Outside the Echo Chamber: Learning from Other Disciplines, Industries, and Art Forms

Naty Hoffman

VP of Technology, 2K

Hi. Thanks to Marc for the great introduction, and thanks to everyone in the audience for entrusting me with an hour of their time which could have been spent sleeping in or at Disney World – I hope you won't regret that choice.

?

First I’ll answer some questions that might be raised by the somewhat ambiguous title. Which echo chamber? Who is doing the learning?

One of the echo chambers I'm thinking of is the computer graphics community. Although our field is by its nature extremely cross-disciplinary, I feel we too often focus on incremental improvements in each other's work. Some of the biggest jumps in our industry have occurred from bringing knowledge and insight from the outside – e.g. the Cook-Torrance model which was adapted from an optics paper. I'll be talking about some insights I've seen from other academic disciplines, though (given my personal experience) these tend to be focused more on practice than theory.

(photo copyright 2012 Jason RM Smith, used with permission / Cropped from original)

The other echo chamber I have in mind is the game industry. I strongly feel that people in my industry – and in this talk I'm focusing primarily on people who work on the technical and creative sides of creating game visuals – would benefit from looking outside our industry for inspiration.

(photo from "GSC12 SF Signage" album by Flickr user "Official GDC", used under CC-BY-2.0 / Cropped from original)

Imaging Science

One discipline I think computer graphics should learn more from is imaging science, which focuses on the problem of image acquisition and reproduction. Imaging science includes…

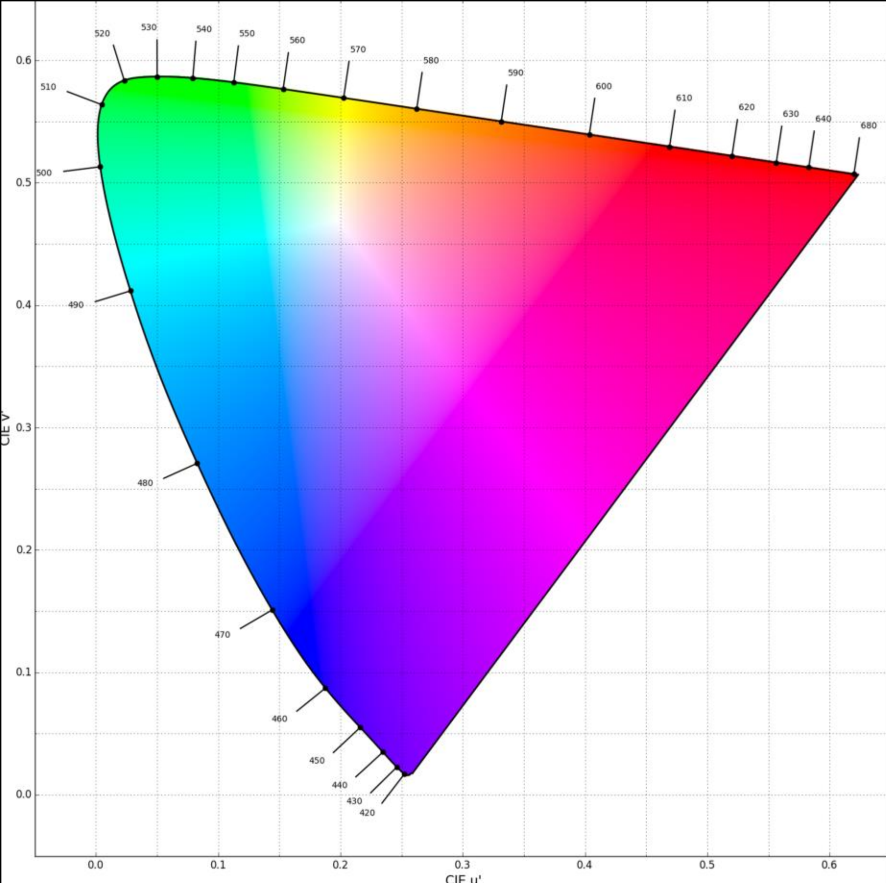

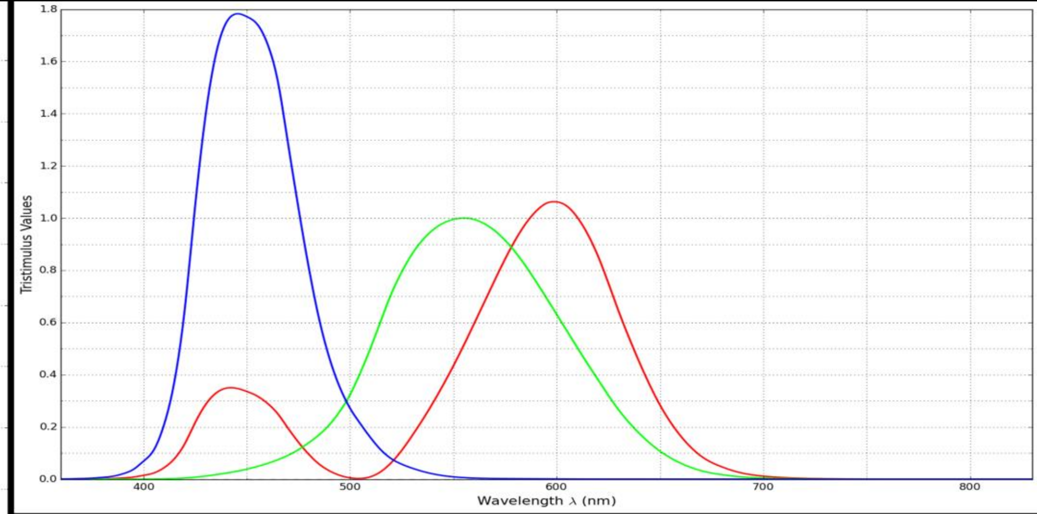

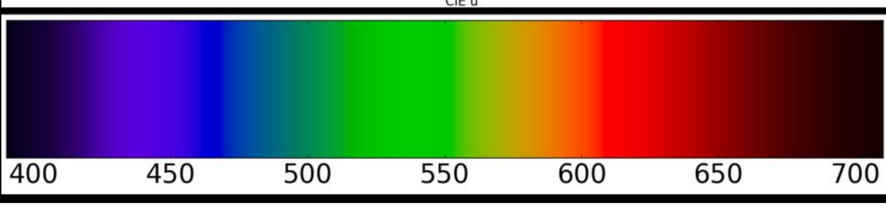

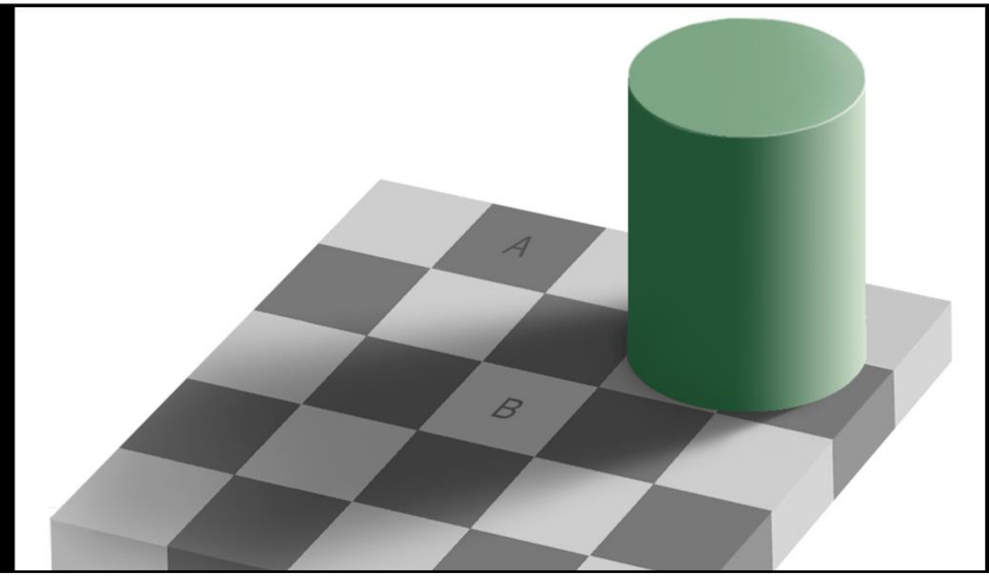

…color science and the study of the human visual system…

(File "Grey square optical illusion.png" by Adrian Pingstone used with permission)

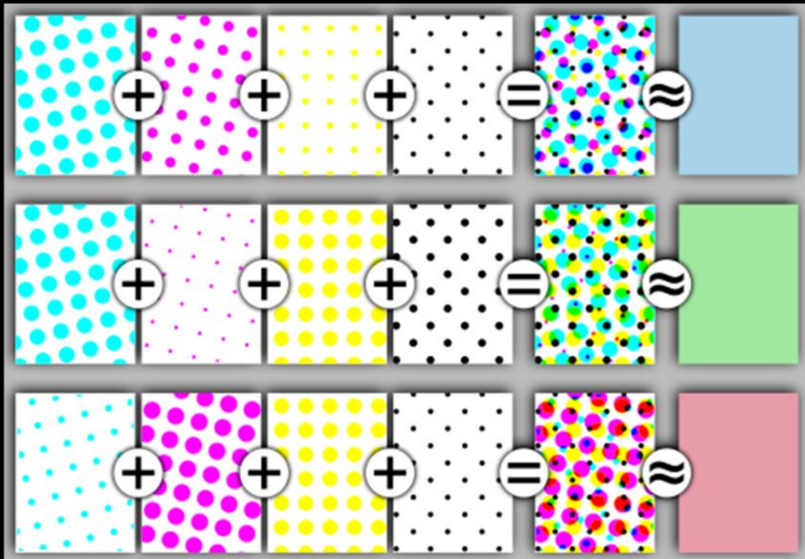

…printing technology…

(File “Halftoningcolor.png” from Wikimedia Commons used under public domain), (File “Bogenoffsetmaschine-2-1980.jpg” by Clemens Pfeiffer of Vienna used with permission)

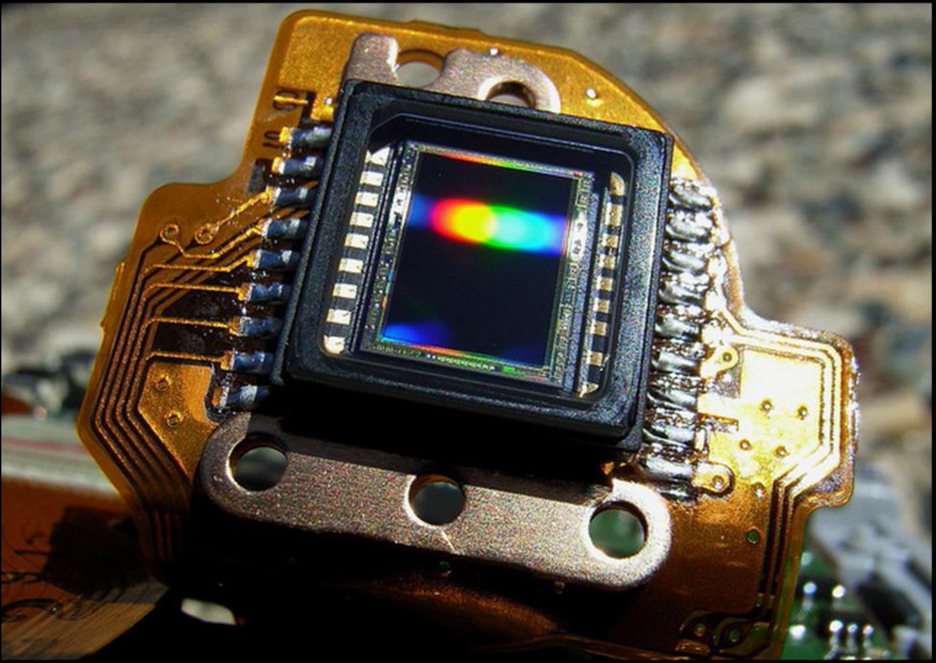

…electronic image capture & display…

(Image "Soul of a New Machine" by Flickr user Steve Jurvetson, used under CC BY 2.0), (Image of LG OLED TV by Flickr user "LGEPR", used under CC-BY-2.0 / Cropped from original)

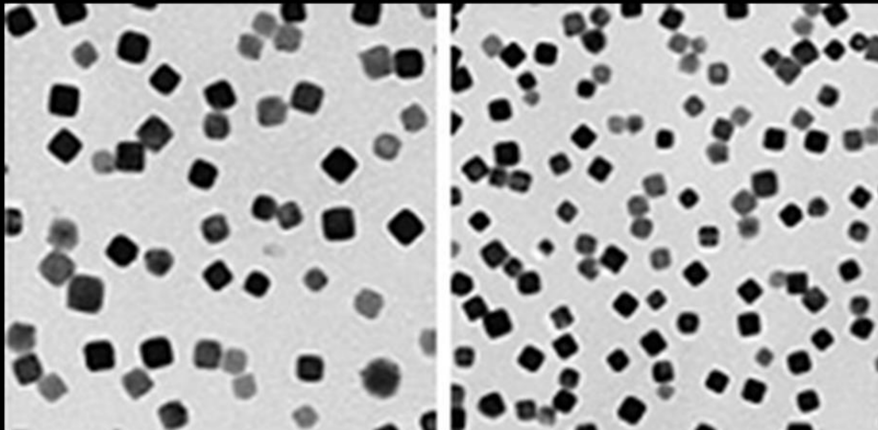

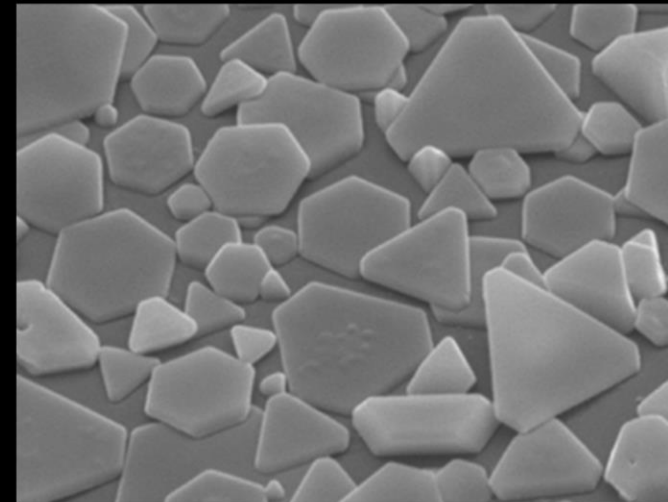

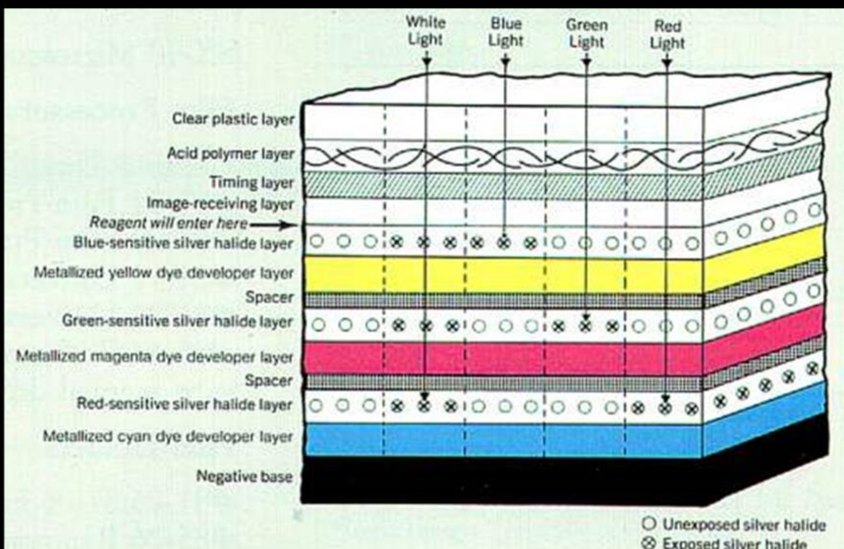

1972. A schematic of SX-70 film during exposure.

…and silver halide film emulsion.

(SX-70 film schematic from Kodak, used for scholarly commentary under "fair use"), (three microphotographs of silver halide crystals from Kodak, used for scholarly commentary under "fair use"), (Image "From silver halide to pixels" by Flickr user Mattias, used under CC-BY-2.0)

IS&T
imaging.org

SID
SOCIETY FOR INFORMATION DISPLAY

SPSTJ

isj

INTER - SOCIETY COLOR COUNCIL, INC.
FOUNDED 1931
SCIENCE • ART • INDUSTRY

They have their own professional societies…

…eminent researchers, dominant research institutions, journals, and prominent textbooks (I personally recommend “Digital Color Management”, I’ve learned a lot from it)…

(Images of researchers and book covers used for scholarly commentary under “fair use”)

… and of course, there are conferences specific to the field. I've been to CIC (the Color & Imaging Conference), and I would highly recommend it to people who are interested in learning more about color & imaging science. They have a very good courses program, in addition to the paper presentations & keynotes.

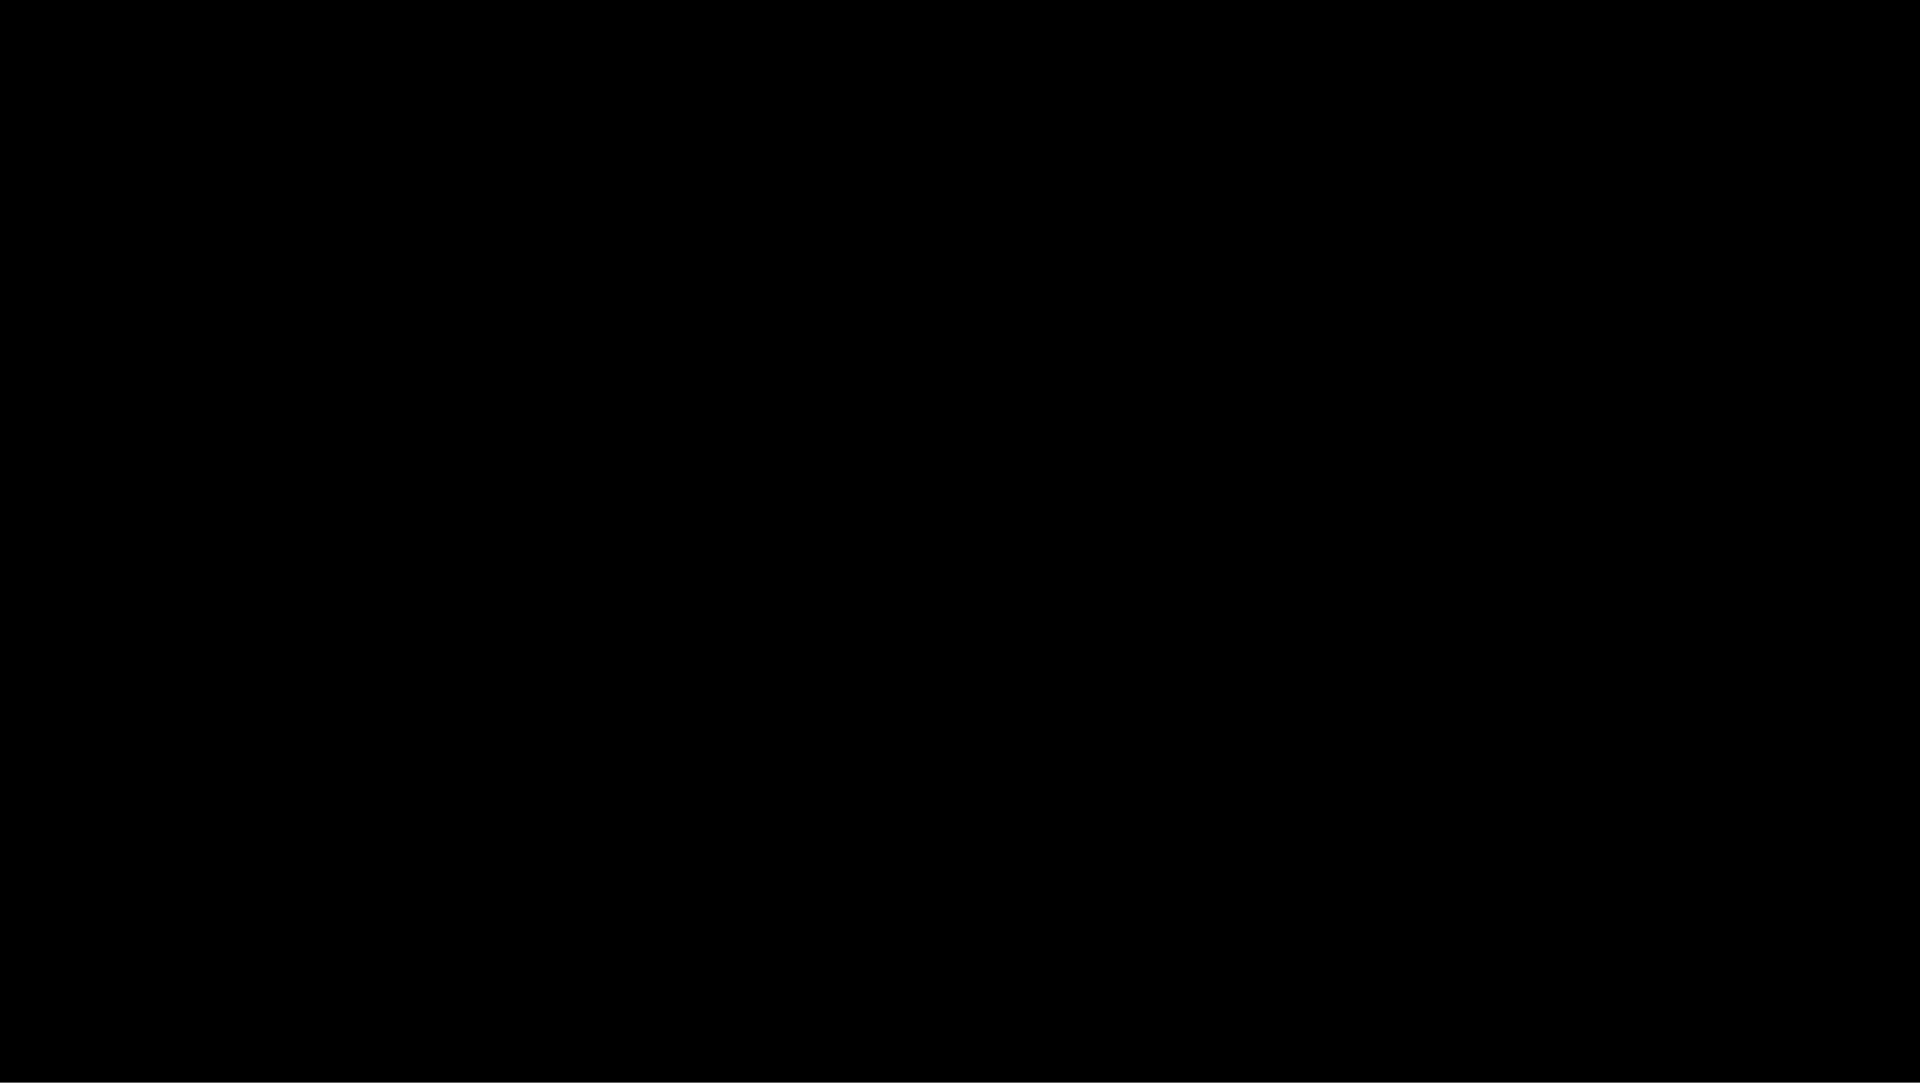

Some researchers (Like James Ferwerda) and practitioners (typically color experts at VFX & feature animation houses) straddle both worlds, both otherwise there’s not a lot of interaction between the two domains. Which is a pity because I think the central problem of imaging science…

Image Reproduction

...image reproduction – is especially important to computer graphics.

An important special case of image reproduction is reproducing the appearance of an actual scene (for example this sunlit meadow) onto a display device of some kind, which can be a…

(Image “Hillside” by Flickr user Nicholas A. Tonelli, used under CC-BY 2.0 / Cropped from original – note, this image is used on many of the following slides as well)

…computer monitor in an office,…

(Image “From the office” by Flickr user Jan Kraus, used under CC-BY 2.0 / Cropped from original and “Hillside” image pasted into frame)

…a television set in a living room,…

(Image “living room post-rearrangement” by Flickr user anneheathen, used under CC-BY 2.0 / Cropped from original and “Hillside” image pasted into frame)

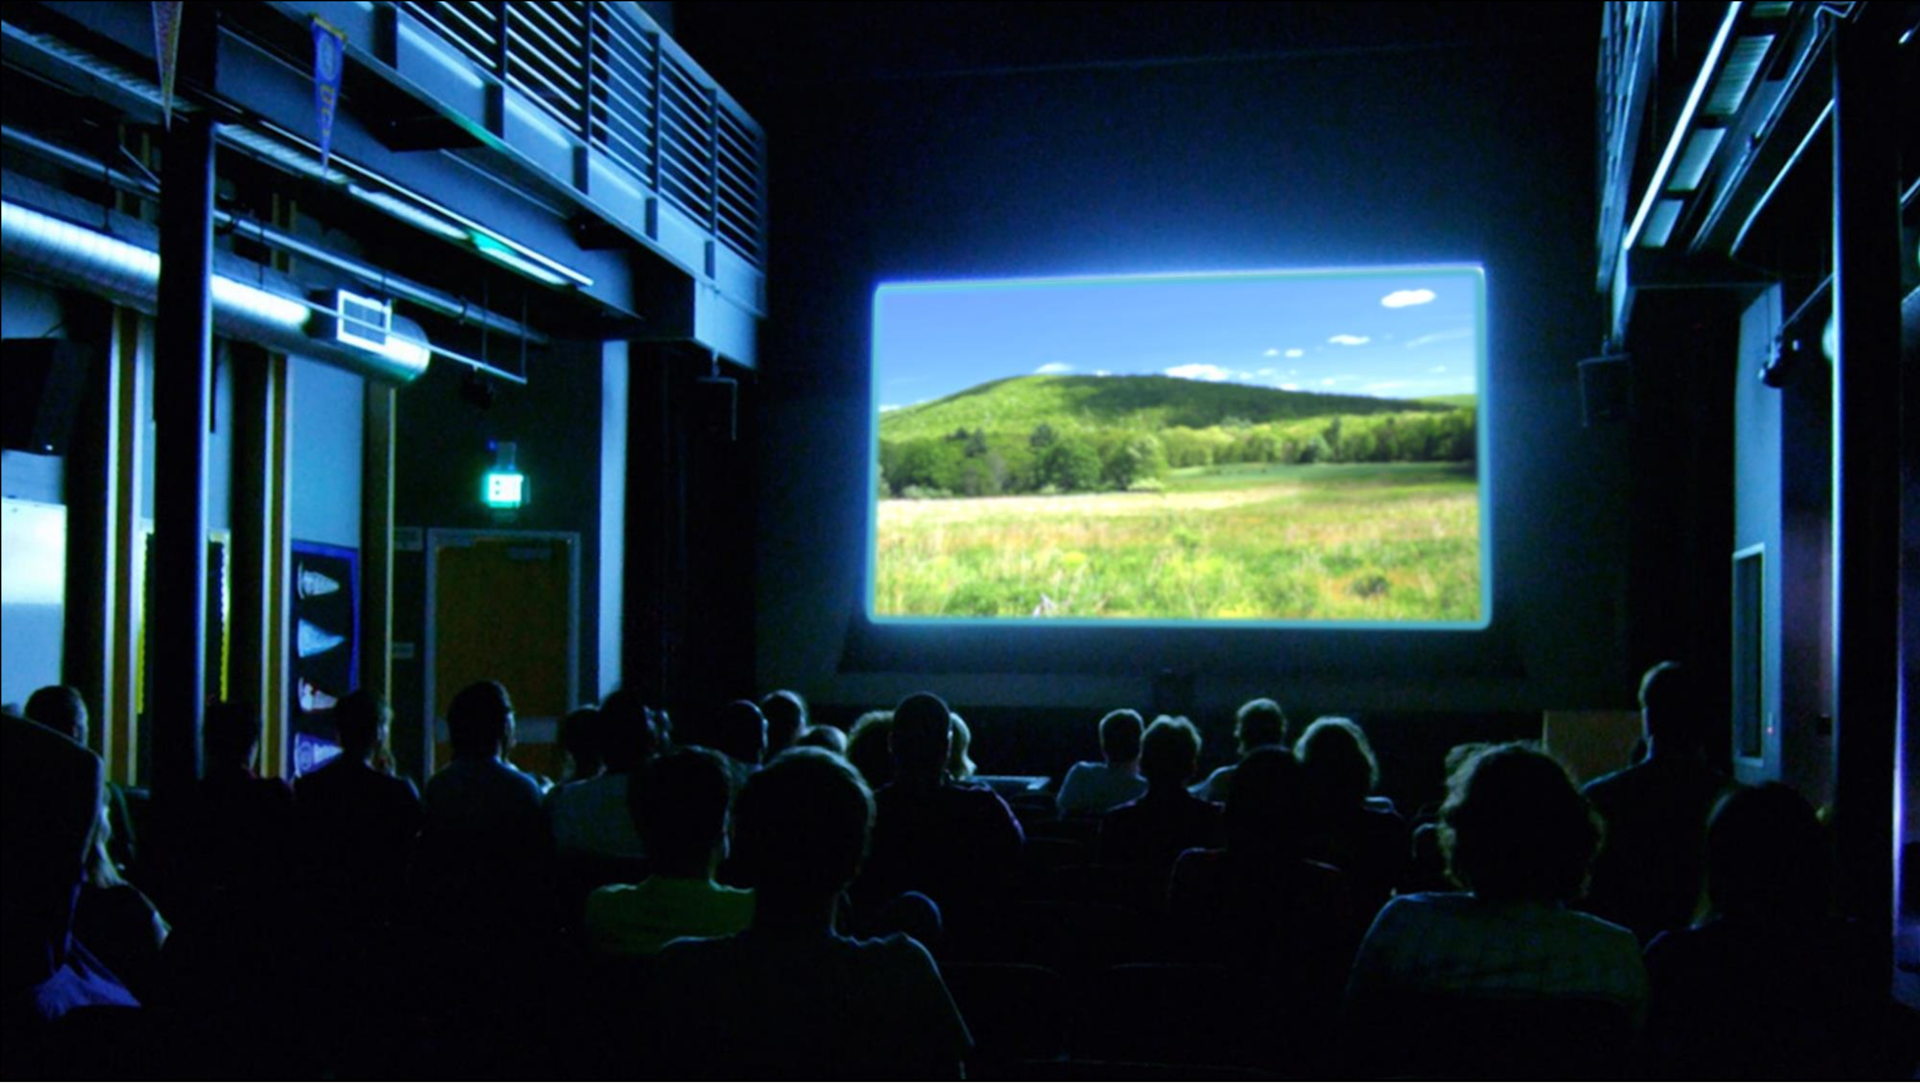

…a projection screen in a movie theater,…

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

…a framed photographic print on a desk, or one of many other possibilities.

A key insight is that image reproduction doesn't make sense without the context of a well-defined…

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

(image of standing 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.

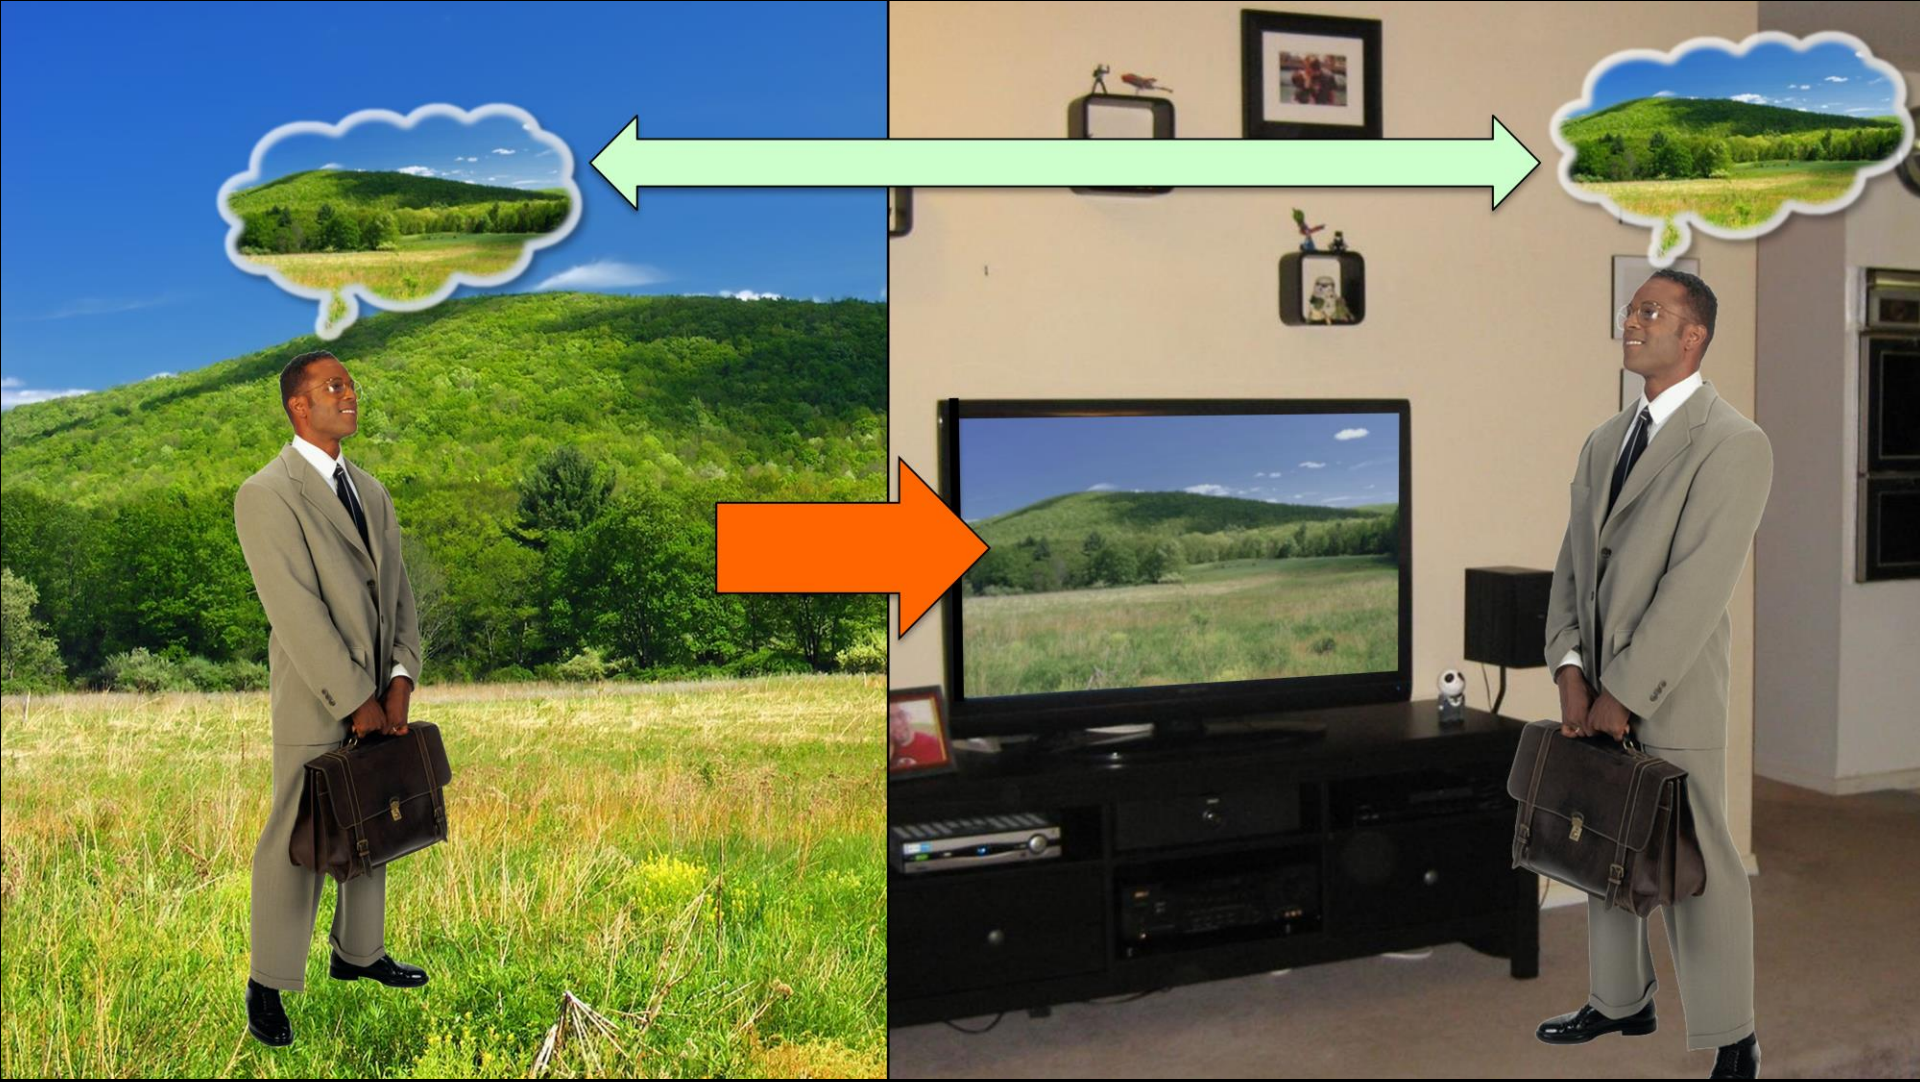

In other words, you need to try to create an image on the final display such that the viewer is getting the same (or as close as possible) perceptual impression as if he or she were present in the original scene. If you think about it, it’s surprising that this is possible to any degree, considering that there are typical several orders of magnitude difference between the absolute light intensities and contrasts in the original scene and the ones which 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.

$$
\begin{aligned}
L_r(x,\omega_r) &= L_e(x,\omega_r) + \int_\Omega L_i(x,\omega_i) f(x,\omega_i,\omega_r)\cos\theta_i d\omega_i \\
&= L_e(x,\omega_r) + \int_\Omega L_r(x',-\omega_i) f(x,\omega_i,\omega_r)\cos\theta_i d\omega_i
\end{aligned}
$$

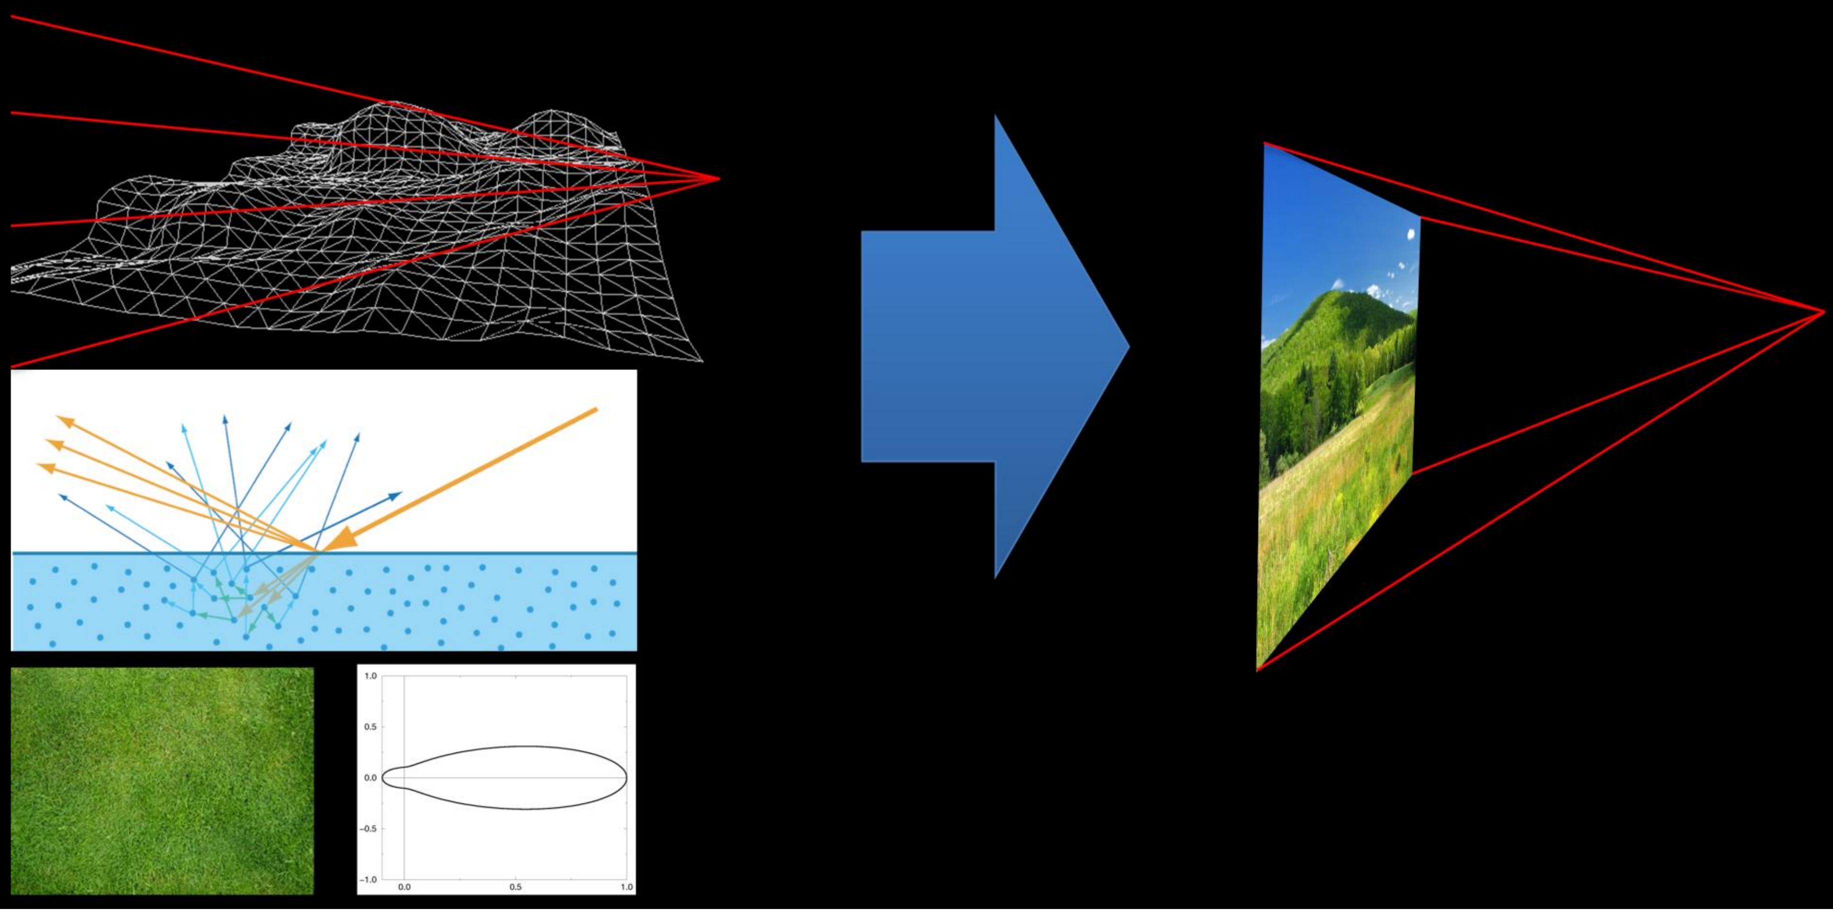

Thus defined, image reproduction is a very important – but typically neglected - part of the process of creating computer-generated imagery. In CG, we have a virtual scene with meshes, high order surfaces, physically based surface shaders, texture maps, global illumination algorithms, volume shaders for participating media, etc. – and we try very hard to compute the radiometric values that would have entered our virtual camera in our virtual scene if the scene and camera were real. But that's only half the job!

(mesh, grass texture and phase function images used for scholarly commentary under "fair use")

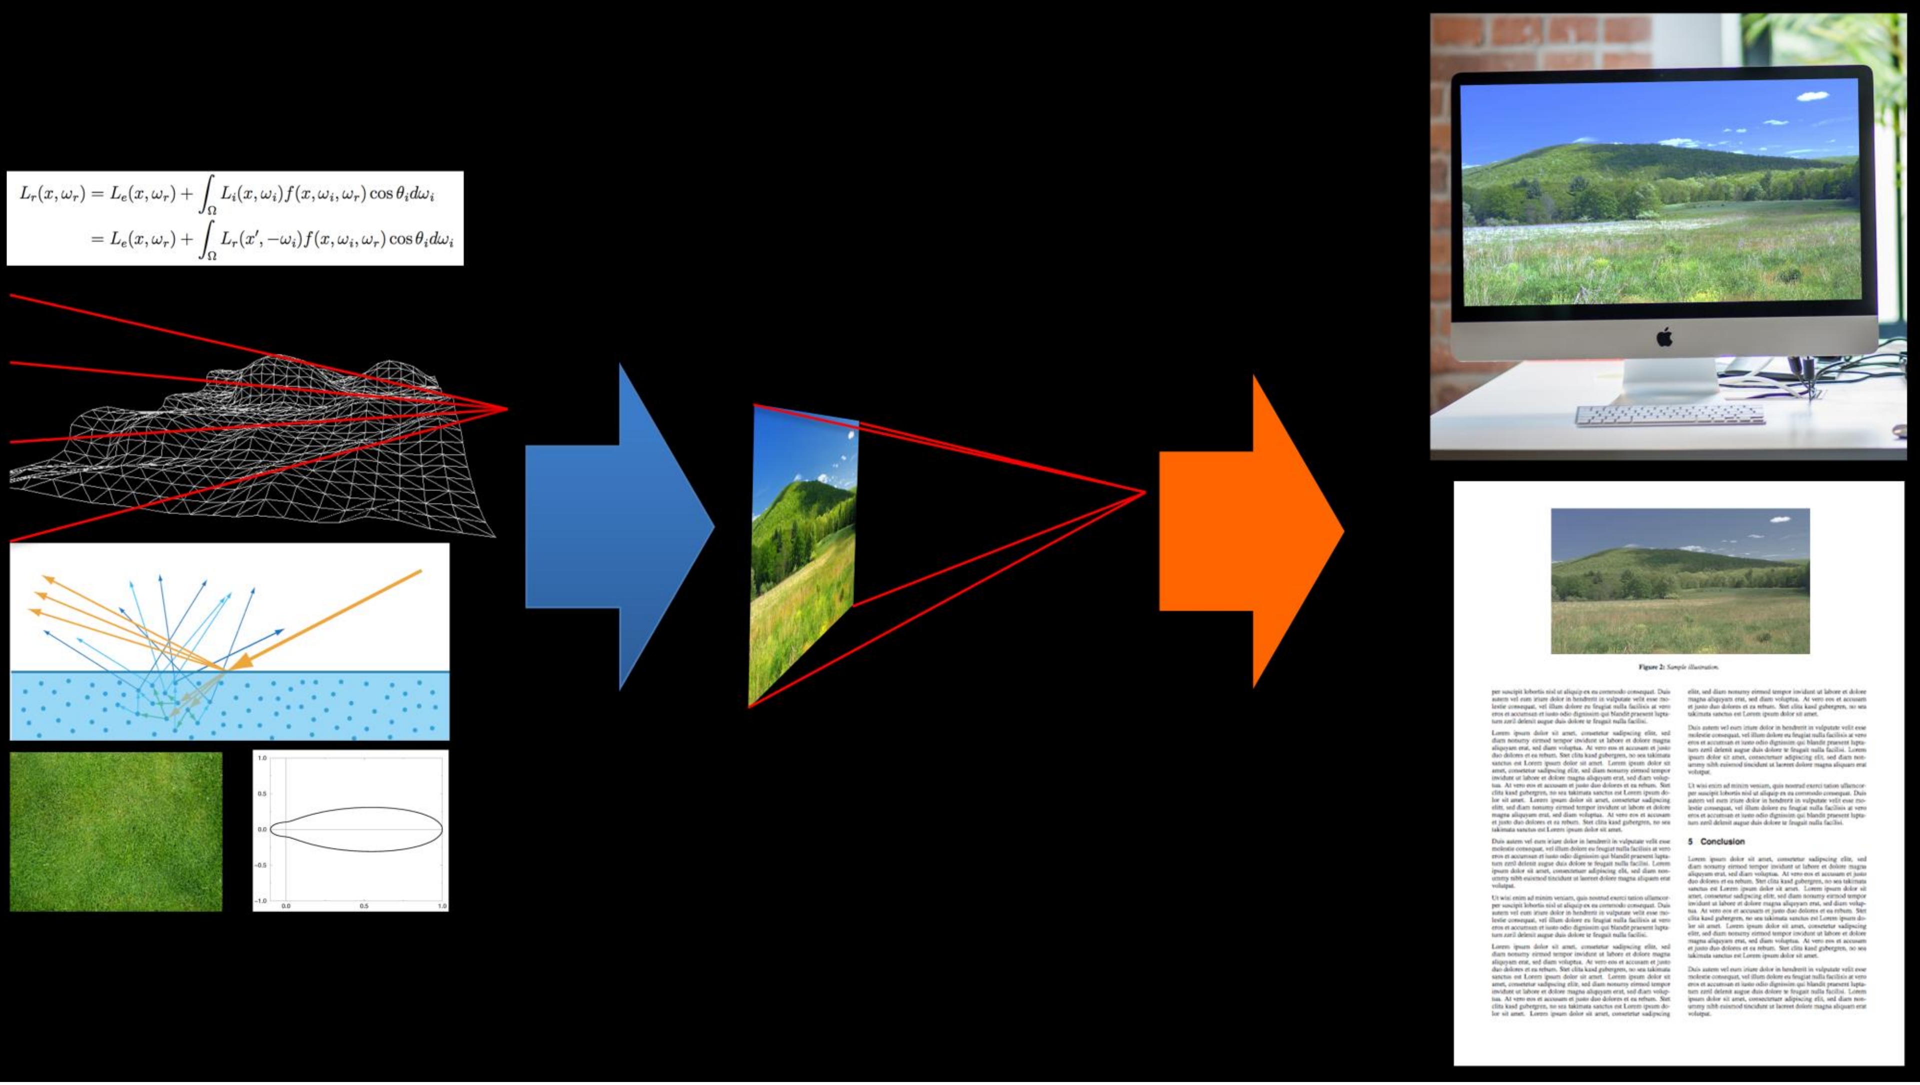

Our final output isn't a bunch of radiometric scene values, it's an image for some display device which is viewed under particular viewing conditions, whether it's a computer screen in an office, or a printed page in conference proceedings. The steps that take the scene values to the final output tend to get a lot less attention. Graphics practitioners working in film & TV do typically have well-reasoned processes for this step, but that's less often the case in the game industry and among graphics researchers.

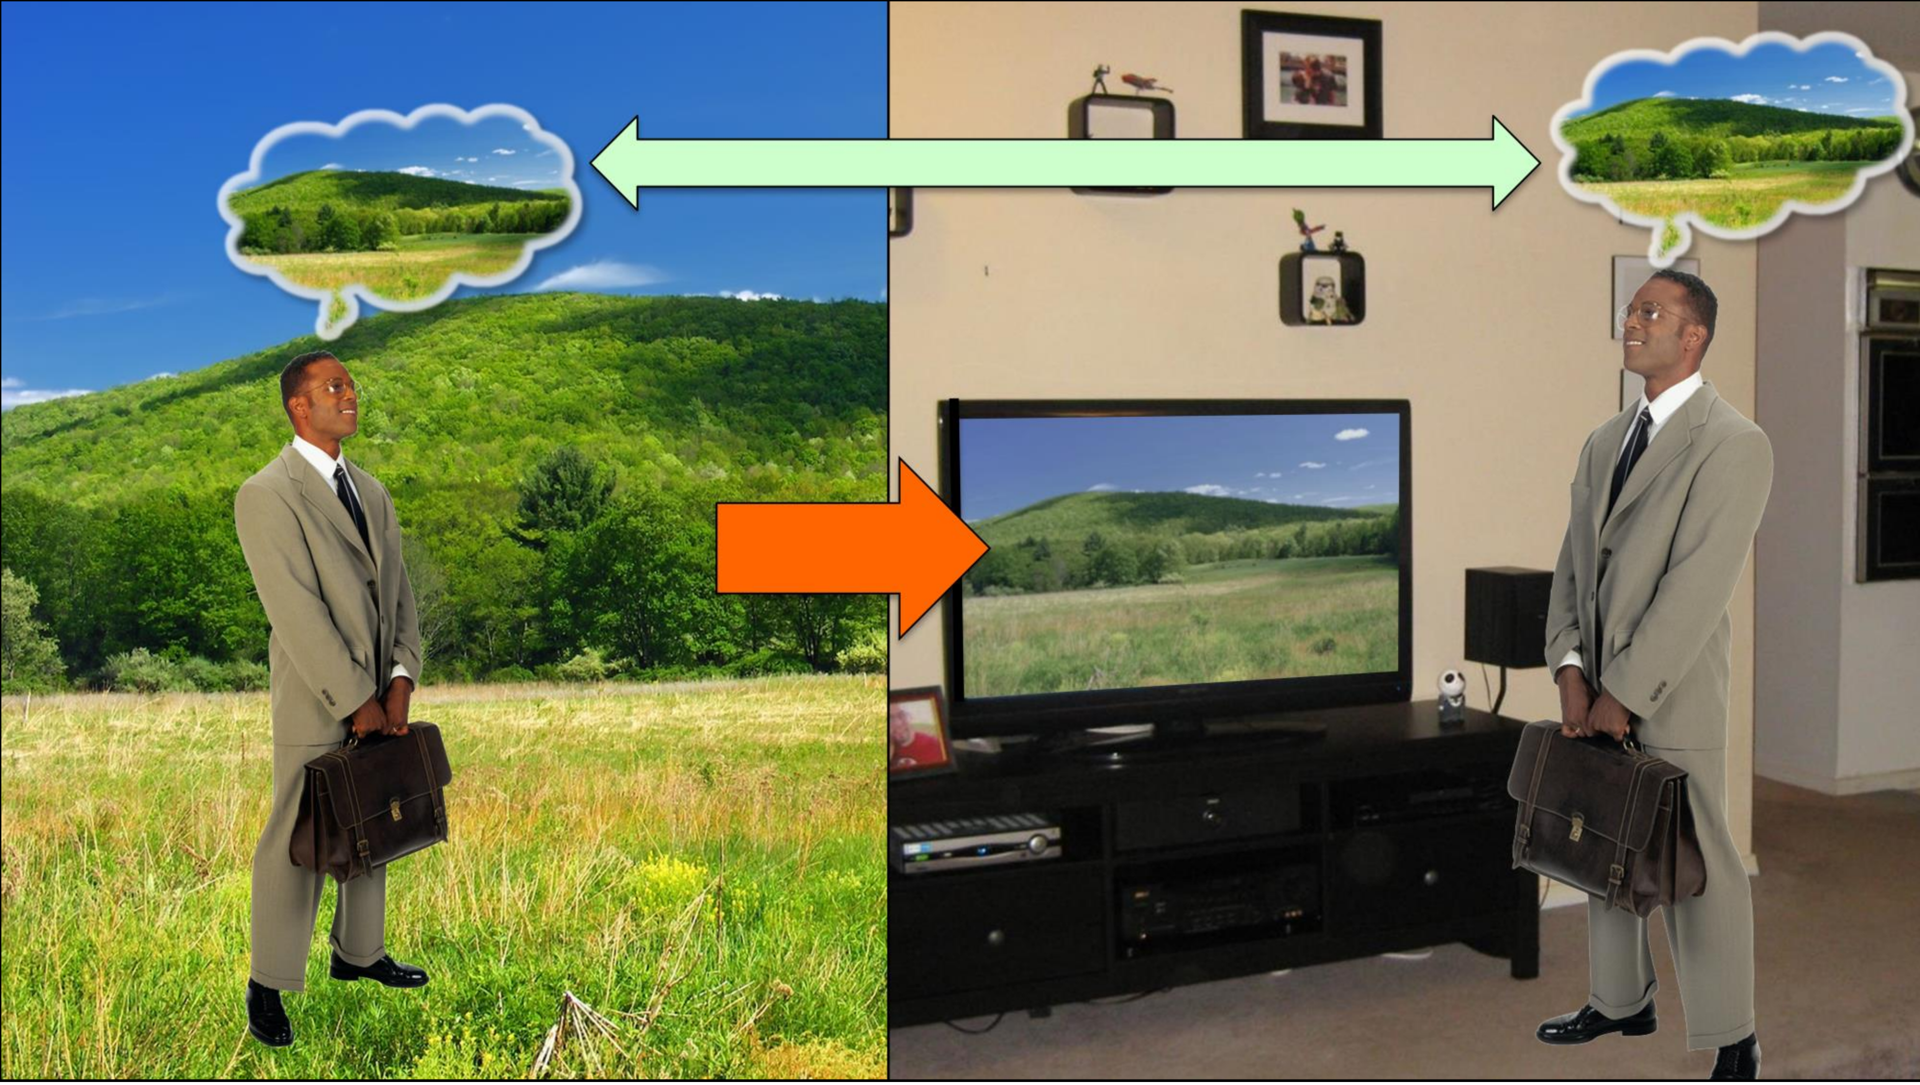

In my opinion, thinking about this problem in the way that imaging people would think about it can be very helpful – in our community we tend to think of this step a bit nebulously as “tone mapping” or “HDR-to-LDR conversion”. And it helps to realize that this problem is not a new one…

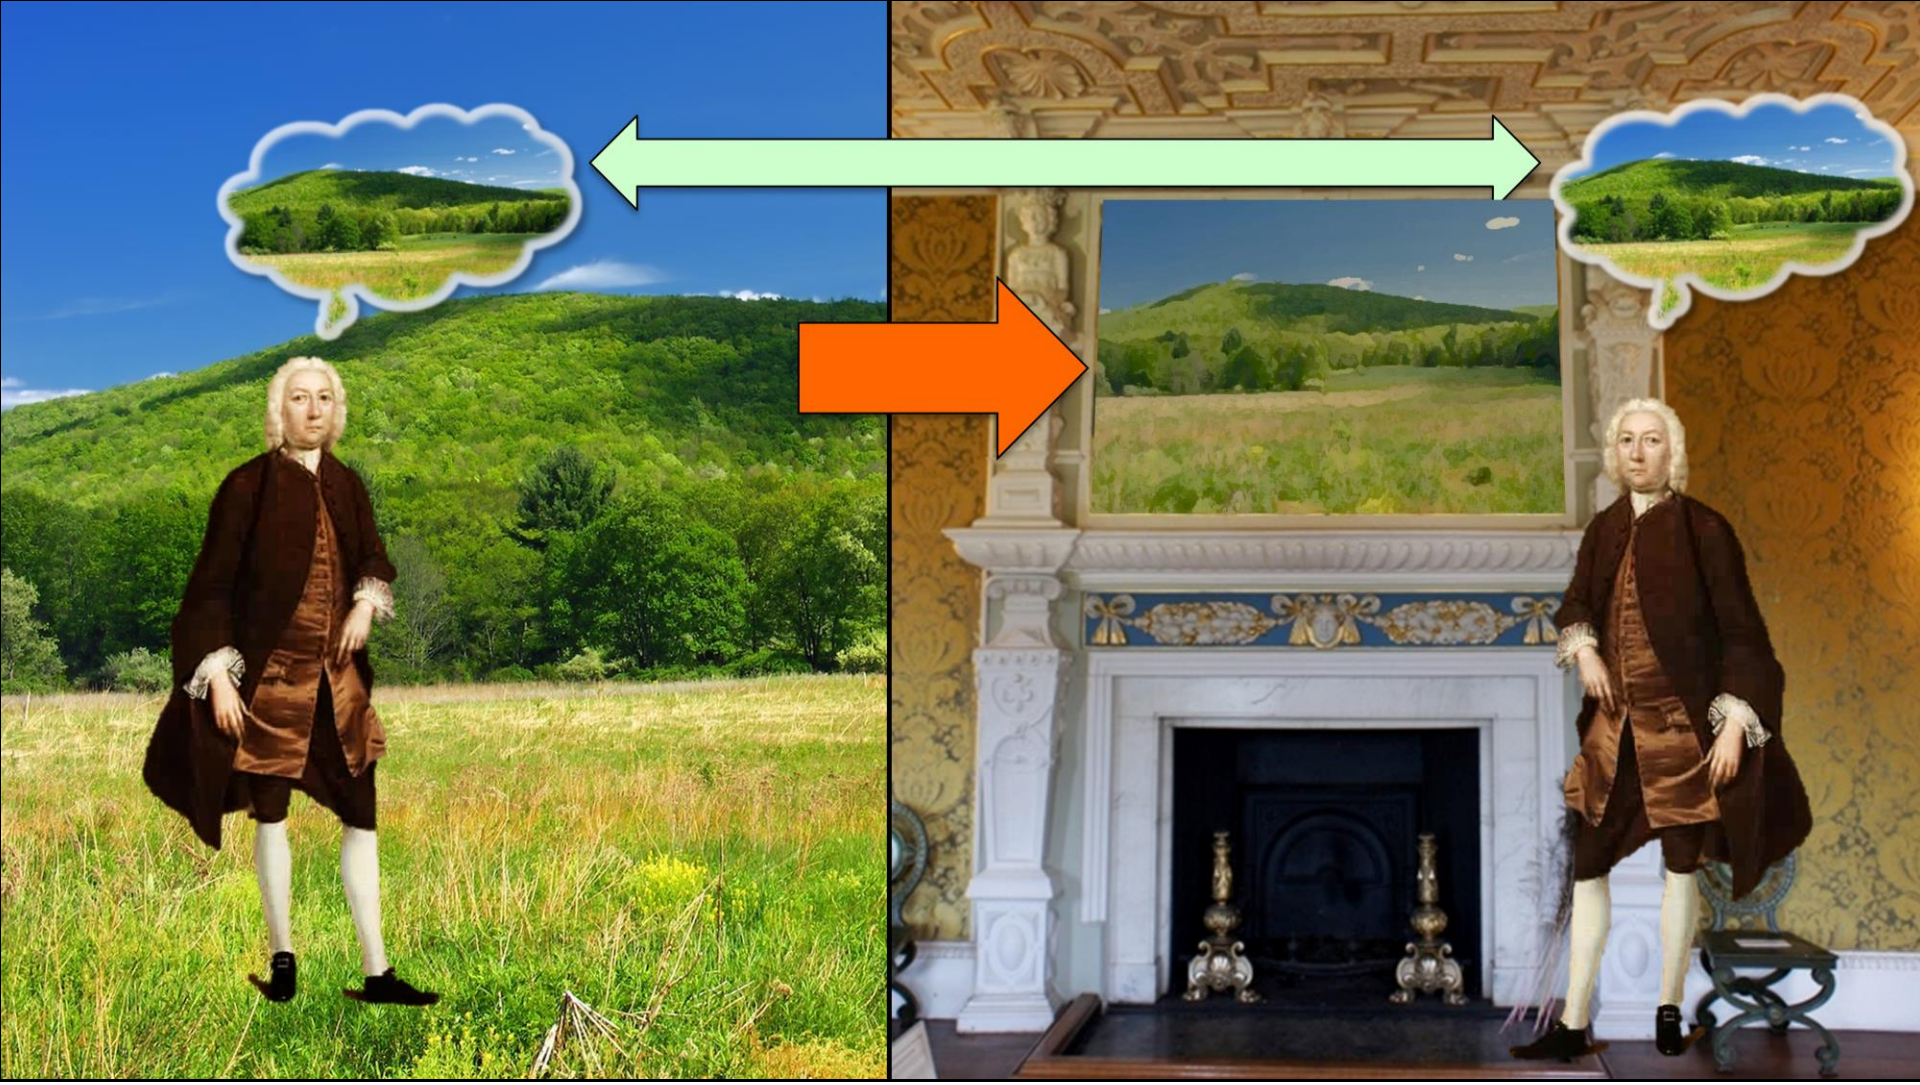

…quite the opposite. Renaissance painters had to solve it too.

(“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 and “Hillside” image pasted into frame)

As we can see in this painting - “The Artist in His Studio” by Rembrandt – the Renaissance painters were able to effectively address the image reproduction problem, with the use of techniques such as chiaroscuro. This painting gives a sense of extreme contrast, far beyond what the paint & canvas actually allow. This slide and the next four are from Josh Pine’s excellent talk in the 2010 SIGGRAPH color course, used with his permission.

(“The Artist in His Studio” by Rembrandt van Rijn, ~1628, public domain)

“The Night Watch”, done by Rembrandt 14 years later, shows the perfection of the chiaroscuro technique.

(“The Night Watch” by Rembrandt van Rijn, ~1642, public domain)

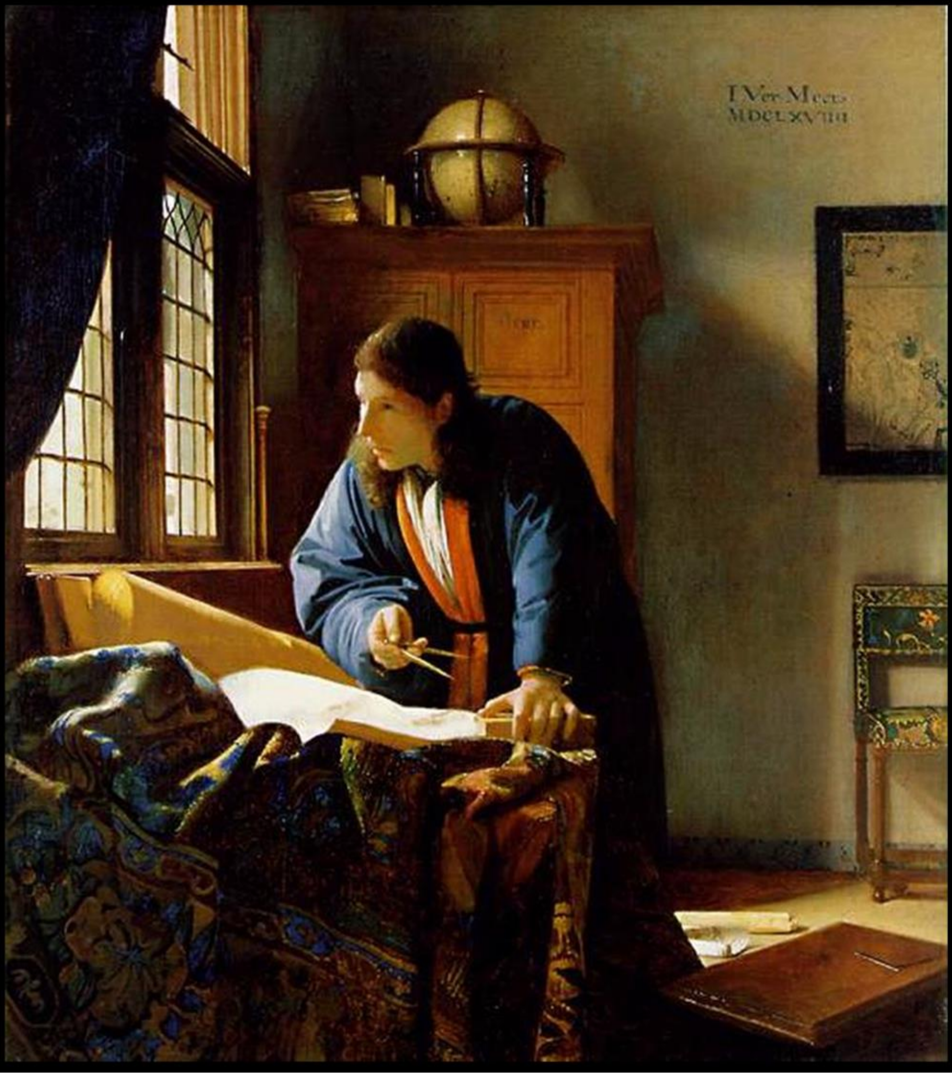

“The Geographer” by Vermeer – painted 27 years after “The Night Watch” – shows a similar command of light and shadow.

(“The Geographer” by Johannes Vermeer, ~1669, public domain)

Skipping over almost 300 years of art history brings us to Edward Hopper, and his well-known “Nighthawks”. Note the impression of extreme contrast between the lit diner interior and the night-time street outside.

(“Nighthawks” by Edward Hopper, 1942, used for scholarly commentary under “fair use”)

Another painting by Hopper – “Summer Evening”, painted in 1947. Painters solved the image reproduction problem by the use of instinct, careful observation, and training. However, by this time a more mechanical method of solving the same problem had become quite prevalent, and was rising in importance as an artform – causing painting to go in a decidedly non-representational direction.

(“Summer Evening” by Edward Hopper, 1947, used for scholarly commentary under “fair use”)

Photography

Photography used technology to solve the image reproduction problem – of course a lot of skill was involved. Not only that of the photographers themselves, but also of the engineers and perceptual psychologists who had carefully engineered this technology to perform good image reproduction. This technology (and the artistry employed in using it) also has some lessons to teach us.

Unlike painting, we expect photography to perform correct image reproduction – in the sense shown here, defined by imaging science – as a matter of course, by default. But this actually took a lot of work.

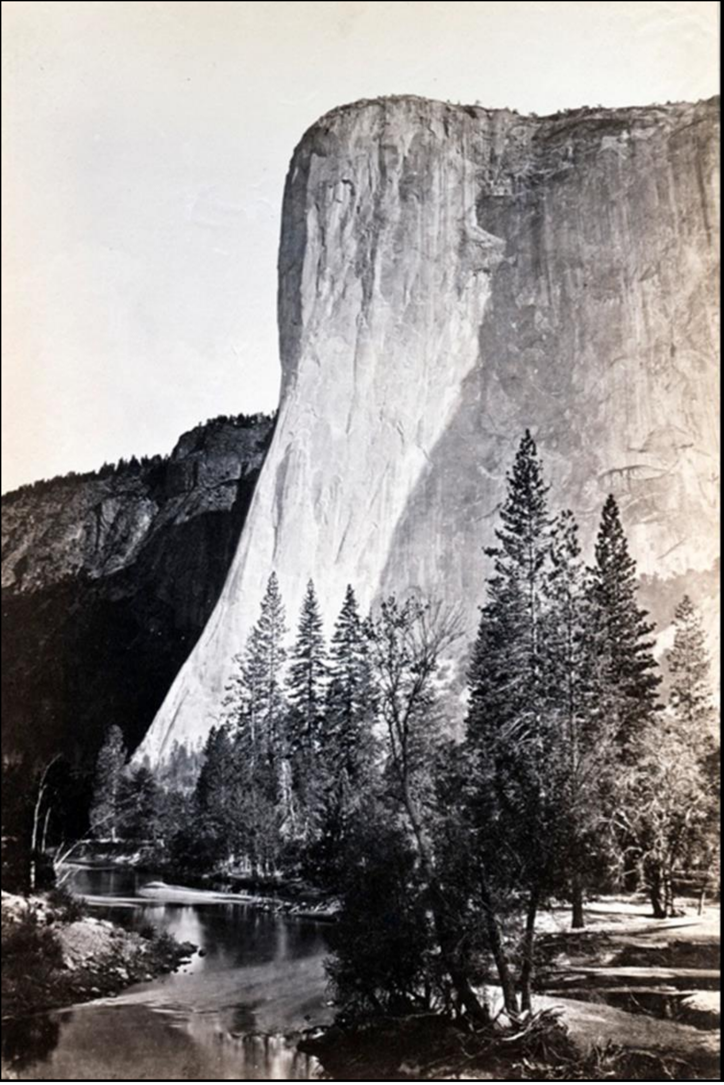

1866
Carleton
Watkins

1956
Ansel
Adams

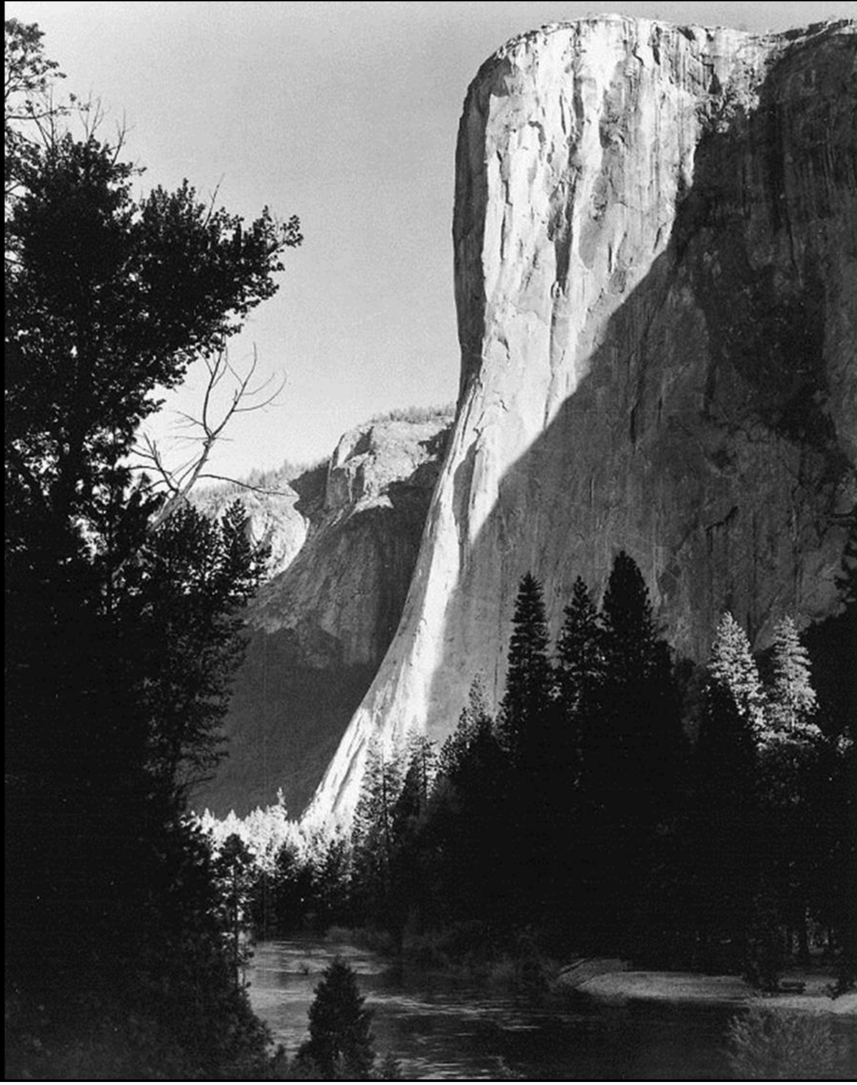

To illustrate the advances in photographic technology, let's look at this pair of photographs. Not only are they both of the same subject (El Capitan at Yosemite National Park), they are photographed from very similar angles and lighting conditions. Furthermore, each of these photographs was taken by the most prominent landscape photographer of that time, both of whom were best-known for their photographs of the American West and of Yosemite in particular. To help in comparison, I'll remove the distracting brownish tint from the left photograph...

("El Capitan" by Carleton Watkins, c1866, public domain). ("El Capitan, Sunrise", by Ansel Adams, c1956, used for scholarly commentary under "fair use")

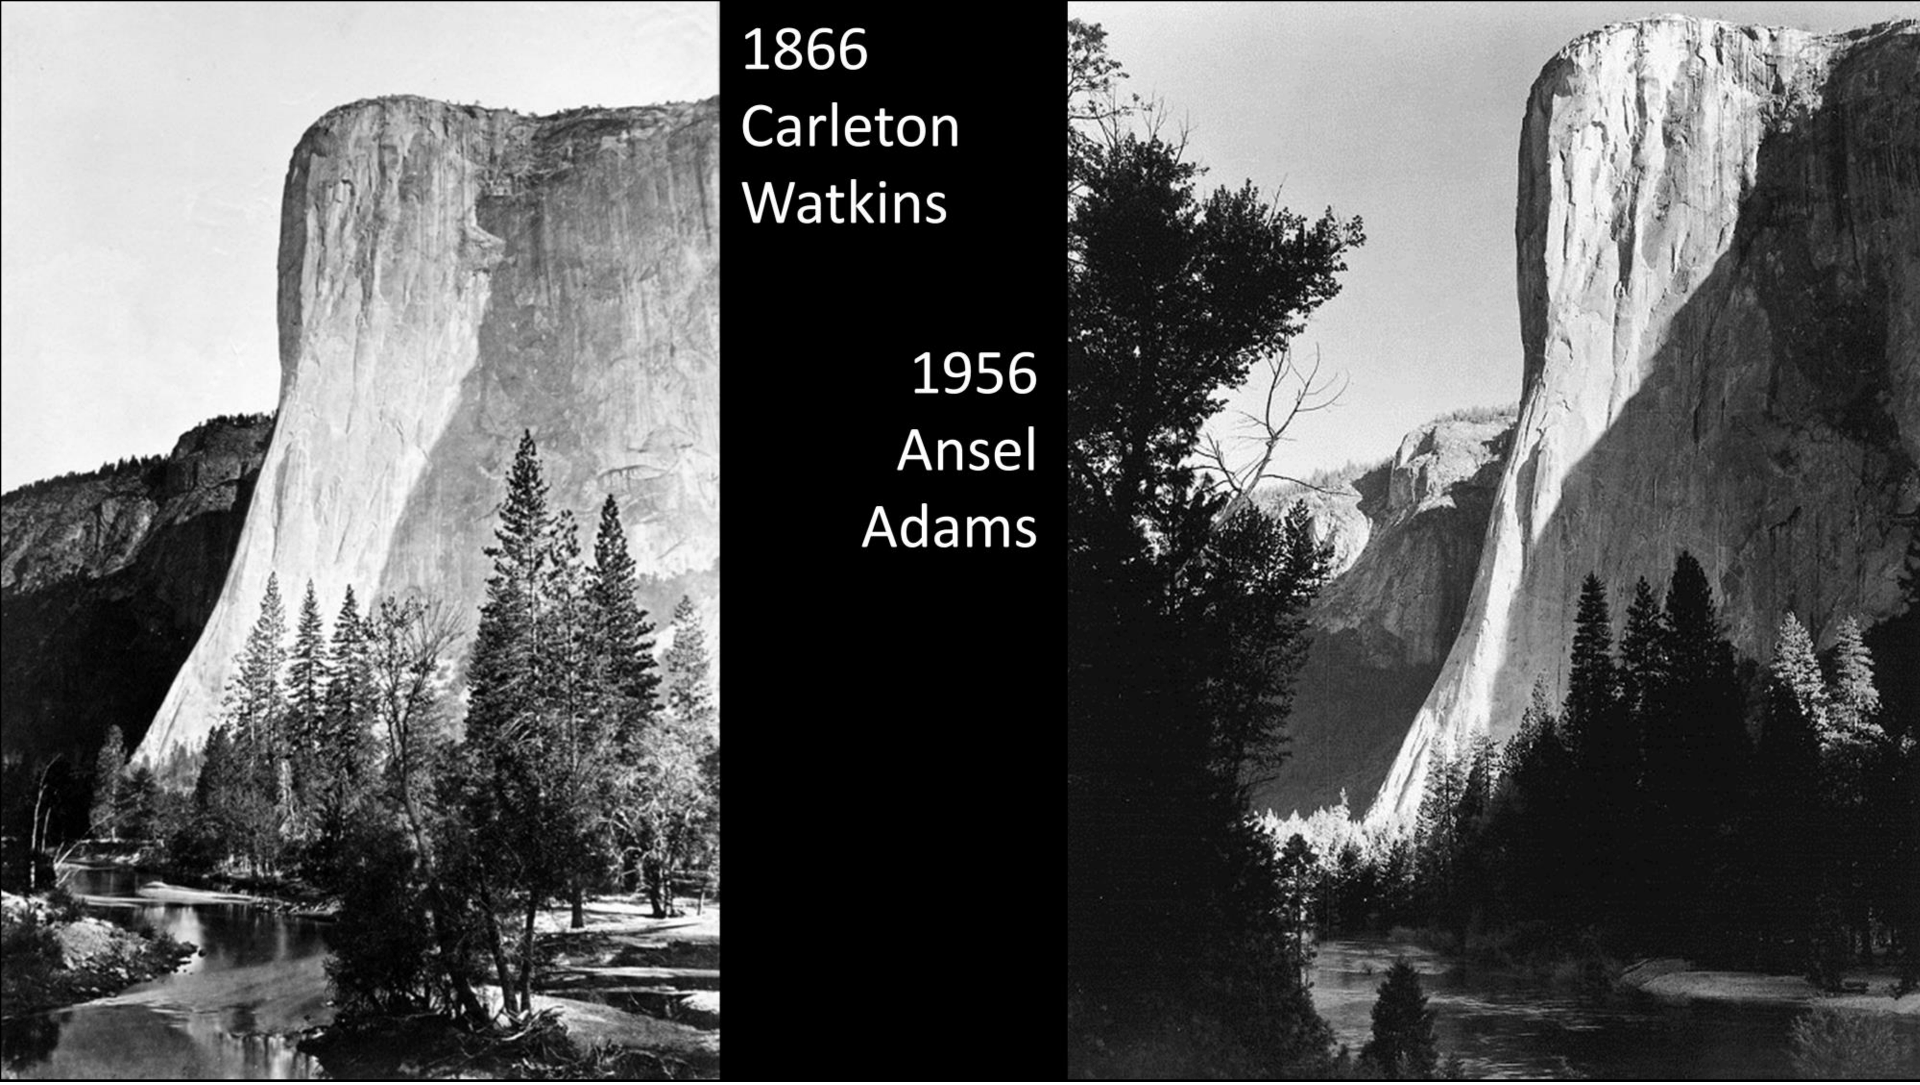

…and it’s clear that the photograph on the right looks more like a real scene. Note that these were each taken by the greatest landscape photographer of the time, using the best available equipment and every ounce of their considerable skill on a subject with which they were each extremely familiar. The improvement in image reproduction is due to 90 years’ advance in photographic technology (which had other benefits; Watkins’ equipment needed 12 mules to carry it, while Adams’ equipment fit in his backpack).

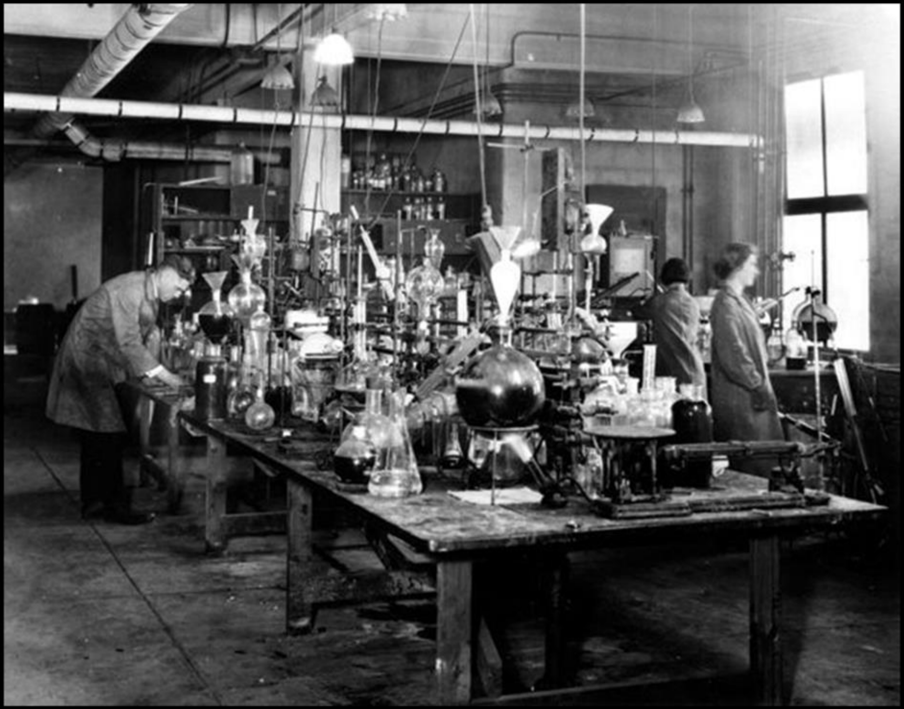

Kodak Research Laboratories, 1920

Many of these technological advances were developed at Kodak Research Laboratories, which in 1916 was one of the first industrial research facilities to be founded in the US. The Kodak labs focused a large part of their efforts on the problem of photographic image reproduction.

(Image of Kodak Research Laboratories, 1920, public domain).

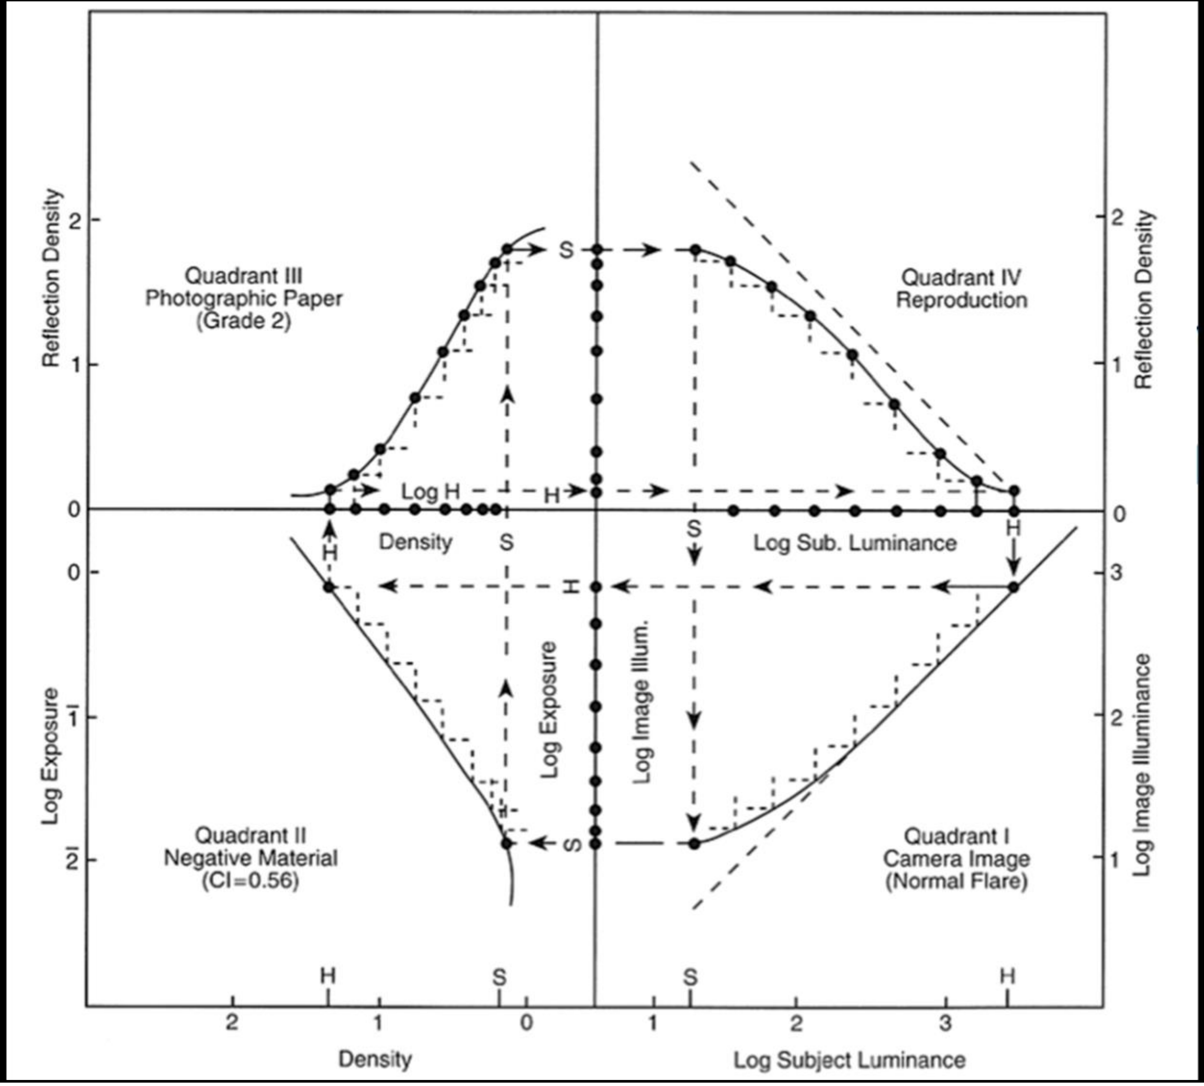

To illustrate the results of their work, we will look at this "Jones Diagram", which shows how the various stages of the photographic process concatenate to produce a global tone reproduction function. All the axes are log-scale, which is useful since visual perception is approximately logarithmic. Slope corresponds roughly to contrast increase / decrease (e.g. a slope of 1 preserves contrast, a slope higher than 1 increases it and a slope lower than 1 decreases it).

(Jones Diagram figure used for scholarly commentary under "fair use")

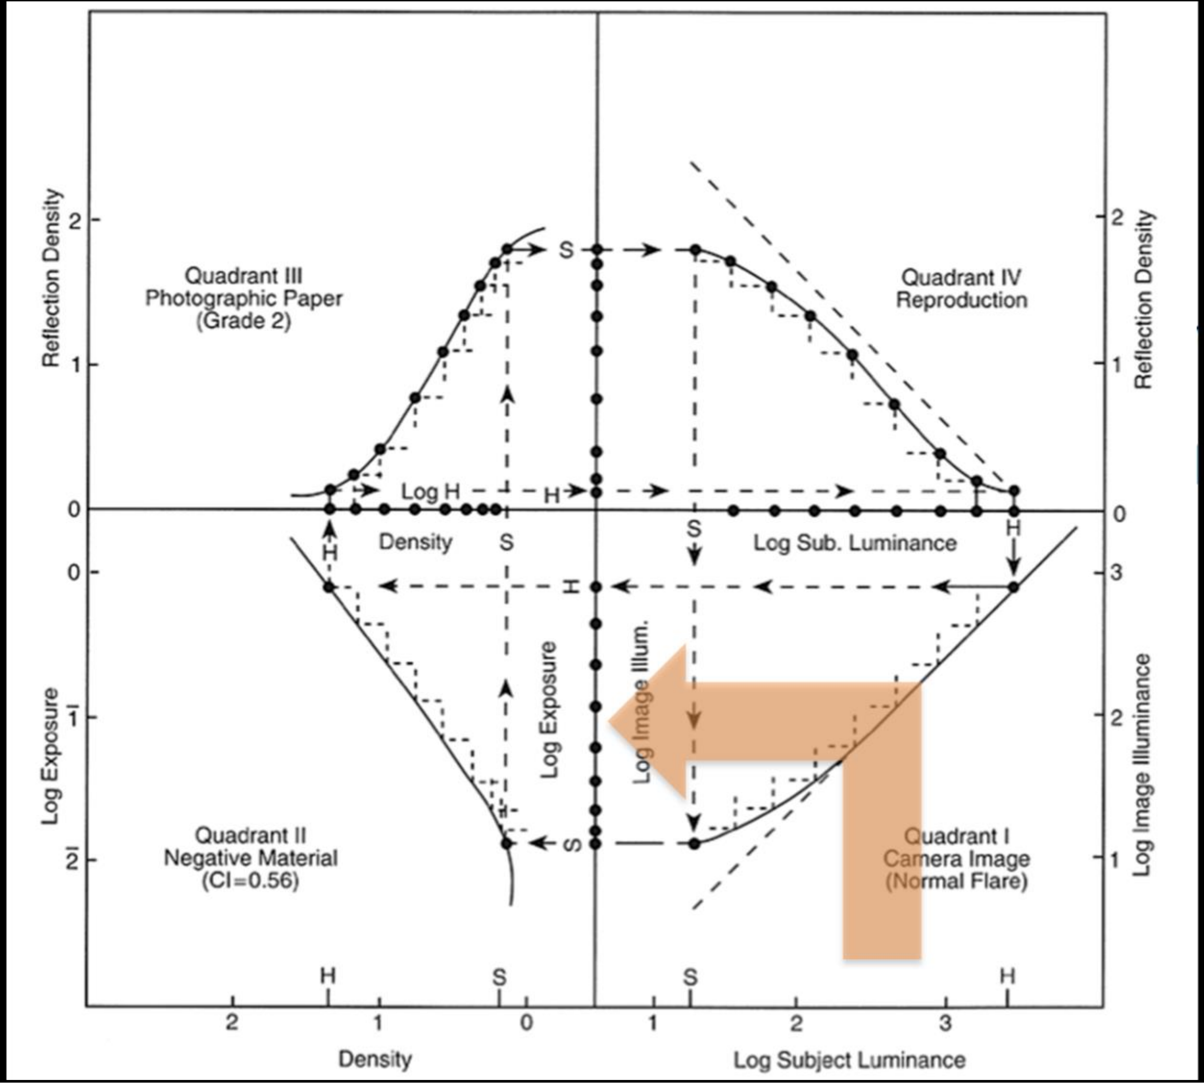

Light coming off the subject gets some flare added (due to imperfections in the lens and/or camera) before it reaches the image plane. Although this is a simple addition, on a log-log graph it dstorts the curve – the slope decrease for low luminance values implies that dark colors are reduced in contrast significantly but others are almost unaffected.

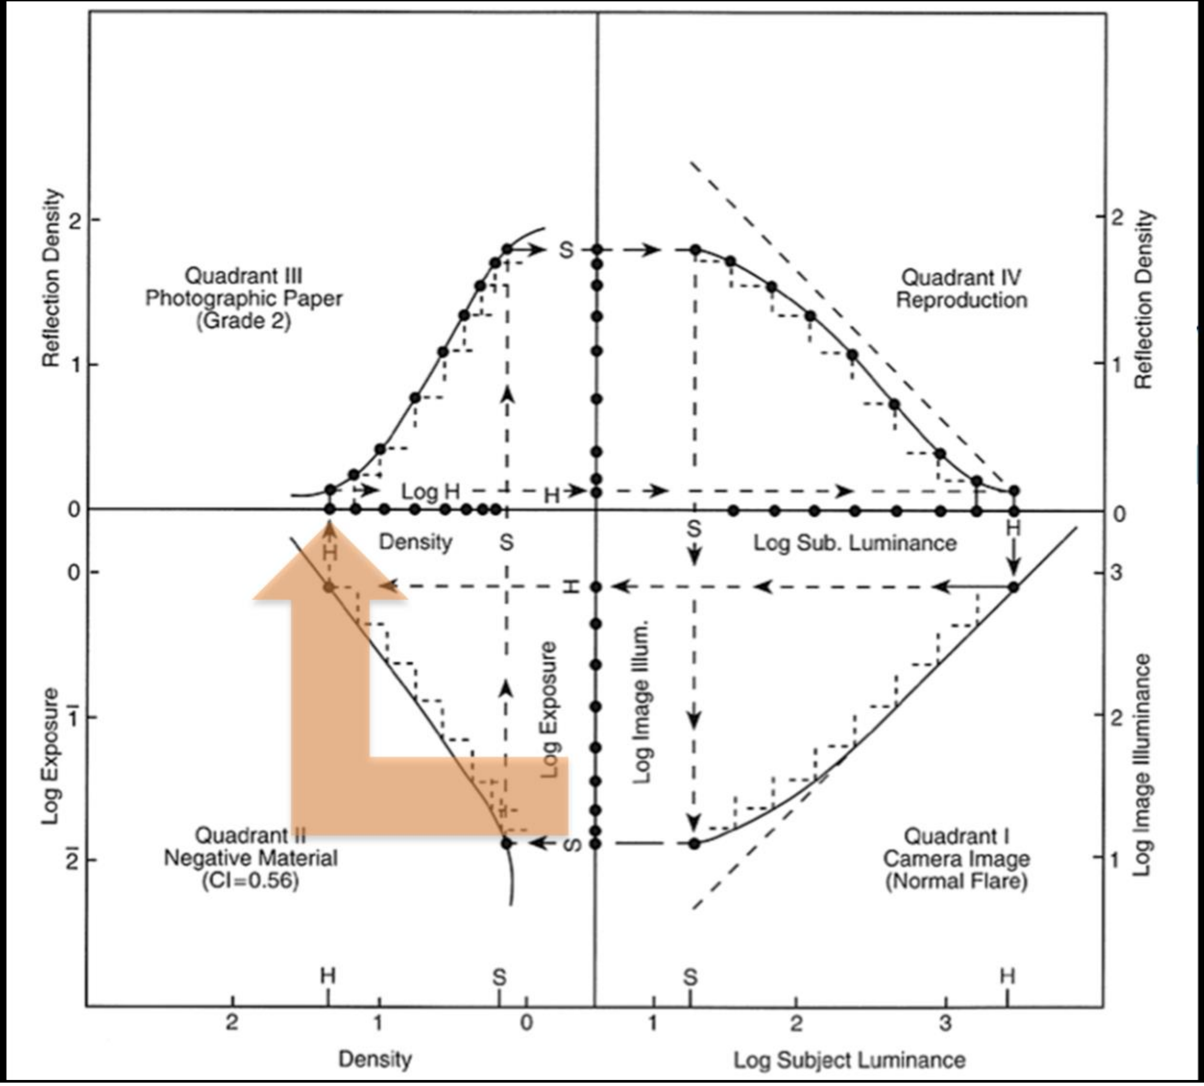

The negative changes density is response to the image plane exposure – this is approximately a straight power relationship (straight line on a log-log graph) though it tails off for the dark colors.

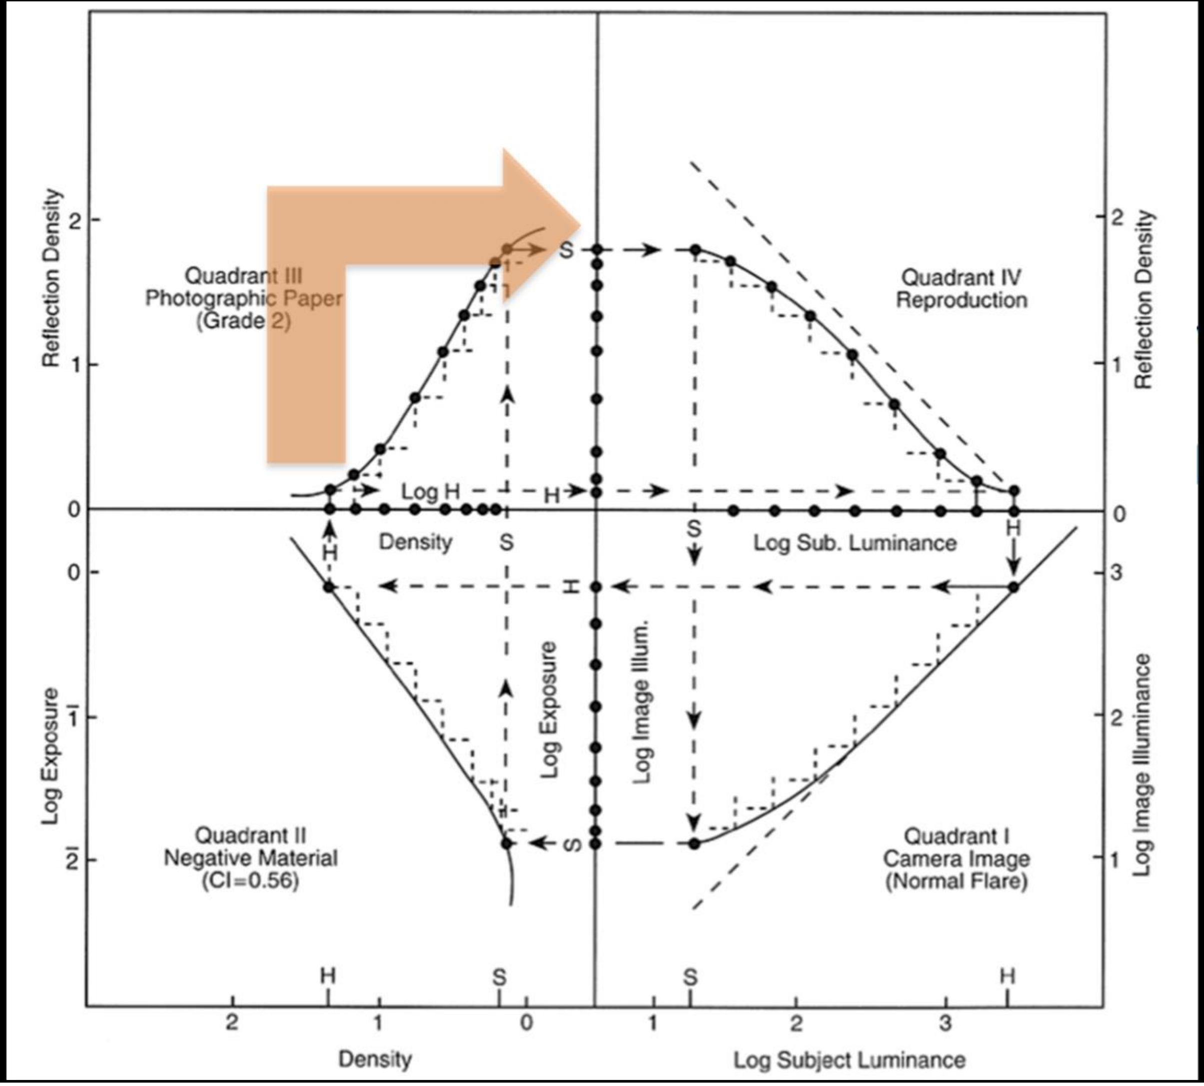

During the printing process, the negative density has a complex non-linear (even on the log-log graph) relationship with the reflection density of the paper (e.g. log reflectance, or log brightness of the illuminated print).

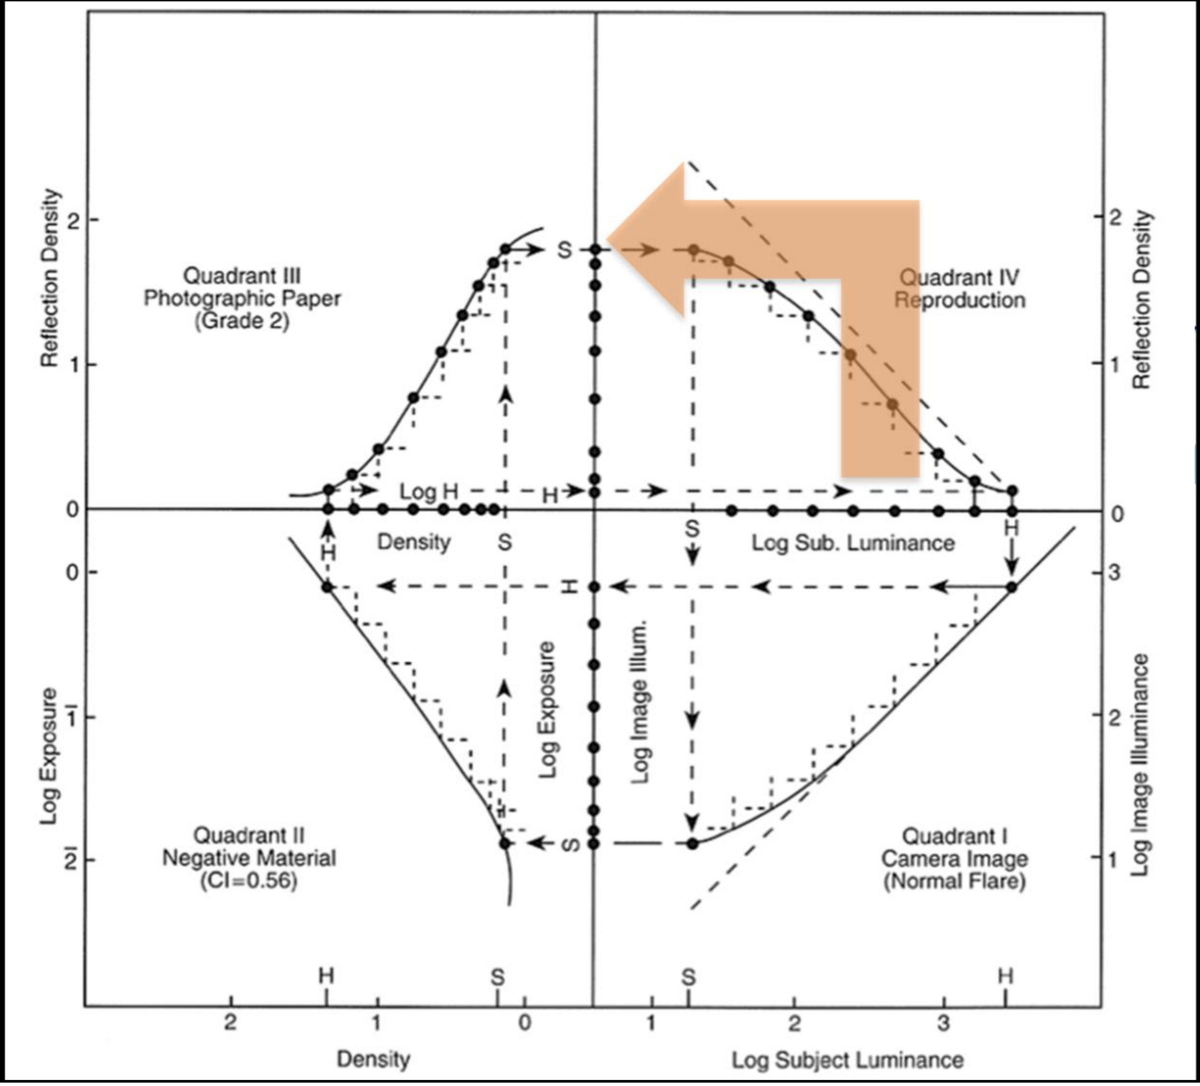

The concatenation of these three curves forms the full image reproduction curve – the curve relating scene brightness to display (print, in this case) brightness.

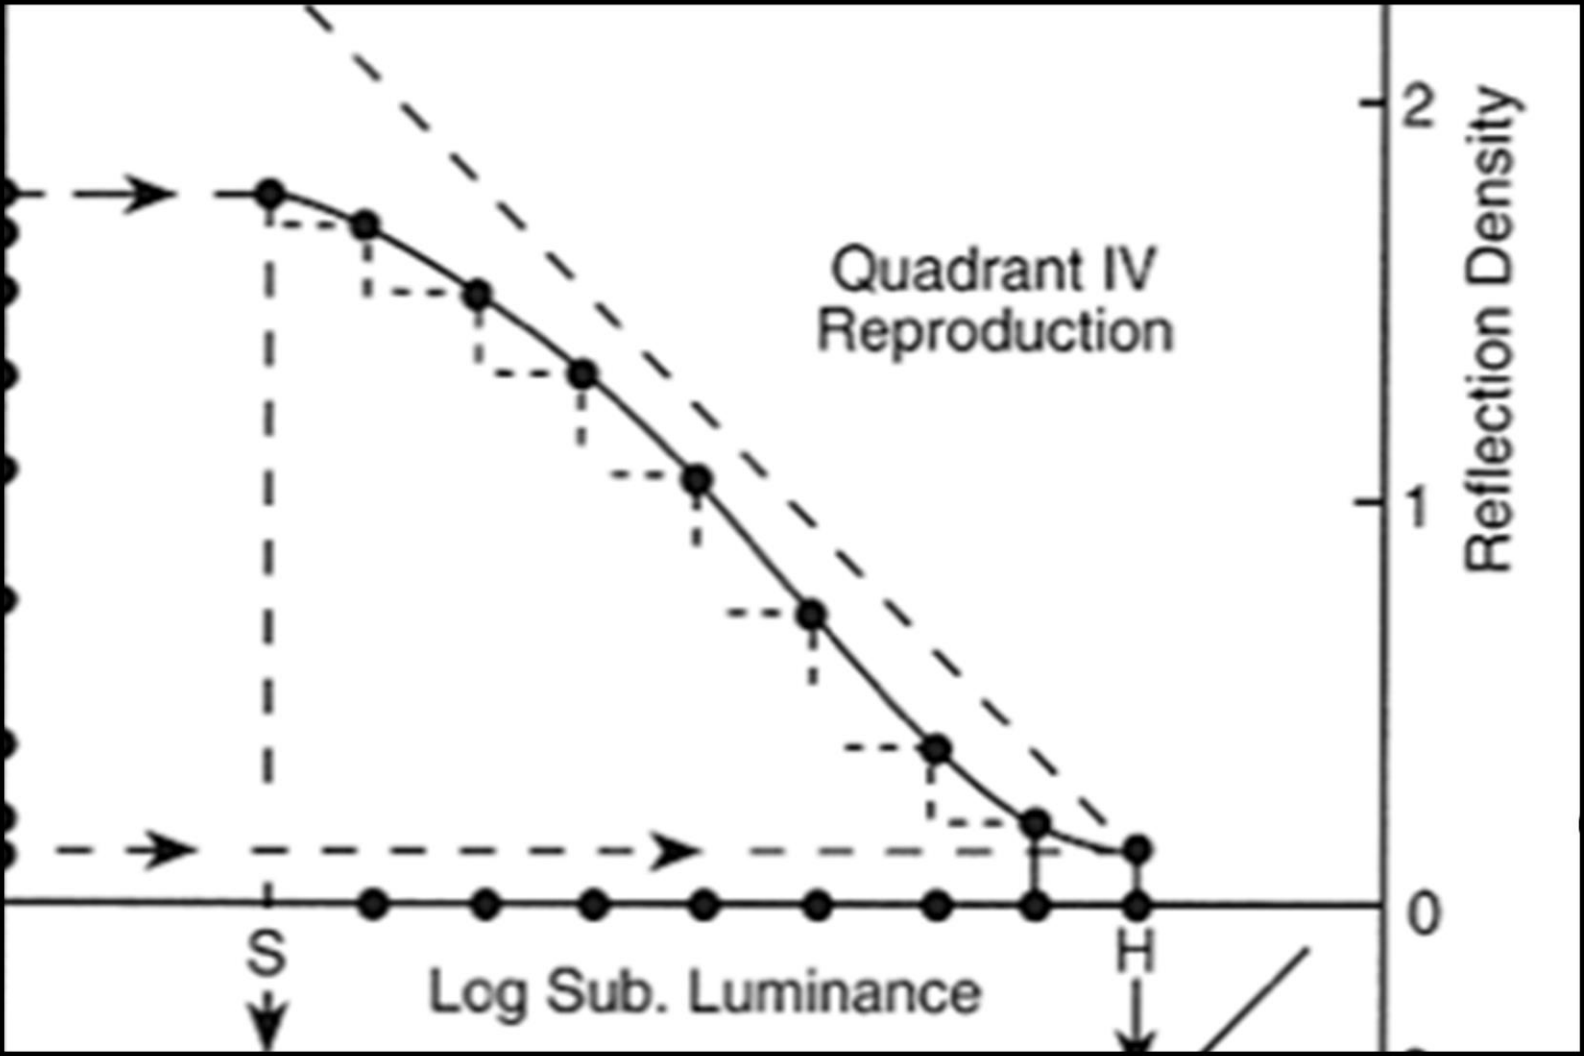

A closer look at the overall tone reproduction curve. We can see it is flat (reducing contrast) at the top and bottom – these regions are called the toe and shoulder respectively - and a pretty straight line in between (we can't tell what the slope is because the units on the axes are different). Although this is for black-and-white film, color film isn't that different – the shape of the curves tends not to vary much across the visible spectrum.

“The Film Curve”

People talk a lot about the “film look” or the “s-shaped film curve”. And it might seem that the shape of this curve is caused by some accident of film photochemisttry, and that a digital product using similar curves is picking a “film look” as a stylistic choice.

That’s not actually an accurate way of looking at it.

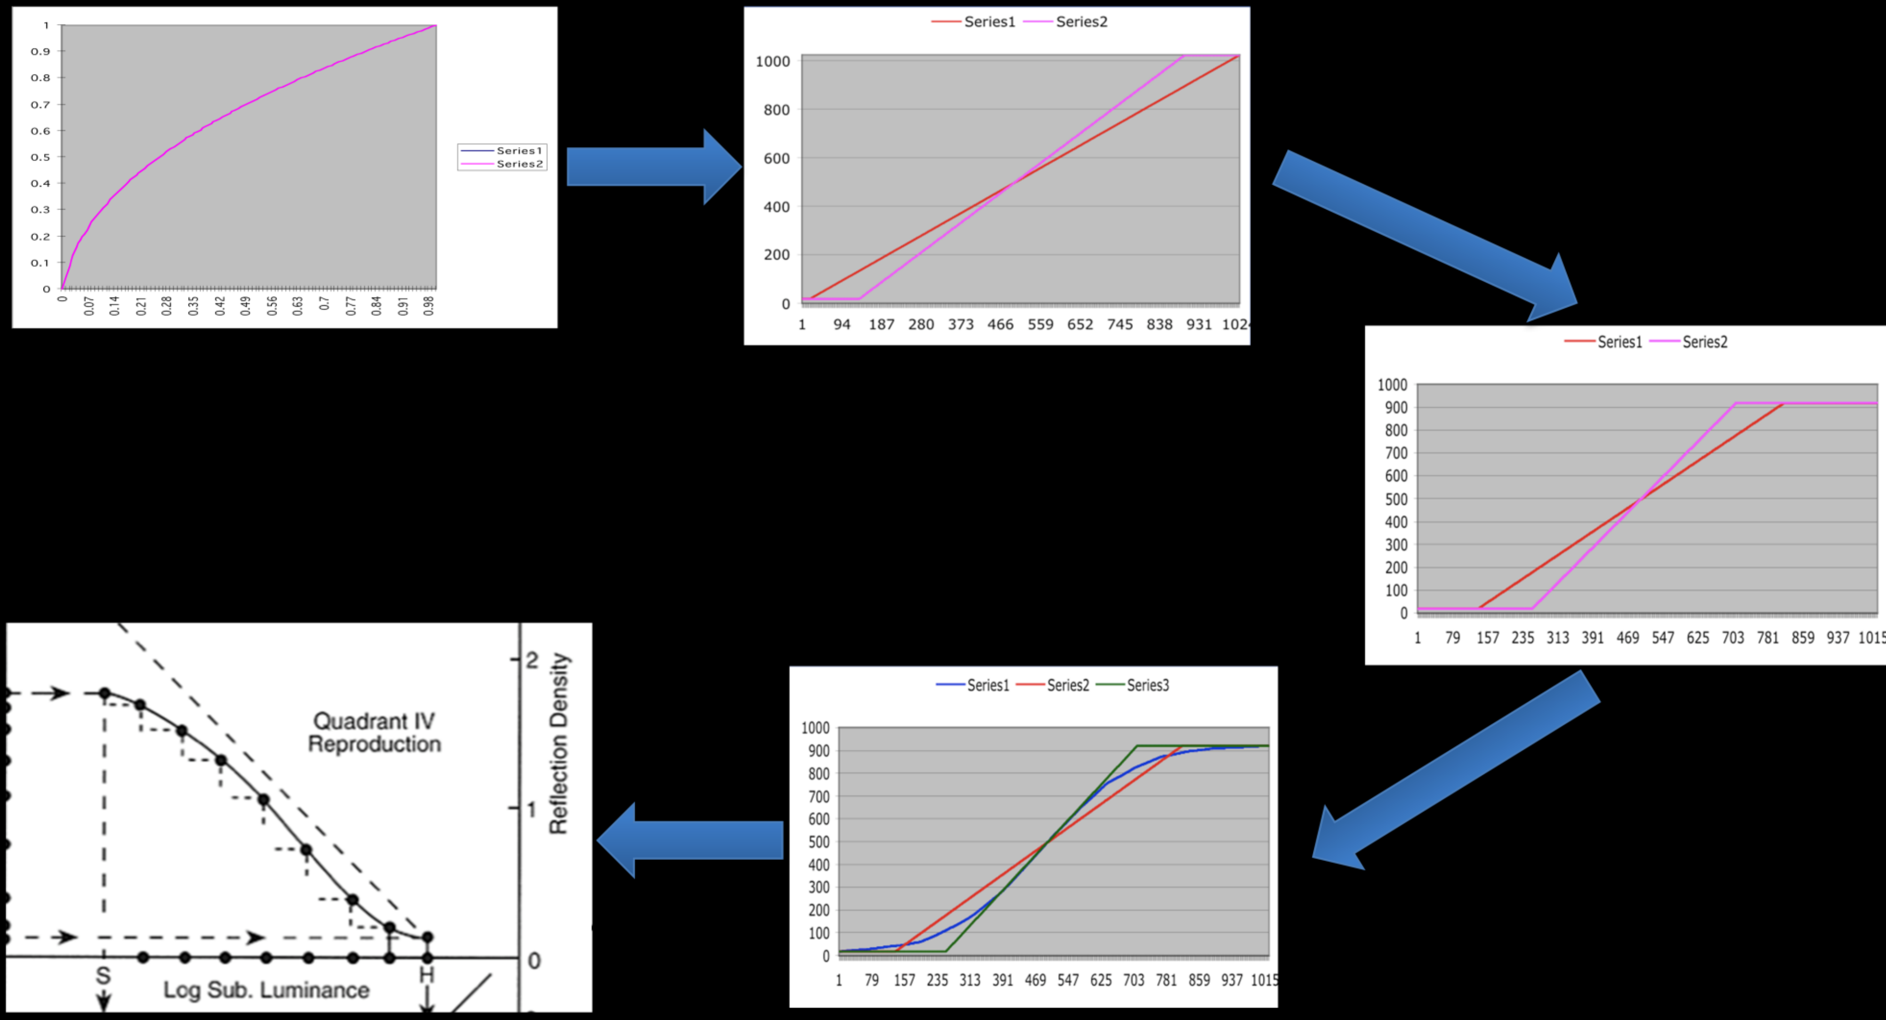

In these slides from his SIGGRAPH 2010 talk – which I recommend you check out – Josh Pines showed that if you build up an image reproduction transform from first principles, you basically end up with the same curve that film has. There is a very real sense in which this curve is optimal. There are variations – not every film stock is exactly the same, after all – but the basic behavior is the correct one for a global "scene to screen" tone transform. Film doesn't happen to do what it does – it was carefully engineered for the best image reproduction characteristics. And this is why digital cameras are also designed to produce similar curves.

(Figures from SIGGRAPH presentation "From Scene to Screen" by Joshua Pines, 2010, used with permission).

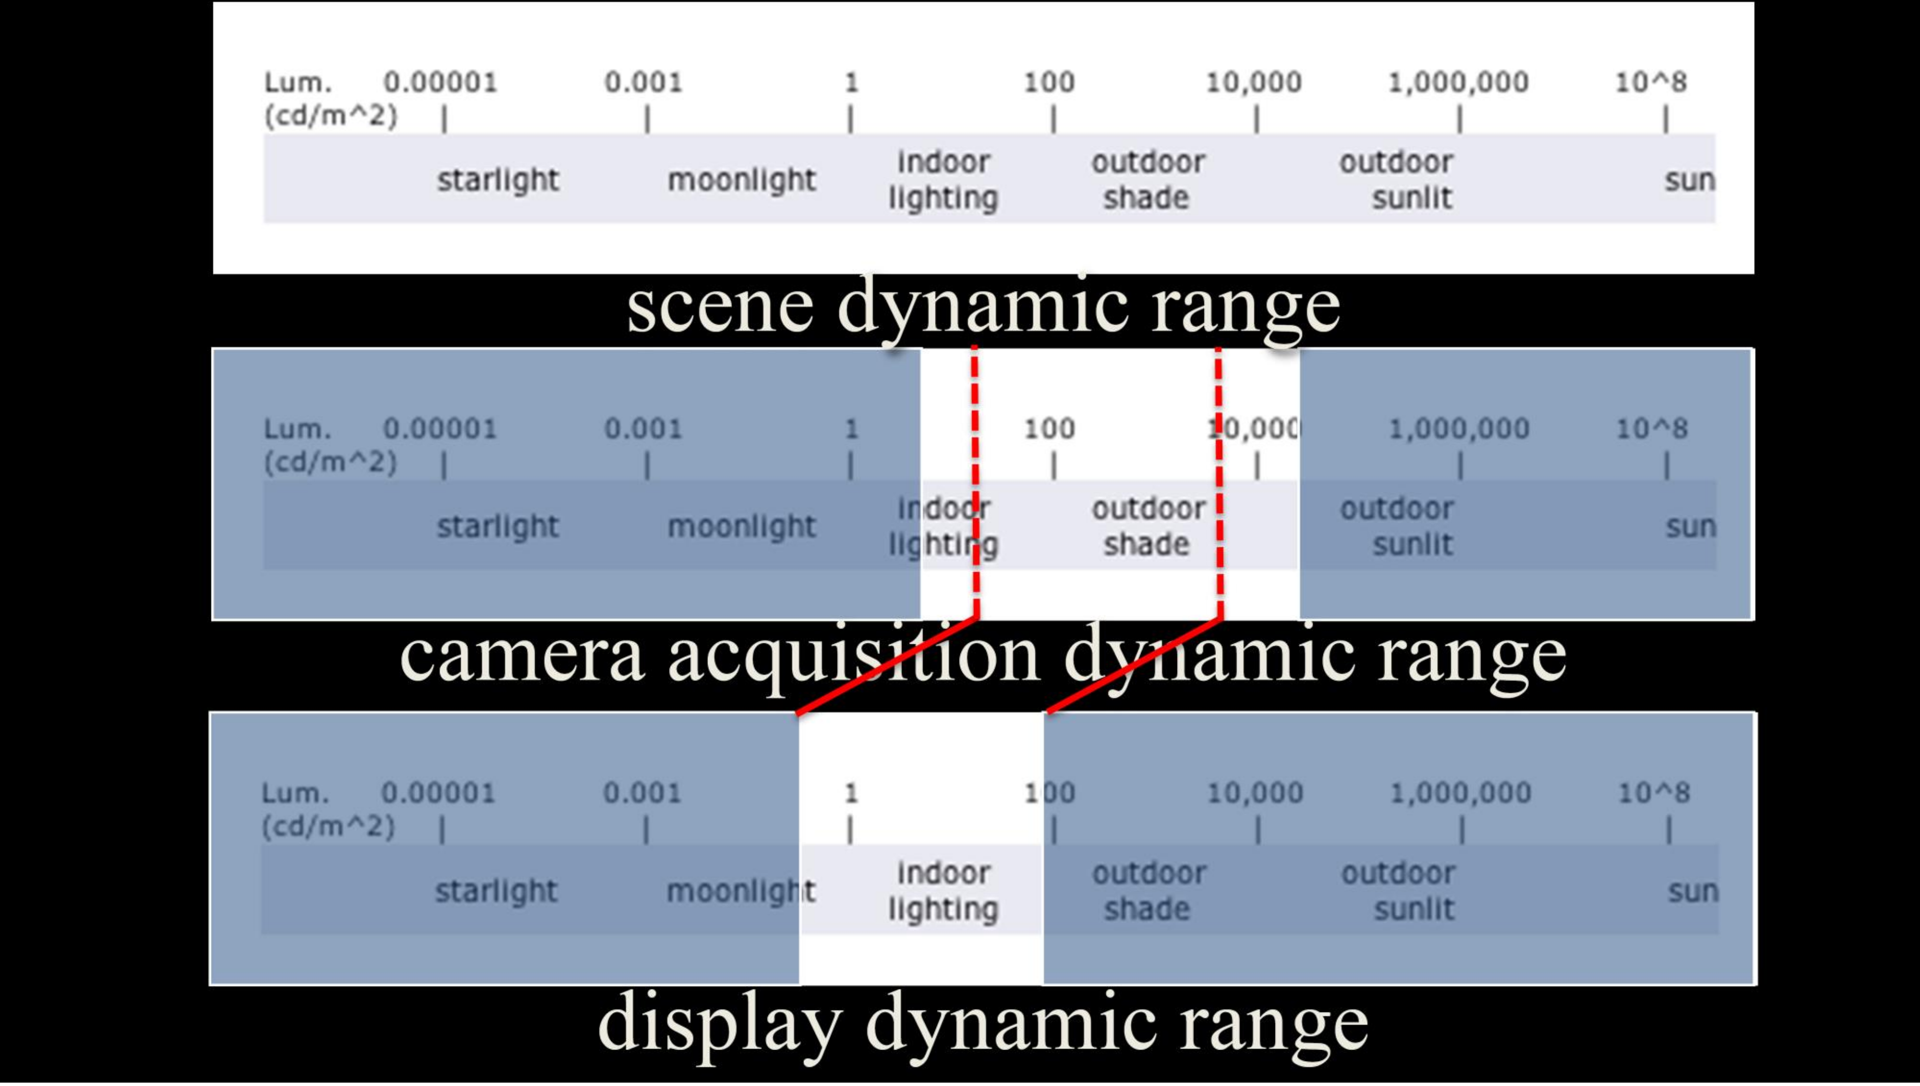

Here's another slide from Josh's presentation, showing an important limitation of this tone transform – as a static, global curve, it has no hope of capturing the entire range of information in the scene. In fact, it can't even capture the entire range that cameras acquire, though the compression in the toe and shoulder region help capture a lot more information than would otherwise be preserved. The reason it works so well in practice is that typically the important information in the scene is in a relatively narrow band – and the sigmoid curve is very good at preserving and reproducing this information in a perceptually accurate way.

(Figure from SIGGRAPH presentation "From Scene to Screen" by Joshua Pines, 2010, used with permission).

Exposure

And this brings us to another learning from photography – exposure. A large part of the art of photography is to use exposure – camera aperture, filters and shutter speed, to place the important information in the scene on the right part of the curve. And it turns out that this "important information" is not purely subjective – there is a method to it.

Zone	Description
0	Pure black
I	Near black, with slight tonality but no texture
II	Textured black; the darkest part of the image in which slight detail is recorded
III	Average dark materials and low values showing adequate texture
IV	Average dark foliage, dark stone, or landscape shadows
V	Middle gray: clear north sky; dark skin, average weathered wood
VI	Average Caucasian skin; light stone; shadows on snow in sunlit landscapes
VII	Very light skin; shadows in snow with acute side lighting
VIII	Lightest tone with texture: textured snow
IX	Slight tone without texture; glaring snow
X	Pure white: light sources and specular reflections

In 1939, Ansel Adams and Fred Archer codified the Zone System. It's a method for setting the exposure such that mapping regions of the image are mapped to print density values based on their albedo – their underlying color independent of lighting. But if we think about it – this basically means that the Zone system uses exposure to divide the lighting out. I will now restate the Ansel Adams Zone system in a somewhat outrageously simplified manner:

(Figure of "Zone System" by Ansel Adams used for scholarly commentary under "fair use").

Zone System (over)simplified

```
exposure = 1 / lighting
```

Although very simplified, there is an important kernel of truth here – and in very simple lighting situations, it is pretty much literally true. The goal of image reproduction – to have the brain's impression of the display be as close as possible to it's impression of the original scene – is reached if objects are mapped to display values according to their material colors, independent of lighting. The brain tends to "discount the illuminant".

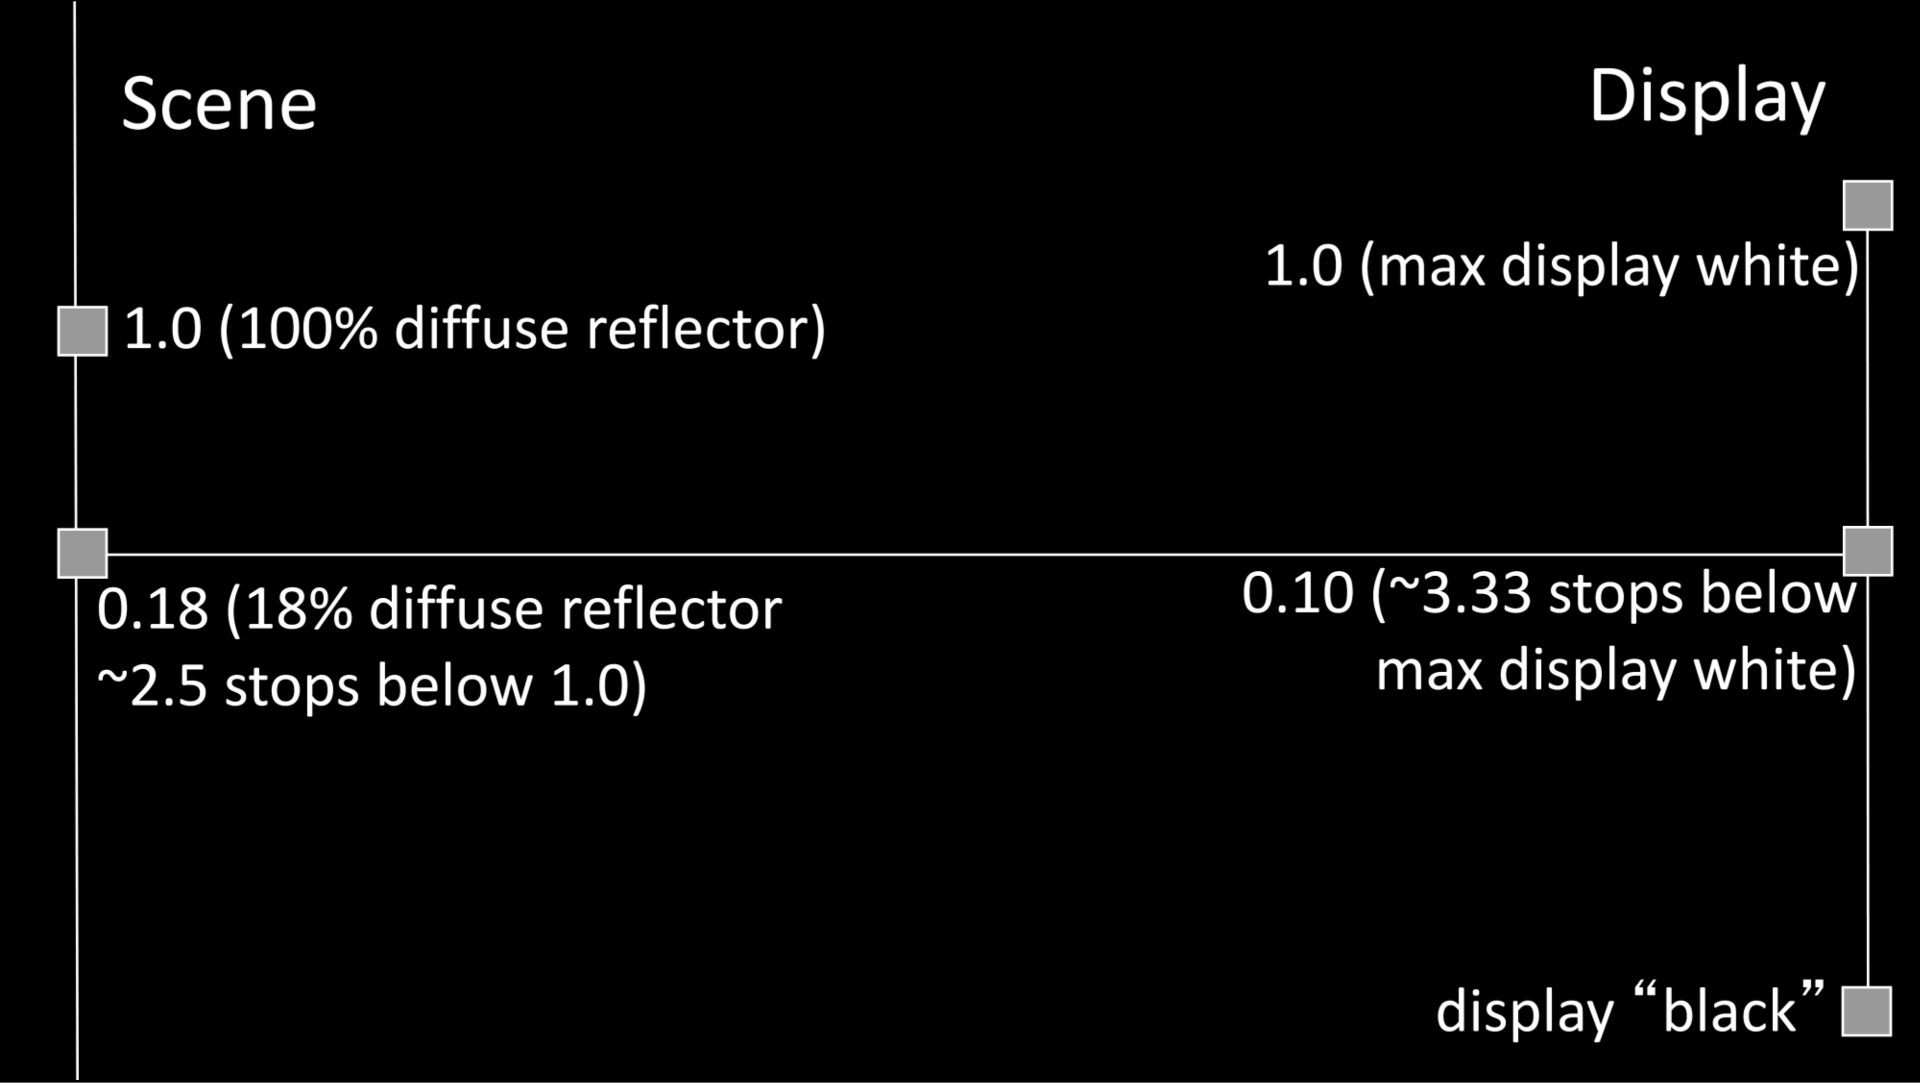

The film tone reproduction curve is designed around this assumption. Here is another one of Josh’s slides, showing properties of good tone transforms. On the left we have scene values, which are unbounded on the top and bottom (this is a log scale, so 0 is at negative infinity). On the right we have (linear, not gamma) display values, bounded on both the top and bottom. Note that here a white Lambertian surface is assumed to have a scene value of 1.0, which is consistent with the “exposure is the reciprocal of lighting” simplification. One rule is that an 18% reflector should map to 10% of maximum display white. Note that the 100% diffuse reflector maps a little lower than display white, to leave some “headroom” for specular highlights and light sources.

(Figure from SIGGRAPH presentation “From Scene to Screen” by Joshua Pines, 2010, used with permission).

$$\bar{L}_w = \frac{1}{N} \exp\left(\sum_{x,y} \log\left(\delta + L_w(x,y)\right) \right)$$

The fact that exposure should counteract the lighting is the problem with things like using scene log average pixel luminance to drive exposure. I'm not saying it's a bad idea in cases where you have limited information, such as auto-exposure systems in cameras – it's likely the best you can do if you don't have information about the scene - if you don't know whether a grey pixel is a dark object brightly lit, a bright object in shadow, or something in the middle.

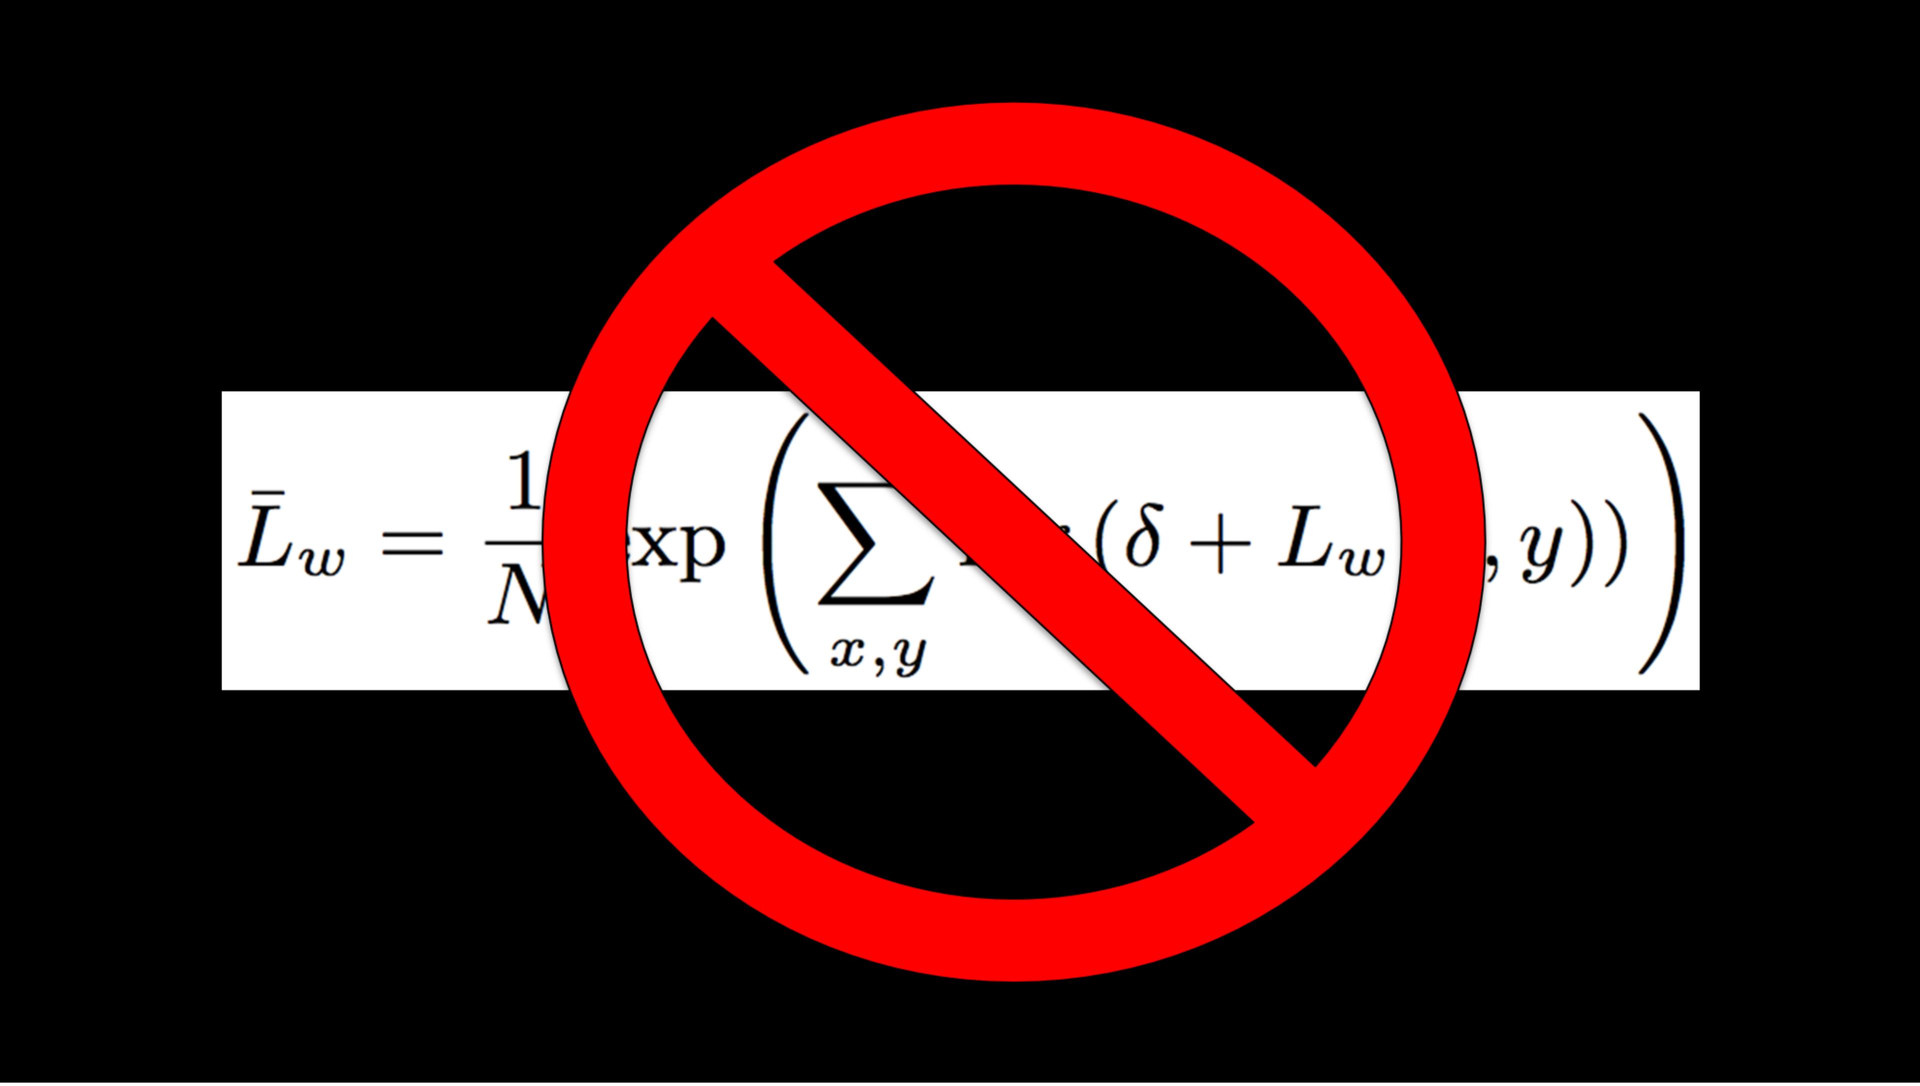

But in CG - and in games in particular - we do have that information, and it’s important to use it. If you have to have an autoexposure system – base it on average lighting luminance, not pixel luminance. But there are good reasons to go with a more manual system if you can – and in many applications, including most games, you can. One of these reasons relates to why my “set exposure to the reciprocal of lighting” simplification of the Zone System was somewhat over-simplified.

```
exposure = 1 / lighting
```

Which lighting?

Which lighting should the exposure be the reciprocal of? Scenes rarely have the same lighting values everywhere, and exposure is a global value over the scene. And of course you don’t want a per-pixel exposure which cancels out the lighting – you would get a very flat-looking scene in that case.

exposure = 1 / lighting

Dominant lighting

You want to use the dominant lighting. This is a bit more of a value judgment, and why I think it's better to have artists set exposure values when possible. If you have a simple scene, like outdoor flat desert during the daytime, then the dominant lighting is pretty clear. An interesting result is that for CG scenes, you can just set both the dominant lighting and exposure multiplier to 1 – which is a nice simple approach I recommend using when possible.

But what about a scene like this, with varying lighting? Here's where the judgment call comes in. If the important content is where the red oval is – let's say it's a movie and a main character is standing there – then you expose assuming that lighting is the dominant lighting and you get an image like this.

("Smoky Tunnel" HDR image from "Mark Fairchild's HDR Photographic Survey", copyright Mark D. Fairchild, available for use for non-commercial presentations / exposed and tone-mapped on this and subsequent slides).

But if the important content is to the left, under the shadows, you need to expose for that.

There is a limit to this approach - if you try to expose for the area under the tunnel, the rest of the image is distractingly overexposed. In situations like this, the ability to control exposure locally (which film photographers used to do by dodging and burning when developing the print) comes in handy. However, such local control is hard to automate – this is doubly true in games, where the images are changing unpredictably. This is where we need to move on to our next source of inspiration.

Cinematography

Movies are closer to games than photography – games have moving images like movies, techniques like dodging and burning aren't applicable (we'll see an interesting exception to this later). Movies have another big difference from photography (this one doesn't apply to games) – it's a lot harder to control exposure since you can't change the shutter speed (you can kinda do that by changing the shutter angle, but that affects motion blur and the overall look a lot and is only used for very particular effects).

Since movies couldn’t control exposure locally and even global control of exposure was somewhat limited, they have relied a lot on lighting control. If this scene was being shot for a movie, they would get around this problem by putting additional hidden lights in the tunnel, bringing it’s brightness closer to the outside scene. It’s time for another gross oversimplification…

Movie lighting = reducing dynamic range

There’s obviously a lot more to it than that, but this is a very common use of movie lighting. People talking in a room with a bright sunlit day outside, and the range is too wide? Put lights in the room or a darkening gel on the windows. This kind of lighting control is very applicable to games, which can control lighting more easily than live-action films (CG animated features enjoy a similar advantage).

Here we have three more slides (last ones!) from Josh Pine's presentation. They show some great examples of cinematography – how the combination of carefully controlled lighting and the sigmoid curve tone reproduction operator combine to create wonderful imagery with a dynamic range that looks a lot higher than the image actually is. There is a huge amount of knowledge in the cinematography community, and most game developers (even professional game lighters) aren't familiar with it. I think there is an opportunity here. This image is from the film "Blade Runner", cinematography by Jordan Cronenweth…

(Image from the film "Blade Runner", 1982, used for scholarly commentary under "fair use").

…"Minority Report", cinematography by Janusz Kamiński…

(Image from the film "Minority Report", 2002, used for scholarly commentary under "fair use").

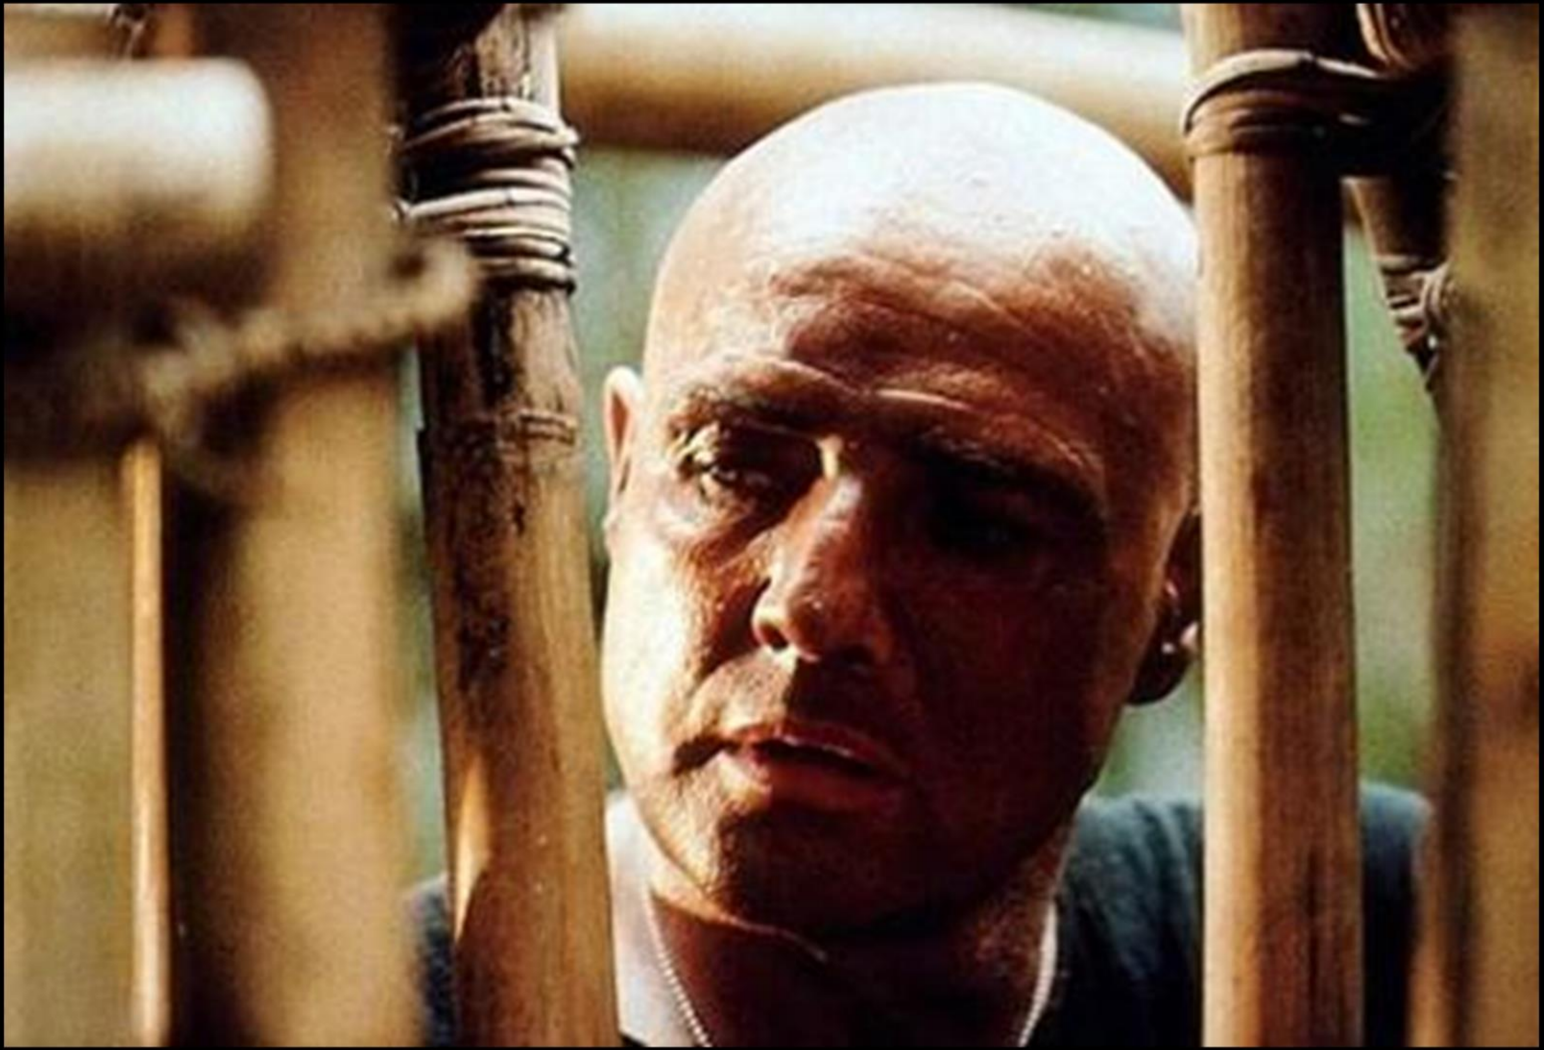

…and “Apocalypse Now”, cinematography by Vittorio Storaro.

(Image from the film “Apocalypse Now”, 1979, used for scholarly commentary under “fair use”).

Color

Another area where I think game development can learn form film is in the use of color. This has both technical and creative aspects.

From the technical aspect, I think game developers need to learn from film production color management practices. Traditionally, in our workflows and practices game developers have gone the “cowboy” route – we’ll tweak it until it looks good. Calibration? We don’ need no stinkin’ calibration!

(“The Herd Quitter” by Charles Marion Russell, 1897, public domain).

I think we need to be a bit more rigorous and scientific about how we process color values.

In film & broadcast, people who are mastering images professionally use professional-grade reference monitors that are carefully calibrated to industry standards. For critical evaluation, they even have rooms with walls painted a specific shade of gray and specific intensities and colors of lighting. This is true even for relatively low-budget productions, such as commercials. In the game industry, this is pretty unusual even in productions with budgets of tens of millions of dollars.

(Images of Cine-tal monitor, Hubble spectrometer and grading room used for scholarly commentary under "fair use").

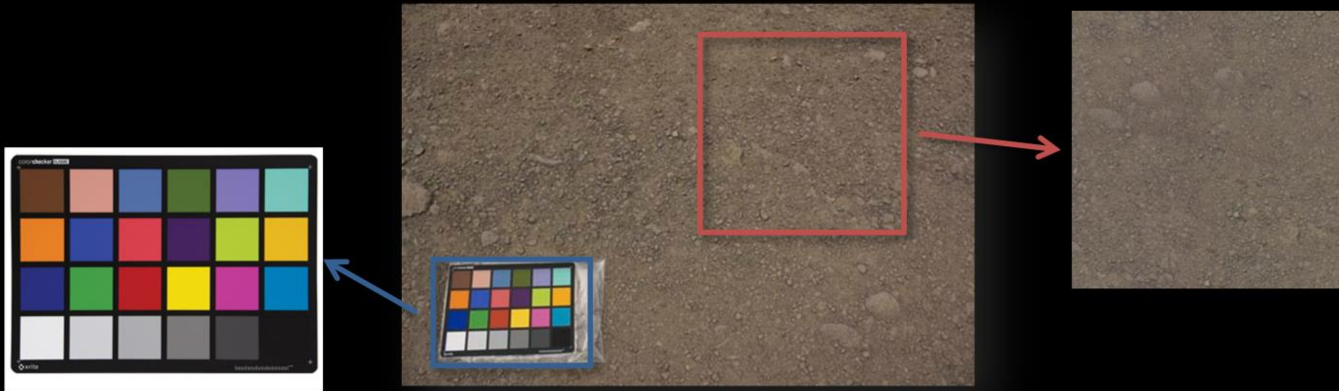

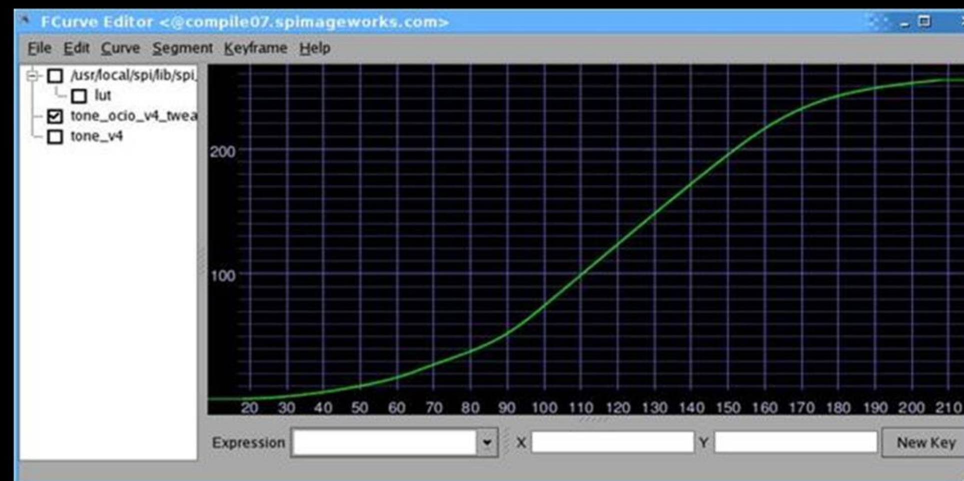

Color management is not just calibration and control of the viewing environments used to author content. It also refers to calibrated capture of source textures (especially important with physically based shaders) and performing operations like diffuse texture painting and skybox painting in the appropriate color spaces.

(Image of color checker on soil from SIGGRAPH presentation "Calibrating Lighting and Materials in Far Cry 3" by Stephen McAuley, 2012, used with permission). (Images of Wacom tablet and OpenColorIO curve used for scholarly commentary under "fair use").

Intentional Use of Color

There are also creative aspects to how color is handled in movies. How the color relates to the plot, character development and mood is carefully thought out, and intentional to a degree rarely seen in games.

A relatively recent example (in film history terms) is the way Pixar uses “color scripts”, like this one for “Up”, to plan out the color palette of the whole movie carefully very early on in production. You can follow the emotional arc of the entire movie from this series of rough color sketches.

(Image of “Up” color script by Pixar used for scholarly commentary under “fair use”).

One of my favorite examples of use of color in movies (and a damn good movie overall – see it if you haven't yet), is the 2001 French film "Amélie". The movies' director, Jean-Pierre Jeunet, was inspired by…

(Image of "Amélie" film poster used for scholarly commentary under "fair use").

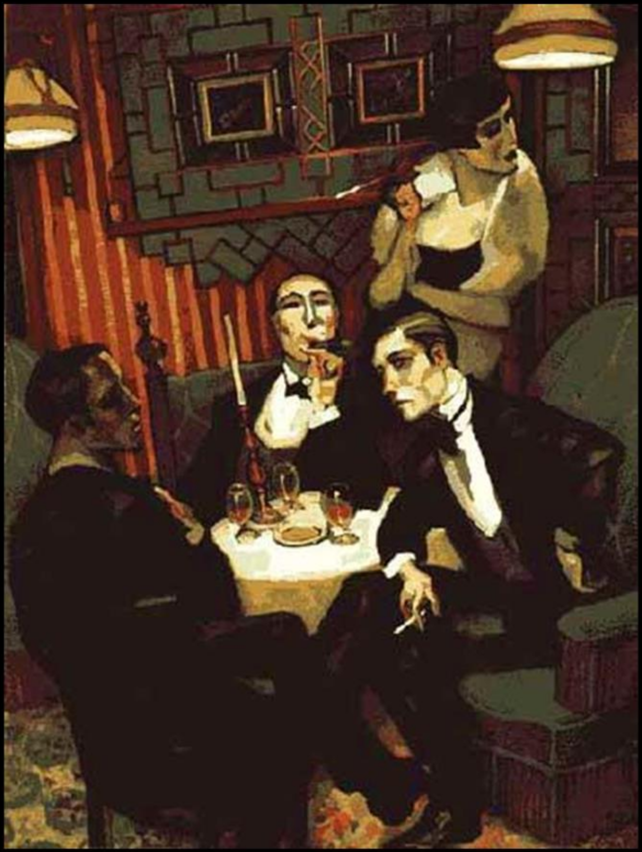

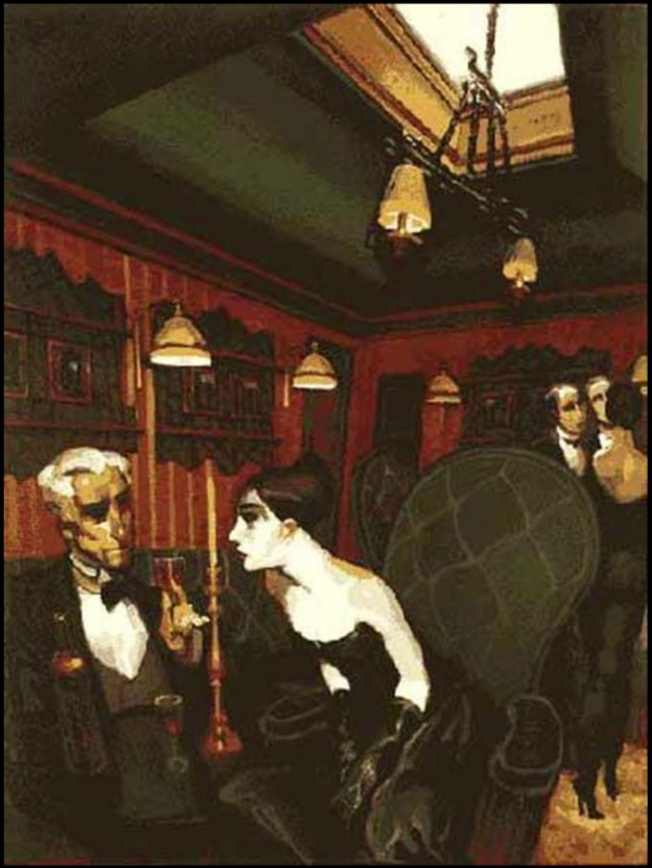

…the work of the Brazilian painter Juarez Machado.

(Images of paintings by Juarez Machado used for scholarly commentary under "fair use").

There is a very consistent color palette running through the movie…

(Images from the film "Amélie" used for scholarly commentary under "fair use" in this and subsequent slides).

…

...

…


des 2 MOULINS
BRASSERIE
BAR BRASSERIE
GRIMBERGEN
BUFFET CHAUD et FROID
RUE LEPIC

…

…, which clearly speaks to the influence of Machado.

In some cases, during post the director and cinematographer realized they had gone too far with the red & green and they needed a small counterbalancing splash of blue…

which they added in post.

Another example.

In this movie, the lighting, sets and costumes all support the palette. This kind of intentional color design running throughout production is one lesson I think games can learn, but there is another lesson here. In Amelie, the finishing touch to the look…

…was the digital color grading done as a post-process.

Here's another example where the sets were designed around this palette…

..and the color grading completed the look.

In this case…

…a window was used to isolate the area that the director wanted to draw the viewer’s eye to…

…and give that area a stronger treatment to make it stand out.

This is an interesting example of local control, like still photography’s dodging and burning, coming in some way to cinema.

Color grading

Color grading in particular is an area where games have been learning from movies, and I think we can learn more. While Amelie was an early case of digital color grading being used in film, it wasn't the first.

Pleasantville came out in 1998, 3 years before Amelie. At that time, color manipulation in movies was typically done photochemically – by varying the exposure time of the different color layers, or via other methods such as bleach bypass. This movie which had color and black-and-white characters interacting couldn't have been done in that way, so all those scenes were processed digitally.

(Image from film "Pleasantville" used for scholarly commentary under "fair use").

The first feature film to be fully processed digitally was O Brother, Where Art Thou? A Coen Brothers film that came out in 2000 –just one year before Amelie. The film was shot in areas of lush greenery, but they needed it to look faded, almost sepia-toned, to evoke the dustbowl era. Here too photochemical processes did not suffice and the entire movie went through a digital color grading process.

(Images from film “O Brother, Where Art Thou?” used for scholarly commentary under “fair use”).

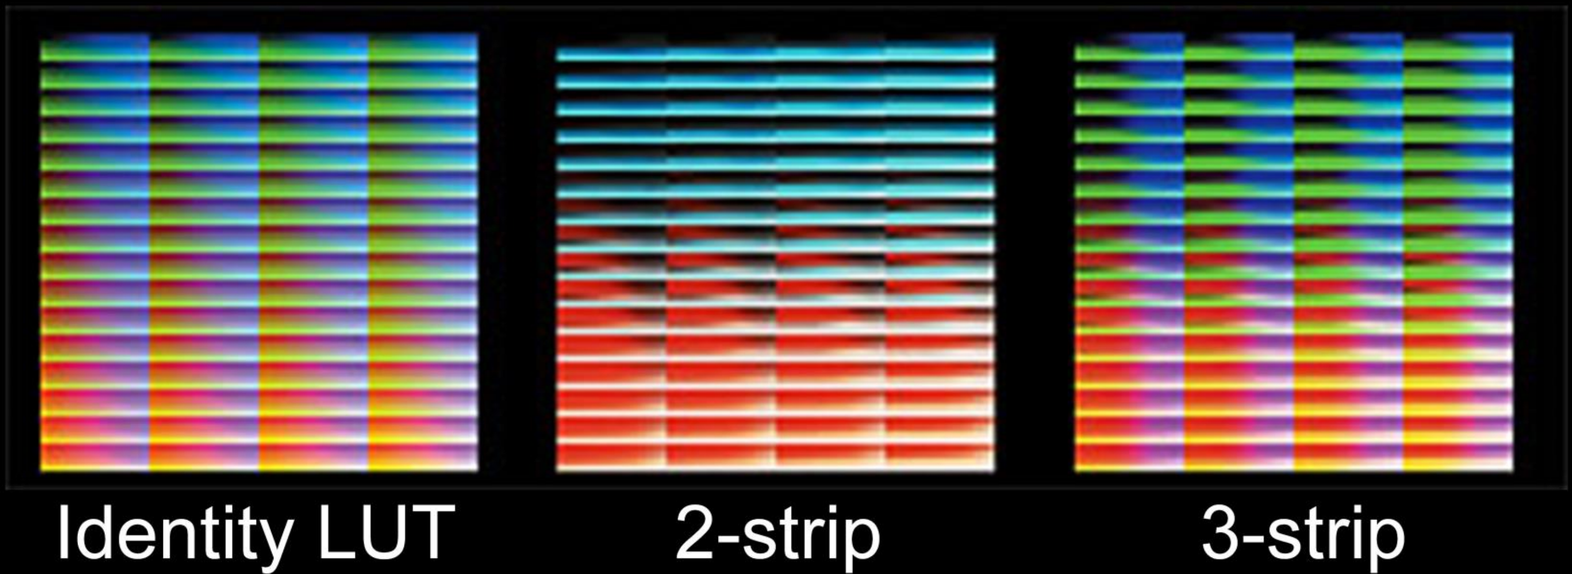

This processing is not always mathematical, often 1D or even 3D lookup tables are used (where red, green and blue were used as the 3 lookup indices, enabling arbitrary color transformations). The LUTs we see here (these images are 3D LUTs sliced into 2D images for viewing) were used in the 2004 film “The Aviator” to mimic old color film processes such as 2-strip and 3-strip Technicolor.

(“The Aviator” LUT images from the “American Cinematographer” 2005 article “The Color Space Conundrum: Part Two” used for scholarly commentary under “fair use”).

Here’s an example of the 2-strip Technicolor simulation – before

(Image from film “The Aviator” used for scholarly commentary under “fair use”).

…and after. Color grading can produce effects that are either subtle, as in Amelie, or profound, as in The Aviator.

(Image from film "The Aviator" used for scholarly commentary under "fair use").

Games have been doing color grading with 3D LUTs authored in Photoshop for a few years now (I helped pioneer this technique when I worked on “God of War 3”, and film color grading was one of my inspirations). However, we can take this inspiration further. Feature film color grading uses tools that are much more powerful and expressive than Photoshop, and the people using them are extremely experienced in doing this particular work. I think we can learn from their methods to do a better job, but via 3D LUTS it is also possible to have feature film colorists grade a game by authoring LUTs for it, using the same tools they used to grade a film. When I worked on “Call of Duty: Black Ops 2” we did exactly that, and it worked quite well.
(Images of color grading control surfaces used for scholarly commentary under “fair use”).

Other Potential Learnings

- Film Production Design (sets, costumes)
- Film Matte Painting (for game skyboxes)
- Architecture (for game spaces – The Witness is already doing this).

We're running out of time, but I'd like to at least mention other areas where I think games can learn from other areas. There are a few more areas where I think we can learn from the film industry's penchant for visual storytelling, by designing spaces and costumes. Film matte painting is another interesting area which can have applications for game skyboxes and other distant scenery. And finally, there is a recent example of a game company hiring architects to design game spaces – this is the kind of thing I think we should be doing more of.

And finally, I'd like to close by mentioning that 2K is looking for graphics programmers in a variety of roles - let me know if you're interested or know of someone who is.