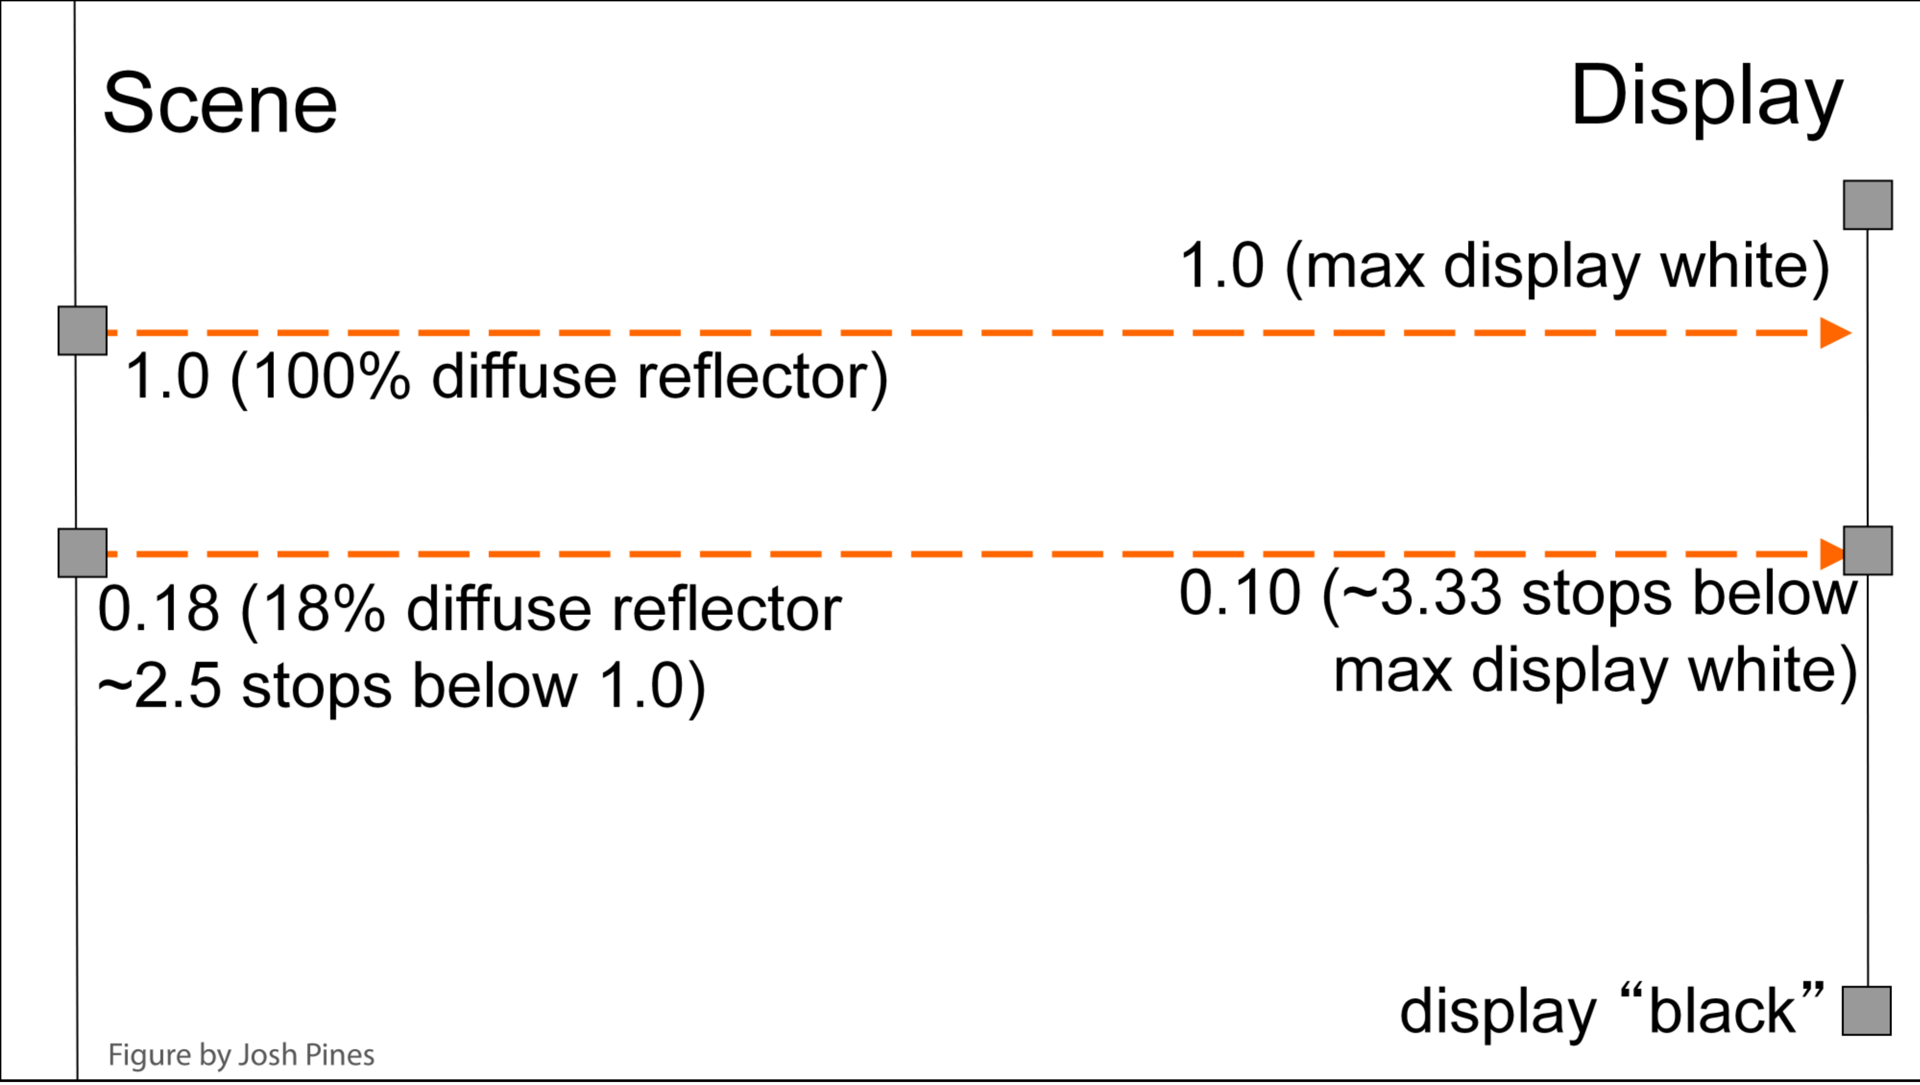

Another Josh Pines figure, showing how well-designed tone transform curves map scene values to display values. Note that scene values are proportional to reflectance; typical lighting and exposure follow the principles of the Ansel Adams zone system, primarily canceling each other out and leaving display values proportional to reflectance. This ties into the human visual system “discounting the illuminant” – our brain tends to ignore lighting and tries to reverse-engineer reflectance instead (sometimes wrongly, which is the basis of many popular optical illusions and the “what color is this dress” social media controversy).

Lighting and Color in AR & VR Games

In the next part of the talk, I'll discuss some considerations for lighting and color in games, VR and AR games in particular.

Before I talk about VR games, I'll show a few examples of games designed for 2D displays. In fact my first example, Candy Crush Saga, is a 2D game in the old sense of the term. Abstract games with 2D sprite graphics are still very successful, especially in the mobile space. They typically don't have much lighting, but color plays an important part.

(Image from the videogame "Candy Crush Saga" used for scholarly commentary under "fair use")

Flower is an example of a game with a visually rich 3D space but abstract gameplay and not much narrative. Color and lighting are used in various ways to guide the player. There are a lot of games like this coming out of the indie space.

(Image from the videogame "Flower" used for scholarly commentary under "fair use")

Although turn-based strategy games like "Civilization" moved from sprites to polygons years ago, the basic layout of the game world and the player's control of the camera is mostly in two dimensions. Lighting and color are primarily used to ensure the detailed game world is clearly readable to the player, enabling them to make complex decisions quickly.

(Image from the videogame "Civilization V" used for scholarly commentary under "fair use")

The popular 1st- and 3rd- person shooter genres range from linear single-player cinematic experiences to action-packed multiplayer games (like Evolve, shown here). Shooters often alternate non-interactive cinematic sequences with interactive gameplay where the camera is fully under the player's control. Lighting and color serve the game's art direction, though visibility and legibility has to be ensured so that the player can easily identify distant enemies. Lighting can also be used to guide the player or provide other cues.

(Image from the videogame "Evolve" used for scholarly commentary under "fair use")

Open-world games like Mafia 3 comprise another popular game genre. These usually have expansive worlds and loosely structured gameplay, with a well-defined narrative and characters to keep the player's interest. Like shooters, open-world games also alternate between non-interactive cinematics and interactive gameplay. Since these games include dynamic time of day and weather, lighting is not as tightly authored as it would be in other game genres, especially during gameplay.

(Image from the videogame "Mafia 3" used for scholarly commentary under "fair use")

Games like “Until Dawn” emphasize highly scripted narrative over interactive gameplay. These games are sometimes described as “interactive movies”. These games are similar to movies in many ways. The player’s control of the camera is limited, and lighting is handled pretty much as it would be in a movie.

(Image from the videogame “Until Dawn” used for scholarly commentary under “fair use”)

VR games are new; game genres haven’t solidified yet. Most existing game genres tend not to work well in VR; player locomotion needs to be handled in very specific ways or mismatch between the visual and vestibular senses will cause acute discomfort—in other words, the player will get motion sickness and toss their cookies. VR game developers are trying out many different approaches with new game genres appearing and old ones returning. Lucky’s Tale (shown) is a colorful platformer, reminiscent of older games.

(Image from the videogame “Lucky’s Tale” used for scholarly commentary under “fair use”)

Top-down, turn-based strategy games like Skyworld tend to work well in VR.

(Image from the videogame "Skyworld" used for scholarly commentary under "fair use")

The presence of a cockpit helps with motion sickness, so many VR developers are working on driving and space dogfighting games. EVE: Valkyrie is the best known of the latter.

(Image from the videogame "EVE: Valkyrie" used for scholarly commentary under "fair use")

Since first person-shooters are so popular, there have been multiple attempts to bring them to VR despite the difficulties. Bullet Train is a recent example, which uses a teleporting locomotion system to avoid motion sickness. It also makes good use of the Oculus Touch controllers.

(Image from the videogame “Bullet Train” used for scholarly commentary under “fair use”)

If VR game development is in its infancy, AR gaming is embryonic. There are a few interesting examples, like Microsoft's adaptation of Minecraft to their Hololens AR system…

(Image from a visualization of the videogame "Minecraft Hololens" used for scholarly commentary under "fair use")

And the “Project X-Ray” demo that was recently shown for the same platform.

(Image from a visualization of the videogame “Hololens Project X-Ray” used for scholarly commentary under “fair use”)

Lighting Tools

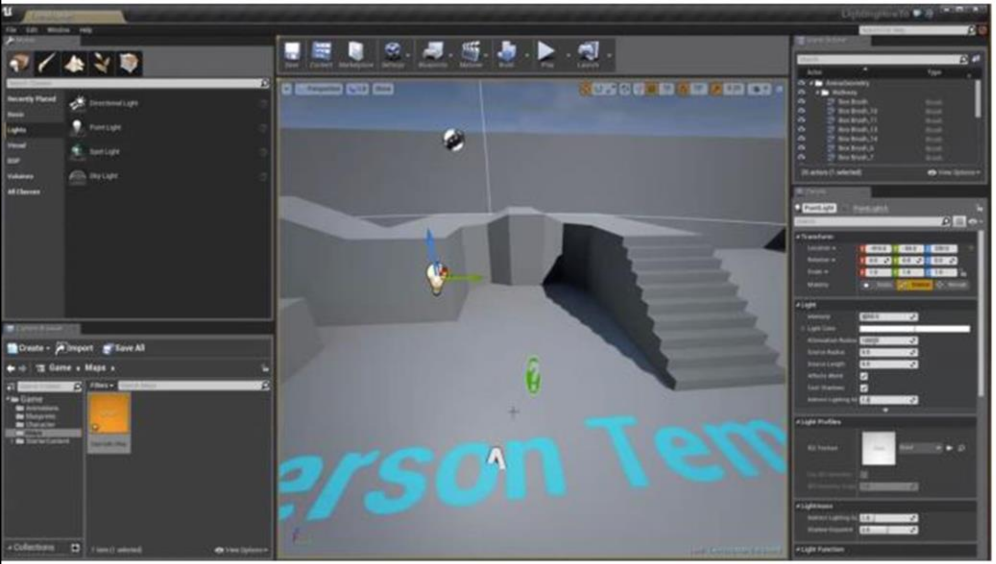

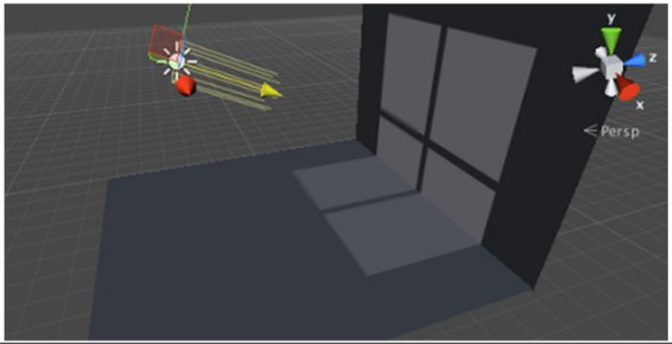

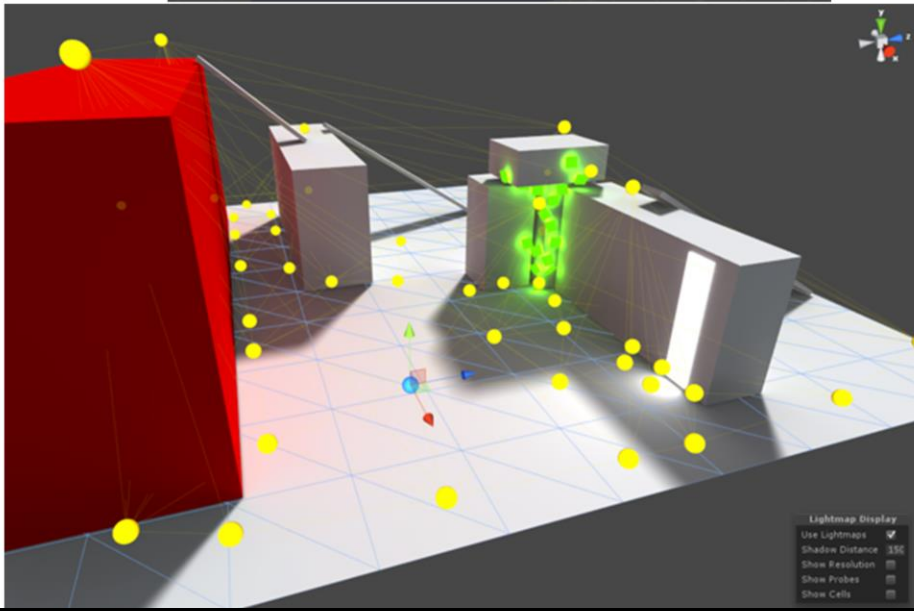

Games typically use custom tools to control lighting in game environments; these are specific to the engine used and tend to differ widely from game to game. Most games emphasize pre-computed, or “baked” lighting, since this allows for highly complex lighting including indirect bounces at low performance cost. Though the trend is towards more dynamic lighting.

(Screenshots from Unreal Engine and Unity editors used for scholarly commentary under “fair use”)

Color Grading Tools

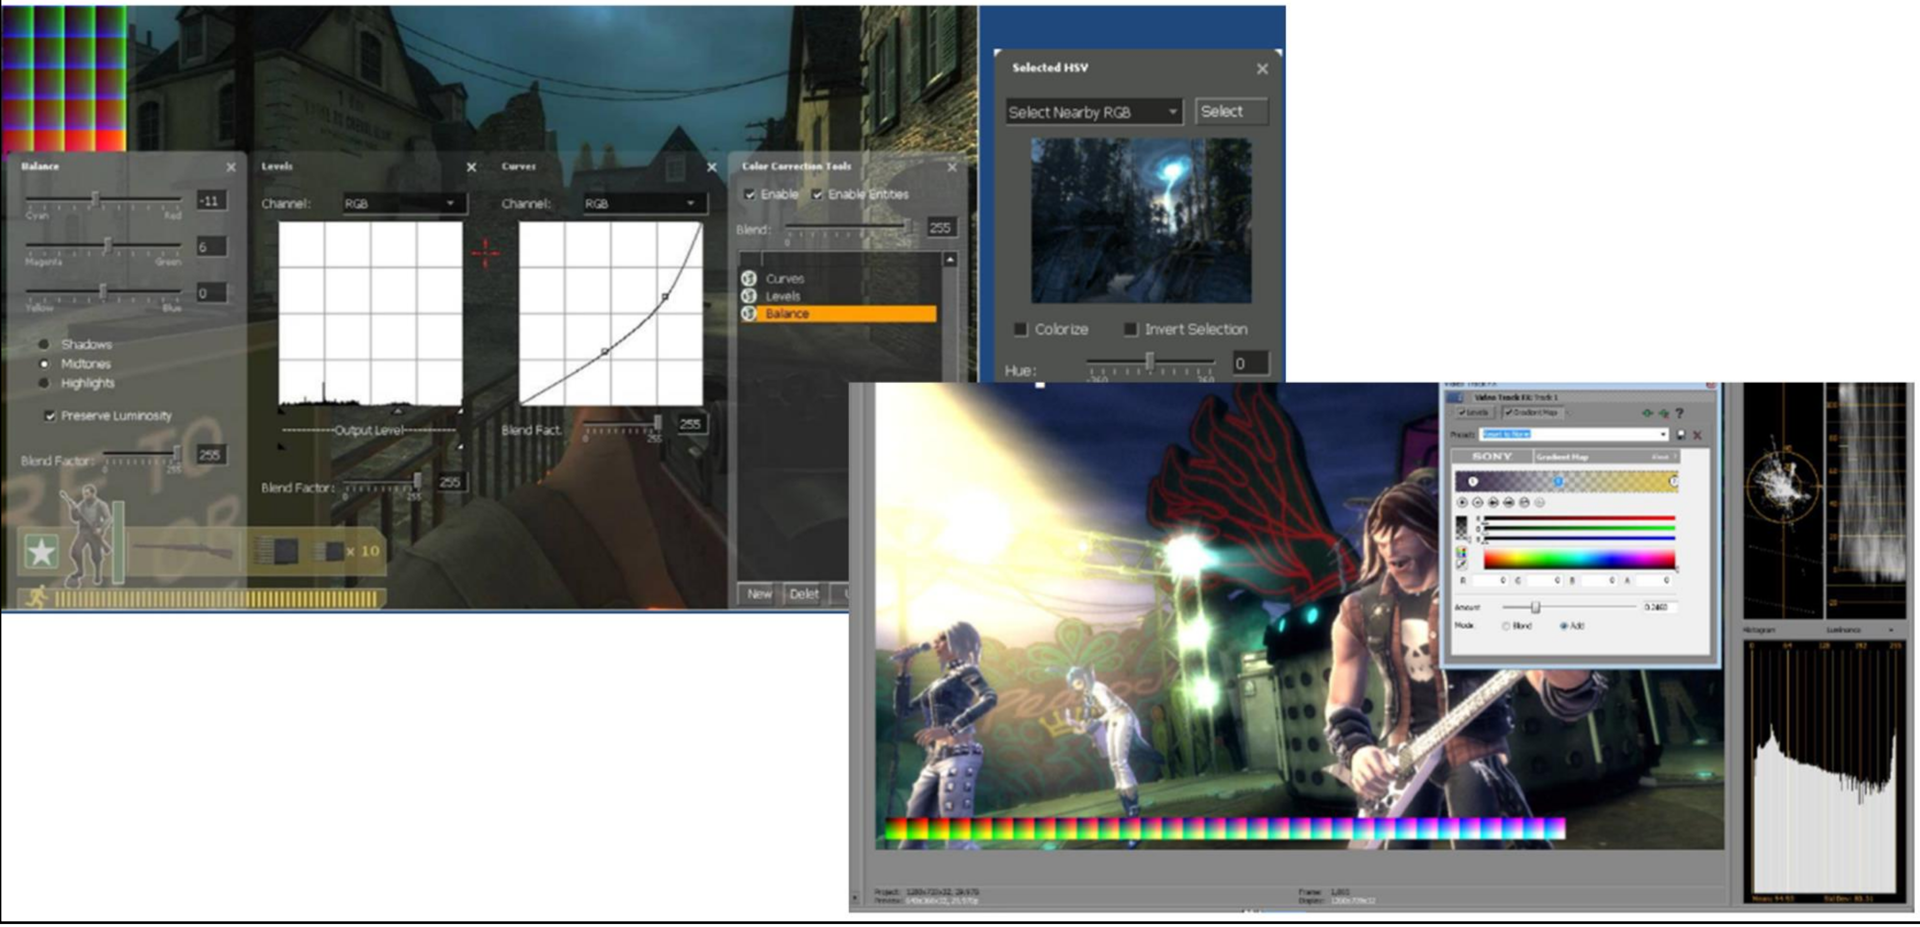

Games are (slowly) following the lead of film and broadcast with respect to color management (Unreal Engine 4.8 uses ACES transforms) and color grading (LUT-based grading in Photoshop has been popular for a few years now, and dedicated grading tools are seeing more use). As in film & broadcast, game color grading is part of post. Like film & broadcast post, game post effects run after rendering, mostly in screen space. However, they run every frame, getting a few milliseconds of the 16 or 33 millisecond total frame budget.

(Screenshots from Valve Source tools and "Guitar Hero" editing in Sony Vegas used with permission)

Motion blur is another common post effect in games.

(Screenshot from "Need for Speed" used for scholarly commentary under "fair use")

Many games use various types of camera lens effects. I'll talk about these a bit more since they are relevant to the VR discussion.

Depth of field needs to be used carefully in games since you don't know where the player will be looking, but it can be useful to visually separate distant background elements

(Screenshot from "Skylanders: Swap Force" used for scholarly commentary under "fair use")

…and bokeh (when we can afford the processing time for it) can give a cinematic look.

("The Witcher 2" Screenshot from bartwronski.com, used for scholarly commentary under "fair use")

Lens flare is often used as visual cue of the scene’s dynamic range…

(Screenshot from “Dying Light” used for scholarly commentary under “fair use”)

…and more specific effects like anamorphic lens flare…

("Killzone Shadow Fall" Screenshot from bartwronski.com, used for scholarly commentary under "fair use")

…or ghosting can help establish a particular cinematic style, or evoke a specific film.

(Screenshot from "Alien Isolation" used for scholarly commentary under "fair use")

One effect that has grown especially popular recently is a kind of "dirty lens" effect. Although it doesn't behave like an actual dirty camera lens (more like a dirty pane of glass in front of the camera), it's useful as a dynamic range cue or to establish a certain atmosphere.

(Screenshot from "The Witcher 3" used for scholarly commentary under "fair use")

…some games take this effect to extremes, in order to make a strong visual statement.

(Screenshot from "ZombiU" used for scholarly commentary under "fair use")

Lighting and Color in VR Games

VR games have many specific considerations. I won't go into things like locomotion and discomfort since that's outside the scope of this talk. But there are many aspects relating to light and color that need to be handled differently.

For one thing, all those lens effects I was just showing, no matter how successful they may be in 2D, tend to produce very bad experiences in VR. Even if implemented without obvious stereo problems (always an issue for 2D post effects), they distract from the experience and should be avoided.

Lighting VR Games

- Some things that are good in 2D games are better or more important in VR
  - Lighting that guides the player to look in the right direction
  - The player affecting lighting (moving lights, casting shadows, etc.) via 1:1 motion control (head tracking, tracked controllers)

Using lighting to guide the player in the right direction is more important in VR, because the player has a more “wrong” directions to look at. Users have found that being able to directly affect lighting with 1:1 motion (waving flashlights or torches, tracked head or hands casting shadows or being reflected in surfaces is extremely compelling in VR.

Lighting VR Games

- And some don't work so well in VR
  - Texture and normal maps for anything other than fine detail
  - Approximate antialiasing techniques

The ability to observe with small head parallax enables the observer to notice errors in specular highlights, reflections, etc. that wouldn't be noticeable at all in 2D or even non-head-tracked stereo. Texture maps have a tendency to look like wallpaper, normal maps for large-scale detail look wrong (fine detail is OK). Antialiasing techniques need to produce accurate results because aliasing is much more noticeable—due to large angular size of individual pixels, differences between eyes breaking stereo fusion, and other factors. I won't talk about what works and what doesn't work in AR because not much is known about that yet.

Preview (even author) VR content in VR!

Due to all these differences, it's important to preview VR content in VR as often as possible – ideally you should author or at least tweak content while viewing it in VR.

Acknowledgements

- Jeremy Selan for insight into color and lighting in VR
- Josh Pines for tone reproduction slides

Finally, I'd like to thank two people who helped me put this talk together. Jeremy Selan gave me a lot of good insight into considerations for color and lighting in VR. And Josh Pines provided several great slides for the section on tone reproduction.