



Driving Toward Reality: Physically Based Tone Mapping and Perceptual Fidelity in Gran Turismo 7

Kenichiro Yasutomi
kenichiro.yasutomi@sony.com

Kentaro Suzuki
kentaro.d.suzuki@sony.com

Hajime Uchimura
hajime.a.uchimura@sony.com



SIGGRAPH 2025
Vancouver+ 10-14 August

Hello, everyone.

Today, we'll be talking about “Driving Toward Reality: Physically Based Tone Mapping and Perceptual Fidelity in Gran Turismo 7.”

Today's Speakers



Kenichiro Yasutomi
Lead Technical Artist



Kentaro Suzuki
Lead Graphics Engineer



Hajime Uchimura
Image Processing Engineer



Here are today's speakers.

From left to right: Kenichiro Yasutomi, Kentaro Suzuki, and Hajime Uchimura.

Why Talk About Tone Mapping?

- Not Just About Shading
 - Tone mapping may seem unrelated to physically based shading at first.



First, let me briefly explain why we talk about tone mapping in this course.

It might seem unusual to talk about tone mapping in a course about physically based shading.

Why Talk About Tone Mapping?

- Not Just About Shading
 - Tone mapping may seem unrelated to physically based shading at first.
- Physically Based Rendering $\neq$ Perceptual Realism
 - The final image is shaped by tone mapping and the display, not by raw rendering.

However, no matter how precise the physically based rendering is, what the user actually sees is the result of tone mapping and the display device.

Why Talk About Tone Mapping?


- Not Just About Shading
 - Tone mapping may seem unrelated to physically based shading at first.
- Physically Based Rendering $\neq$ Perceptual Realism
 - The final image is shaped by tone mapping and the display, not by raw rendering.
- Part of the Physically Based Rendering Pipeline
 - We treat tone mapping as a core part of the physically based rendering process.
 - It is based not only on physics, but also on psychophysics and perceptual psychology.



Therefore, we treat tone mapping as an important part of the physically based rendering pipeline.

This time, we will explain our approach, called physically based tone mapping, which maintains consistency with the physically based rendering pipeline.

Why Talk About Tone Mapping?



POLYPHONY™
DIGITAL

This slide shows the structure of our Physically Based Tone Mapping.

At the center is our main system, called GT7 Tone Mapping, which helps create realistic and consistent images.

On the right are post effects like glare, auto exposure, and auto white balance.


These are closely connected to tone mapping and are not treated as separate steps.

Together, all of these make up what we call our Physically Based Tone Mapping.

The entire system is based on how human vision works — using ideas from physics, psychophysics, and psychology.

In this talk, we'll explain each part in more detail.

Polyphony Digital



- A software development studio fully owned by Sony Interactive Entertainment Inc.
 - Established in 1998
 - Originally the "Gran Turismo" development team within Sony Interactive Entertainment
- Tokyo, Fukuoka, and overseas locations



Let me introduce our studio.

We are a software development studio fully owned by Sony Interactive Entertainment.

Originally, we were the team that created *Gran Turismo* within Sony Interactive Entertainment.

In 1998, we became an independent company, while still remaining within the Sony group.

Today, we have offices in Japan — in Tokyo and Fukuoka — as well as in several other countries.

Gran Turismo Series

- A globally acclaimed driving simulation franchise
- Over 90 million copies sold worldwide (as of Nov. 16, 2022)



POLYPHONY™
DIGITAL

Our main title is Gran Turismo, a globally acclaimed driving simulation franchise.

Gran Turismo 7



- Driving simulator for PlayStation®4/5
 - Advanced car physics
 - More than 500 car models (as of July 14, 2025)
 - Realistic visuals
 - Supports PlayStation® VR2
 - Dynamic time and weather
 - Regular free content updates

POLYPHONY™
DIGITAL

The latest title in the Gran Turismo series is Gran Turismo 7. It's a driving simulator for PlayStation®4 and PlayStation®5.

Note:

It features advanced car physics and more than 500 car models as of July 14, 2025.

The visuals are realistic, and it supports PlayStation® VR2.

It also includes dynamic time and weather, among many other features.

As of 2025, the game continues to receive free content updates every one or two months, with new additions still being made on a regular basis.

Captured in-game



This image was captured in-game. This is a racing scene on a real-world circuit.
You can see the realistic lighting and detailed car models.



Here is an interior view from the seat.
We simulate the dashboard and the cockpit with high accuracy.



This is how the game looks in VR with PlayStation® VR2.
You can turn your head freely and feel as if you are really inside
the car.

Agenda

- Introduction
- Light and Human Perception
- Tone Mapping: Theory and Recent Work
- Tone Mapping: Implementation of "GT7 Tone Mapping"
- Extending Tone Mapping with Perception-Aware Physically Based Post-Processing
- Summary

Now, I'd like to show you today's agenda.

First, we will give a short introduction.

Then, we will talk about how light works — in the real world, on displays, and how we see it as humans.

After that, we will go over the theory and compare some well-known tone mapping methods.

Then, we will talk about how we implemented tone mapping in Gran Turismo 7.

After that, we will talk about extending tone mapping with perception-aware, physically based post-processing.

Finally, we will give a short summary.

Introduction

The Visual Goals of Gran Turismo

- We focus on reproducing the realism perceived by the naked eye.
 - Important for a driving simulator
 - Accurate color reproduction for car manufacturers
 - Visual fidelity for racing tracks
 - Clarity of the real world as a core goal
- Our rendering pipeline is built to support this goal.



Realistic cars



Realistic tracks

POLYPHONY™
DIGITAL

We aim to reproduce the world as we see it with our own eyes. This is very important for a driving simulator like Gran Turismo.

Car manufacturers often ask us to match the real car colors exactly.

They care a lot about how their cars look in the game.

Racing tracks also need to look just like the real ones.


In the end, what we care about most is how close we get to the real world.

To make this happen, we design our rendering pipeline with this goal in mind.

Rendering in Gran Turismo

- The rendering pipeline is based on physical values.
 - Real-world values for lighting and materials
 - No “artist tweaks” just to make it look nice (as much as possible)
 - If the result looks wrong, fix it in the pipeline.
- Some of these ideas exist in physically based rendering, but for us, “real-world reproduction” is the key.

Real-World Reproduction Pipeline



POLYPHONY™
DIGITAL

We use a rendering pipeline that is based on physical values from the real world.

For lighting and materials, we use values taken directly from real-world data.

These values are used either as inputs to the system or as references.

We do not change these values just to make the image look better. If the result looks strange, we fix it within the pipeline instead.

Some of these ideas are also used in general physically based rendering.

But for us, real-world reproduction is especially important.

In our real-world reproduction pipeline, the rendering result goes through tone mapping and the display device.

These steps create a mapping between the physical world and the player's perception.

So, what the player sees is not the physical value itself, but the mapped result.

That is why consistent tone mapping is important for real-world reproduction.

Lighting in Gran Turismo 7

- Natural light
 - LUT-based sky dome lighting using physically based atmospheric simulation [Suzuki and Yasutomi 23]
 - In Gran Turismo 7, time and weather can change in real time and randomly. Both affect the lighting in the game.
- Artificial light
 - RGB values are calculated from actual spectral measurement data.
 - For car headlights and similar lights, we use HDR photographic data.



Let's talk about lighting in Gran Turismo 7.

First, natural light.

We use a sky dome system based on LUTs and physically based atmospheric simulation.

This was shown in detail in our GDC 2023 talk.

In Gran Turismo 7, both time and weather can change during gameplay. It can go from day to night, or from clear skies to rain skies — both in real time and randomly.

These changes all affect how the lighting looks in the game.

Next, for artificial light, we calculate RGB values from real-world spectral measurement data.


And for lights like car headlights, we use data measured directly using HDR photography.

Note:

Our atmospheric simulation:

- The light source is based on the ASTM G-173 reference solar spectrum.
- Atmospheric scattering is simulated using real observation-based models.

Example of Lighting | A Deep Blue Sky



This is an example of how our sky dome system reproduces a deep blue sky at sunset.

The atmospheric simulation creates a smooth color gradient based on physical models.

Example of Lighting | Sharp Reflections from the Sun



Here you can see strong reflections on the car surface. These highlights come directly from the simulated sun and depend on the viewing angle and material.

Example of Lighting | Headlights and Taillights



This image shows headlights that were rendered using real measured light data.
We used HDR photography to capture the correct brightness and color.

Materials in Gran Turismo 7

- Our BRDF model is based on the Disney Principled BRDF [Burley 12].
- For car materials, we measure real automotive paint samples and fit the results to BRDF parameters.
 - BRDF fitting using a differentiable renderer [Uchimura 23].
 - Car manufacturers care about whiteness, blackness, and saturation.

Measurement
Equipment



Mini-Diff V2 [Synopsys 25]



X-Rite MA-T12 [X-Rite 25]

POLYPHONY™
DIGITAL

Now, let's look at materials.

Our BRDF model is based on the Disney Principled BRDF.

For car materials, we measure real car paint samples and fit the data to BRDF parameters used directly in the game.

This was also presented at CEDEC 2023.

Car paints often have much more vivid and extreme colors than natural objects.

Manufacturers care deeply about how white, black, or saturated the paint looks — resulting in a very wide color gamut.

Note:

Our material models are a slight extension of the Disney Principled BRDF, as described in detail by [Hirai and Takano 22].


To obtain accurate results, we use two different types of measurement equipment. Each device has different characteristics.

Mini-Diff V2: Capable of dense and accurate BRDF measurements but captures only RGB (tri-stimulus) data, limiting spectral accuracy.


X-Rite MA-T12: Measures color spectrally with high color accuracy, but BRDF sampling is sparse and less precise.

Material Fitting


- We measure actual paint samples.
- Shader parameters are numerically optimized using Mitsuba 3 [Jakob et al. 22] to reproduce the BRDF.
 - Reconstructed in Rec. 2020



Rendered result using
fitted BRDF



Measured real paint samples



Rendered result using
fitted BRDF



Here is a more concrete example of our material fitting process.

We begin by measuring actual paint samples.

The red plate shown in the center is one of those real samples.

Then, we use Mitsuba 3 to numerically optimize the shader parameters.

This process helps us reproduce the BRDF as accurately as possible.


The red spheres at the bottom left and right are both rendered results.

They use the BRDF that was fitted from the measured data.

As you can see, they look quite similar to the real sample.

The resulting colors are spread across a very wide color gamut.

Car Paint Colors Go Beyond the Standard Gamut



- Over 5,000 real car paint colors
- This chart shows the distribution of base paint colors.
- Many colors are clearly outside the Rec. 709 gamut.

Now, let's look at the color distribution of real car paints.

We have measured more than five thousand real automotive paint colors.

This chart shows how these base colors are distributed.

The triangle with a white dotted line shows the Rec. 709 gamut. Many of these colors are far outside it because car makers often choose bold, saturated colors.

This makes wide color gamut rendering essential.

Example of Materials | Vivid, Pure Red



This is an example of a vivid red material.
The color is highly saturated and is clearly outside the Rec. 709 gamut.

Example of Materials | Shiny Black



Here is a shiny black material.

Black paint is especially sensitive to reflections and surface smoothness.

Tone Mapping in Gran Turismo


An important bridge between the physical world of rendering and the user's perceptual world.



Now, let us talk about tone mapping in Gran Turismo.
Tone mapping is a crucial bridge between the physical world of rendering and the user's perceptual world, making it essential for our goal of realistic reproduction.

Tone Mapping in Gran Turismo

An important bridge between the physical world of rendering and the user's perceptual world.



- Requirements
 - Physically based world with a wide dynamic range (in both brightness and color)
 - Need to produce optimal output for a wide variety of displays (both SDR and HDR)
 - Reproduce realistic appearances consistently in both formats
 - HDR as a natural extension of SDR
- Gran Turismo Sport (2017)
 - GT Tone Mapping [Uchimura and Suzuki 18]

POLYPHONY™
DIGITAL

Here are the key requirements we must meet.

The world we render is physically based and has a very wide dynamic range.

This applies to both brightness and color.

We also support both SDR and HDR. Our goal is to faithfully reproduce how things look to the human eye in both formats.

To meet these requirements, in Gran Turismo Sport, we introduced our own tone mapping system, which we called GT Tone Mapping.

Tone Mapping in Gran Turismo

- New additional requirements
 - Real-time time and weather changes were introduced in Gran Turismo 7.
 - The HDR display environment evolved over 5 years.
 - 2017: HDR TVs were still in the early stages of development.
 - 2022: LCD, OLED, and Mini LED technologies advanced.



From Gran Turismo Sport to 7, several important changes made tone mapping even more challenging.

First, real-time changes in time and weather were introduced. Second, the HDR display environment evolved a lot.

Note:

In Gran Turismo Sport, players chose the time and weather before the race, and they stayed fixed.

But in Gran Turismo 7, these conditions can change during the race, in real time.

This adds a major layer of complexity for lighting control.

- About displays

Back in 2017, HDR TVs were still new. But by 2022, we saw rapid progress in display technologies like LCD, OLED, and Mini LED.

More manufacturers entered the market, and display characteristics became much more diverse, especially for gaming monitors.

Tone Mapping in Gran Turismo

- New additional requirements
 - Real-time time and weather changes were introduced in Gran Turismo 7.
 - The HDR display environment evolved over 5 years.
 - 2017: HDR TVs were still in the early stages of development.
 - 2022: LCD, OLED, and Mini LED technologies advanced.
- Gran Turismo 7 (2022)
 - GT7 Tone Mapping
 - Perception-aware physically based post-processing

Physically based tone mapping approach



Over these five years, our own understanding of tone mapping has also improved.

In Gran Turismo 7, we introduced a physically based tone mapping system that unifies tone mapping and display adaptation to address previous challenges.

Light and Human Perception

Tone mapping translates physical quantities of light into perceptual ones.

So, let's start with the relationship between them, and discuss how we perceive physical luminance.

Luminance in the Real World




Jasper National Park, Canada

 POLYPHONY™
DIGITAL

Let's look at luminance in the real world for a moment.

Luminance in the Real World



This is Jasper National Park, located about 600km northeast of Vancouver.

Clouds and snow under direct sunlight are among the few natural scenes can exceed 10,000 nits, excluding the sun itself.

Note: The luminance values you see here were measured on-site from “Scapes” in Gran Turismo.

Luminance in the Real World




Historic Villages of Shirakawa-go, Japan

POLYPHONY™
DIGITAL

Here you see Shirakawa-go, a UNESCO heritage village in Japan. The cloudy sky diffuses the luminance and makes the scene dimmer.

Luminance in the Real World



Historic Villages of Shirakawa-go, Japan

POLYPHONY™
DIGITAL

Here you see Shirakawa-go, a UNESCO heritage village in Japan. The cloudy sky diffuses the luminance and makes the scene dimmer.

Luminance in the Real World



Tokyo Gaikan Expressway, Japan

POLYPHONY™
DIGITAL

Next is a cool tunnel on a Japanese highway, it is even darker; the luminance of the lighting really stands out.

Luminance in the Real World



Tokyo Gaikan Expressway, Japan

POLYPHONY™
DIGITAL

Next is a cool tunnel on a Japanese highway, it is even darker; the luminance of the lighting really stands out.

Luminance in the Real World



Streets of Vancouver at night, Canada

 POLYPHONY™
DIGITAL

The last is a night scene.

Luminance in the Real World




Streets of Vancouver at night, Canada

The lighting is exceptionally bright, while everything else is extremely dark.

Dynamic Range

Physical Luminance of the Real World




POLYPHONY™
DIGITAL

As you have seen, the physical luminance in the real world ranges from extremely high to extremely low.

Dynamic Range

Physical Luminance of the Real World



Human Vision System




POLYPHONY™
DIGITAL

We perceive this wide range through two types of photoreceptors — rods and cones.

Combined, they cover a substantial dynamic range.

Dynamic Range

Physical Luminance of the Real World



Human Vision System



POLYPHONY™
DIGITAL


However, the range is the result of physiological adaptation.

According to Kunkel and Reinhard*, the simultaneous dynamic range is estimated at 3.7 log units — roughly a factor of five thousand.

*[Kunkel and Reinhard 2010]

Dynamic Range | Media

Physical Luminance of the Real World



sRGB Display

Cinema Screen

Printed Image

*Depends on ambient lighting

Human Vision System




POLYPHONY™
DIGITAL

Now, these are the approximated ranges of typical media luminance. They are considerably narrow.

Dynamic Range | Media

Physical Luminance of the Real World



sRGB Display

Cinema Screen

Printed Image

*Depends on ambient lighting

Human Vision System




POLYPHONY™
DIGITAL

Naturally, your eyes adapt to the luminance of each medium.

Dynamic Range | Media

Physical Luminance of the Real World



HDR Display (Ideal)

sRGB Display

Cinema Screen

Printed Image

*Depends on ambient lighting

Human Vision System




POLYPHONY™
DIGITAL

HDR displays are supposed to have this large dynamic range,
but ...

Dynamic Range | Media

Physical Luminance of the Real World




This is the reality.

OLEDs have excellent dark tones because they can turn pixels off completely, but their maximum brightness is not so high. The opposite is true for LCDs.

Dynamic Range | Media

Physical Luminance of the Real World



Yet, true black is still elusive because of physical factors — such as reflections from the room on the screen.

On the chart, OLED shows an impressive dark dynamic range. However, in practice, effective contrast depends on the environment, and it cannot be guaranteed.

Adaptation | Perception is Relative

- We perceive relative luminance, not absolute values.
- Maintains **contrast constancy**



POLYPHONY™
DIGITAL

So, how does adaptation work?

The key perceptual outcome is contrast constancy: we perceive relative luminance as constant, not its absolute value.

That's why a photographic image can still appear realistic: even with much lower absolute luminance, it feels like the same scene if relative contrast is preserved.



This image may have looked too dark on the previous slide. Yet, when seen alone, it soon begins to feel natural.

That's contrast constancy at work.

Adaptation | Exposure $\neq$ Adaptation

Exposure Adjustment

- **Quantity-based**
 - Linear and uniform
 - Physically fixed response
 - Contrast changes as a result

Perceptual Adaptation

- **Contrast-based**
 - Nonlinear and nonuniform
 - Dynamic adaptation to contrast **delta**
 - **Reconstructs contrast afterward** in neural circuits




However, adaptation is not the same as adjusting exposure on a camera.

Exposure adjustment scales absolute luminance, and contrast changes only as a result.

Adaptation instead works to preserve contrast itself.

Adaptation | Exposure $\neq$ Adaptation




POLYPHONY™
DIGITAL

Adaptation is nonlinear and nonuniform because retinal receptors and neural circuits work in tandem.


Signals from receptors are always uneven — they adjust sensitivity locally to avoid saturation and preserve contrast. The coherency of our view is the result of neural circuits somehow reconstructing it through a highly nonlinear process.

Adaptation | Exposure $\neq$ Adaptation



Think of it as a lossy codec — and its recovery is never perfect.
Next, we'll see where the limits become evident.

Adaptation | Limits of Contrast Constancy




Here you see the human eye's contrast sensitivity as a function of adaptation luminance — a dimension of our vision as fundamental as the well-known spatial frequency.

Note for experts: This is essentially the inverse of the TVI curve, plotted as sensitivity. It can be regarded as a luminance-dependent cross-section of the full, multi-dimensional contrast sensitivity function in reality.


The curve is based on Blackwell's classical measurements [Blackwell 46], extended to lower luminance with the de Vries–Rose law, and to higher luminance with bleaching models proposed by [Burkhardt 94] and [Stockman et al. 06].

Adaptation | Limits of Contrast Constancy




Contrast constancy holds only in the narrow range where contrast sensitivity stays roughly constant — about 2 to 200 nits.

Adaptation | Limits of Contrast Constancy



It is not unusual for contrast constancy to break down locally, even under natural daylight conditions.


Adaptation | Limits of Contrast Constancy



The decrease in sensitivity is largely due to the signal-to-noise ratio.


Not only at low luminance where the input itself is faint, but also at high luminance, the signal can be buried under internal noise, as receptors lower their absolute sensitivity to avoid saturation — and, importantly, this occurs locally rather than uniformly.

Adaptation | Limits of Contrast Constancy



The dynamic range of sRGB displays lies exactly within the constant range; it is a truly reasonable design.

Adaptation | Limits of Contrast Constancy




By integrating this contrast sensitivity curve, we obtain a contrast response curve that closely resembles the PQ curve. It shows whether the response is linear or not.




Now, let's revisit Jasper National Park.

Adaptation | Input Defines Response



It is an exceptionally high luminance scene, and this bar represents the approximate dynamic range of the scene.


Adaptation | Input Defines Response



When we view a real scene, our visual system operates in this region of the contrast response curve.

On the other hand, when viewing the same scene scaled to an sRGB display, the response shifts downward.

Adaptation | Input Defines Response




When we view a real scene, our visual system operates in this region of the contrast response curve.

On the other hand, when viewing the same scene scaled to an sRGB display, the response shifts downward.

Adaptation | Input Defines Response

Outside contrast constancy — **exposure $\neq$ adaptation**




POLYPHONY™
DIGITAL

Beyond the range where contrast constancy holds, our perception is no longer linear, so exposure cannot approximate adaptation — thus discarding the nonlinear perceptual response of real-world scenes.

Perceptual Contrast Mapping

Compressing loses luminance
(Shoulder: highlight roll-off)




POLYPHONY™
DIGITAL

The “shoulder” in tone mapping can be described as embedding the perceptual response at high luminance.

Perceptual Contrast Mapping

Normalizing contrast includes internal noise
(Toe)



POLYPHONY™
DIGITAL

And the “Toe” represents that dark areas are difficult to perceive due to internal noise in perception.



An image in which all contrasts except the sun are within the dynamic range

This image preserves almost all contrasts except for the sun within its dynamic range.

But it is not immersive — because contrasts that would have been compressed beyond recovery in reality are now visible.

Human perception doesn't see high luminance in this way.



Immersive exposure, without perception-based tone mapping

So, let's raise the exposure — pushing those contrasts beyond perception.

But now highlights blow out, colors distort; instead, clipping is now visible in this range.


Again, human perception doesn't see high luminance in this way — exposure alone reaches its limit here.



Here, the missing nonlinear contrast response has been embedded back into the image.
We can restore a sense of perceptual reality that existed on site — and this is tone mapping.

Note:

When viewing the scene directly, the receptors become less sensitive, so contrast is compressed at high luminance and can be buried in internal noise at low luminance.




Tone mappers often have an S-shaped curve.

First from film, it also arises from matching contrast sensitivity to scene luminance — a perceptual necessity.

Tone Mapping: History and Related Technologies

Painting | Perception to Perception



- **Perception-based tone mapping**
 - Centuries before the term existed
- **Subjective and conventional**
 - Physical luminance unknown

 POLYPHONY™
DIGITAL

Before cameras and before film, for centuries, painting was the only way for humanity to capture, express, and share their visions.

“Valeur” in classical painting was about expressing how things should feel in the light.

It was essentially tone mapping, tuned not for linear fidelity, but for human perception.

Painting | Perception to Perception



Raphael, The School of Athens (1509-1510)



Rembrandt, The Night Watch (1642)



Watteau, The Embarkation for Cythera (1717)




Renoir, Luncheon of the Boating Party (1881)

In art history, a change in painting style could be described as a change in the manner of tone mapping.

Each style had its own answer to the same question:
“How should this scene feel to the human eye?”

Silver-Halide Film | Measurable Vision



- **The scene met a measurable response**
 - Measured, modeled, and engineered
- **The origin of “Photorealistic”**
 - Accumulated empiricism and trust

POLYPHONY™
DIGITAL

Next is silver halide film — this is where things changed.

The image wasn't just crafted — it was measured, modeled, and engineered.

And something emerged — what we now call “photorealistic.”
It came from trusted empiricism, not from first principles.

Silver-Halide Film | Measurable Vision

- **Nonlinear** response, yet **consistent and repeatable**.
 - The tone curve of silver-halide film closely resembles human perception.
- **The D-log E curve** made the photographic response **quantifiable and predictable**. [Hurter and Driffield 1890]
- Essentially **empirical**, yet the benchmark of realism



Silver-halide film had a nonlinear but consistent response. Its tone curve resembled human perception well enough to feel realistic.

In 1890, Hurter and Driffield introduced the D-log E curve, which made this response measurable for the first time. It took decades for it to become a standard industrial practice, but from that point on, tone mapping could be systematically engineered.

It was trusted for a long time not because it was physically accurate — but because it felt right, again and again. Even today, our sense of “realistic” carries this legacy.

CRT (Television) | Display-Driven Encoding



- **The CRT's gamma property looked unnatural**
 - In contrast to silver halide film, which looked natural
- **Starting point for the OETF**
 - Signal encoding to compensate for the display gamma (i.e., the EOTF)

CRTs used in television also had a nonlinear response, known as gamma. However, unlike silver-halide film, this physical characteristic did not resemble human perception at all.

To compensate, engineers applied an inverse gamma function to the signal. This pair of curves — now known as the OETF and the EOTF — mathematically cancels each other out and restores the original signal.

Although it began as a practical compromise to make televisions more affordable, this very idea would later become the foundation of tone mapping.

* The image is credited to "© Raimond Spekking / CC BY-SA 4.0 (via Wikimedia Commons)"

CRT (Television) | Display-Driven Encoding



- **The CRT's gamma property looked unnatural**
 - In contrast to silver halide film, which looked natural
- **Starting point for the OETF**
 - Signal encoding to compensate for the display gamma (i.e., the EOTF)
- **No human perceptual tone mapping**
 - No more than **distortion correction**

Just to be clear — the design purpose of the OETF was distortion correction, not human perceptual tone mapping in the end-to-end image pipeline.


* The image is credited to "© Raimond Spekking / CC BY-SA 4.0 (via Wikimedia Commons)"

CRT (Television) | OETF and EOTF

Let's explain how the OETF and EOTF work on CRTs. This is not the main point, but it cannot be ignored.

CRT (Television) | OETF and EOTF

- **EOTF (Electro-Optical Transfer Function, or Display Gamma)**
 - Due to the CRT's nonlinear physical response




CRT's EOTF

CRT displays had a nonlinear light output — roughly a power-law response — with respect to the input voltage, which is called gamma.

This was not designed for human perception; it was simply a physical property of the CRT.

CRT (Television) | OETF and EOTF

- **EOTF (Electro-Optical Transfer Function, or Display Gamma)**
 - Due to the CRT's nonlinear physical response
- **OETF (Opto-Electronic Transfer Function, or Gamma Correction)**
 - Compensation for the CRT's EOTF to restore a linear response




POLYPHONY™
DIGITAL

Rather than making CRTs linear, engineers applied the OETF to the signal before transmission

CRT (Television) | OETF and EOTF

- **EOTF (Electro-Optical Transfer Function, or Display Gamma)**
 - Due to the CRT's nonlinear physical response
- **OETF (Opto-Electronic Transfer Function, or Gamma Correction)**
 - Compensation for the CRT's EOTF to restore a linear response




to linearize the entire system at a lower cost.

CRT (Television) | OETF and EOTF

As a serendipitous result, the combination of the EOTF and the OETF optimized the **signal distribution** for human perception.

This optimization also works with **quantized digital signals**.



POLYPHONY™
DIGITAL

Interestingly, this signal nonlinearity accidentally aligned well with human perceptual sensitivity — denser codes for darker tones, sparser codes for brighter ones.

Note that this only optimized signal distribution — the physical luminance output was the linear and, unlike film, no perceptual tone curve was embedded.

CRT (Television) | Display-Driven Encoding

- **CRT's gamma response shaped the entire signal path**
 - Once again, in contrast to film, which was input-driven
- **OETF + EOTF = approximate linearity at the end**
 - Signal distribution was accidentally optimized for perception (**Weber-Fechner law**)
 - Linear-to-linear images with early video cameras **didn't look good**
- **Implicit tone mapping was embedded in the content**
 - Little awareness of tone mapping until the early 2000s



In the CRT era, the display's gamma dictated the entire image pipeline. From camera to broadcast, everything bent around that one nonlinear output.

OETF plus EOTF restored an approximately linear signal — and, in addition, accidentally optimized the signal distribution for human perception.

Since early video cameras had an approximately linear sensor response, combining them with CRTs yielded a nearly linear-to-linear pipeline. In theory, this was a technical ideal — clean, accurate, and efficient.

However, in practice, they were often criticized as dull or even cheap-looking.


That mismatch revealed the absence of perceptual tone mapping in television — and perhaps the fact that it was needed.

In fact, tone mapping had already been happening virtually — TV production focused on technical tuning, while film production aimed to emulate the film look — but it was on-site craftsmanship, not an explicit pipeline.


These facts seem to suggest a later tone mapping pipeline.

CRT (Television) | Display-Driven Encoding

- **sRGB inherits the CRT gamma and color space**
 - sRGB is based on **Rec. 709**, which was originally defined for HDTV, and is nearly identical.



OETF (gamma correction)



sRGB color space

POLYPHONY™
DIGITAL

There's one more thing we shouldn't overlook.

sRGB — still the standard today — was designed for CRT displays in the 1990s.

Its primaries match CRT phosphors. Its gamma mimics CRT's voltage curve.

Even 30 years later, modern LCDs and OLEDs often mimic CRTs just to stay compatible.

What started as a physical quirk became a digital convention. And that convention still defines how we see correct color.

Digital Imaging | Unleashing Linearity, Losing Norms



- **Linear input via a digital camera or CGI**
 - Free from embedded tone curves of media or devices
 - Enables a fully linear workflow
- **Explicit tone mapping in the pipeline**
 - Became part of downstream content design
 - Some tone mappers are formally defined (e.g., Reinhard, ACES)

POLYPHONY™
DIGITAL

Digital workflows began with post-production — but eventually, cameras went digital too.

That gave us linear, scene-referred data: clean, reversible, and free from embedded tone curves.

Tone mapping was no longer implicit — it became an explicit step that had to be designed into the pipeline, with responsibility.

That's why tone mappers like Reinhard and ACES didn't just emerge as creative tools — they were responses to a structural need.

Digital Imaging | Unleashing Linearity, Losing Norms

- **Scene-referred linear workflows**
 - Driven by VFX demands in the film industry
 - Tone mappers and other components remained film-centric by default
- **Free from film, image tone shifted to cameras and LUTs**
 - No longer a film retrospective, but norms were lost
 - OCIO and ACES as frameworks for unifying fragmented image formats
- **PBR entered under the same boundary conditions**
 - Tone mapping is part of the content
 - What will we choose?



The filmic look remained a standard in tone mapping — partly because film had earned deep trust as a photorealistic reference, and partly because linear workflows were largely driven by VFX in the film industry.

Visual effects needed to match back to film, so the pipeline inherited its tone.

But from those filmic origins, we moved toward perceptual freedom. Linear workflows made tone mapping a matter of choice: freedom from embedded curves and freedom to design tone in post.

Physically based rendering (PBR) operates under exactly the same assumptions.

And as we already knew from television, linear isn't enough.

Filmic tone curves still make sense. But now we're free to build something else, too.

HDR Imaging | Tone Mapping Meets Physics



- **Goodbye CRT, hello JND**
 - New EOTF: **Perceptual Quantizer (SMPTE ST 2084)**
 - Maps absolute luminance to human perception based on JND (Just Noticeable Difference)
- **Radiant and dark — display evolution**
 - Extends both **peak luminance** and **dynamic range**
 - Wide-gamut **Rec.2020** instead of **Rec.709**
- **Fragmented tone mapping**
 - Multiple implementations are required for consistent output
 - EOTF divergence breaks compatibility

POLYPHONY™
DIGITAL

CRTs are gone — HDR brought a brand-new EOTF: the Perceptual Quantizer.

It maps absolute luminance to just-noticeable differences in human perception — at least in theory, it aligns physical output with how we actually see.

Displays became brighter, darker, more contrasty — and with Rec.2020, more colorful, too.

However, tone mapping faced a serious problem.

PQ and sRGB are not compatible — the same implementation won't produce the same result.

Worse, HDR displays don't follow a single spec — or even two, and the intent behind tone mapping began to fragment.

HDR Imaging | Tone Mapping Meets Physics

- **Display-referred, not signal-referred**
 - PQ defines output luminance **directly** — with no paired OETF
 - Breaks the OETF-EOTF **symmetry** of legacy CRT workflows
- **One EOTF, many displays — by design**
 - Although PQ defines a fixed EOTF, the actual range is **display-dependent**
 - Each display has **its own capability and tone mapping**
- **Adaptive tone mapping: a necessary compromise**
 - Tone mapping adapts to the output format and metadata to maintain compatibility
 - In HDR, creators must take responsibility for display diversity — not just for the image itself.



PQ is fundamentally different from CRT-based systems — it defines absolute output luminance, not signal encoding. There is no paired OETF — the EOTF stands alone, mapped directly to human perception via JND steps.

That breaks the symmetry we had before — the logic of signal + display forming an end-to-end linear system no longer applies.

Each display interprets the PQ signal within its own physical limitations.

While the signal is display-referred in principle, the actual tone mapping is still display-specific.

In the HDR pipeline, tone mapping is no longer one-size-fits-all. Adaptive implementations became necessary.

Media-Driven or Content-Driven?

Media	Tone Mapping Comes from	
Paintings	Perceptual feedback	↑ Tone Mapping Embedded in the Media (Inevitable limitation)
Films	Physical property	
CRT	Just signal efficiency	↓ Tone Mapping Applied to the Content (Fully optional)
Digital	Content-dependent	
HDR		

 POLYPHONY™
DIGITAL

In paintings, tone is naturally adapted through perceptual feedback.

In films, tone curves emerged from material physics aligned with human visual perception.

Conversely, on displays, tone mapping is no longer inherent in the media — it must be implemented as part of the content, not as a limitation.



**Tone mapping is not a cosmetic finish —
Perceptual bridge between radiance and vision.**

Most researchers in physically based rendering might treat the rendered scene as the final product.

However, the final product is not the scene — it's what the viewer sees on a display.

In the digital age, tone mapping has become content-dependent, and to some, it may seem arbitrary.

But perceptually and physically grounded paths do exist.

Tone mapping is the bridge between physical radiance and perceptual experience — not a cosmetic finish, but a perceptual necessity.

Tone Mapping: Theory and Recent Work

Tone Mapping is...

Mapping between different color volumes in general

- “Color Volume Mapping” or “Color Volume Transform”
- “Tone mapping” is the term commonly used in the game industry.

Compression process:

- Luminance compression
 - E.g., linear rendering space => limited luminance TVs
- Chromatic compression
 - E.g., Rec. 2020 => Rec. 709



Tone mapping can be generalized as a mapping conversion between different color volumes.


In the game industry, “tone mapping” is the term commonly used.

Tone mapping includes two compression processes: luminance compression and chromatic compression.

Color Volumes

The **color volume** refers to the three-dimensional range of colors that a medium can represent.

- Printed media
 - Paper and ink with varying color purity
- Display devices
 - SDR TV
 - HDR TV
 - Cinema projector + screen




Color volume is a three-dimensional space that represents the full range of colors a medium can reproduce.

Paper, ink, and display devices each have their own unique color volumes.

HDR-capable TVs offer far larger color volumes than conventional SDR displays.

Beginning of Tone Mapping in Games

- The history of tone mapping in games began with “HDR rendering,” which produces images with a much wider dynamic range than standard displays can show.
- This is why tone mapping is needed.



POLYPHONY™
DIGITAL

The history of tone mapping in games began with HDR rendering.

HDR rendering enabled games to produce images with a much wider dynamic range than before.

The diagram shows how pixel colors can lie outside the standard sRGB color volume of displays.

Tone Mapping in Games (excerpt)

Year	Title	Author
2005 GDC	High Dynamic Range Rendering in Valve's Source Engine	Gary McTaggart (Valve)
2006 GDC	HDR Meets Black & White 2: A Case Study	Francesco Carucci (Lionhead Studios)
2008 GDC	Lighting and Material of Halo 3 (HDR, but w/o tone mapping information)	Hao Chen (Bungie)
2009 GDC	The Rendering Technology of KILLZONE 2 (HDR, but w/o tone mapping information)	Michal Valient (Guerrilla Games)
2010 GDC	Uncharted 2: HDR Lighting	John Hable (Naughty Dog)
2011 GDC	Lighting You Up in Battlefield 3	Kenny Magnusson (DICE)
2012 SIGGRAPH	Calibrating Lighting and Materials in Far Cry 3 (HDR, but w/o tone mapping information)	Stephen McAuley (Ubisoft Montreal)
2013 SIGGRAPH	Graphics Gems from CryENGINE 3	Tiago Sousa, Nicolas Schulz, Nickolay Kasyan (Crytek)
2014 GDC	Next Generation Post Processing in Call of Duty: Advanced Warfare	Jorge Jimenez (Activision)
2016 GDC	Advanced Techniques and Optimization of HDR & VDR Color Pipelines	Timothy Lottes (AMD)
2016 CEDEC	HDR Output, theory and practice by silicon studio	Masaki Kawase (silicon studio)
2017 GDC	High Dynamic Range Color Grading and Display in Frostbite	Alex Fry (Electronic Arts)
2017 CEDEC	Don't let me tell you that SDR is more attractive! ~HDR support for the latest titles~ [in Japanese]	Oushiro Tanaka (CAPCOM)
2017 CEDEC+KYUSHU	Practical HDR Output Support ~Building a Rendering Pipeline~ [in Japanese]	Masaki Kawase (Silicon Studio)
2018 SIGGRAPH ASIA	Practical HDR and Wide Color Techniques in Gran Turismo SPORT	Hajime Uchimura, Kentaro Suzuki (Polyphony Digital)
2018 GDC	HDR Ecosystem for Games	Evan Hart (NVIDIA)
2018 CEDEC	Physically based HDR-rendered fighting games on mobile! ~Android/iOS "TEKKEN" Graphics Example~ [in Japanese]	Shigeki Tomizawa (Bandai Namco Studio)
2018 GDC	"Not-So-Little Light": Bringing Destiny 2 to HDR Displays	Kevin Todisco (Vicarious Visions)
2019 GDC	A Blend of GCN Optimisation and Colour Processing	Timothy Lottes, Jordan Logan (AMD)
2021 SIGGRAPH	Real-Time Samurai Cinema: Lighting, Atmosphere, and Tonemapping in Ghost of Tsushima	Jasmin Patry (Sucker Punch Productions)
2025 Develop:Brighton	Understanding and Modifying the Colour Pipeline in Unreal	Ali Cormack (Studio Gobo)

Here is a brief list of presentations related to HDR and tone mapping in games.

Evolution of Game Graphics

Advanced rendering techniques have become practical:

- HDR rendering
- Physically based rendering
- Support for HDR displays



Games and real-time rendering have always been constrained by the limits of computer processing power.

As computational capabilities have improved, advanced rendering techniques such as HDR rendering and physically based rendering have become practical.

Current Demands for Tone Mapping in Games

Current demands:

- Rendering time must be strictly limited.
- Tone mapping must adapt to the performance characteristics of the display device.

Tone mapping methods in games must be:

- Simple in calculation (able to run in real time).
- Dynamically adjustable.



Game graphics change constantly and must meet strict real-time rendering budgets.


Modern game consoles can probe a display's HDR performance, yet the readings are often inconsistent.

As a result, tone mapping operators now need to be lightweight, real-time capable, and easily adjustable via parameters.

Measuring Display Device Performance

Gaming consoles have an HDR calibration system.

- Max Full-Frame Luminance (MaxFFML)
 - Peak brightness across the entire screen
- Max Window Luminance (MaxTML)
 - Peak brightness in a smaller windowed area
- Minimum Luminance (MinTML)
 - The darkest visible luminance level



[HGiG 19][Ignite 25]

POLYPHONY™
DIGITAL

These days, display devices have also evolved. HDR displays are becoming common.

In gaming consoles, we have an HDR calibration screen.

On PlayStation®, for example, the display's luminance capacity is adjusted by the user, and the performance indices are sent to the game.

Games can adjust their output image using these values.

However, these performance indices are unpredictable.

Some displays have a peak luminance of 357 nits, while others have 5,436 nits.

This is why the tone mapping function in games needs to be flexible.

Global vs. Localized Tone Mapping

Global Tone Mapping

- Apply tone mapping pixel by pixel
- The same process is applied to every pixel on the screen.
- Easy to implement
- LUTs (Look-Up Tables) work very well for this approach.

Localized Tone Mapping

- Apply tone mapping using localized image context
- There are correlations between neighboring pixels.
- Requires more computation



Before comparing tone-mapping methods, we'll first define global vs. local operators.

A global operator applies one curve to every pixel.

A localized operator adapts the curve using nearby pixels.

That extra context improves detail but costs more GPU time.

Per-channel Mapping vs. Color Volume Mapping

Per-channel Mapping

- Apply the tone mapping function to each channel separately:

$$R = f(R)$$

$$G = f(G)$$

$$B = f(B)$$

- Easy to implement
- Can cause hue twist
- No gamut mapping

Color Volume Mapping

- Apply a single function to the combined (R, G, B) values:

$$\begin{pmatrix} R' \\ G' \\ B' \end{pmatrix} = f \begin{pmatrix} R \\ G \\ B \end{pmatrix}$$

- More complicated
- Allows more flexible hue control
- Gamut Mapping included




Games use two tone-mapping styles: per-channel and color-volume.

The per-channel style maps R, G, and B separately. This is simple but prone to hue twists.

The color-volume style maps the combined RGB value. It offers finer hue control.

Color Clipping Artifacts



With HDR wide-gamut rendering, color clipping artifacts can sometimes appear.

In this night scene, the red and blue lights are artificial light sources and cause severe color clipping.

Some tone mapping methods include a gamut mapping effect to avoid these artifacts.

Gamut Mapping

- Color coordinate ranges are limited by the medium.
 - “Gamut”
- Gamut mapping is required when converting between different color volumes.
- If no gamut mapping is applied, color clipping artifacts may occur.

All color volumes have their own color coordinate ranges, which are called the gamut.

Gamut mapping maps colors into that range while preserving hue and appearance.

Selecting Priorities in Gamut Mapping

- When mapping to a smaller color volume, you must always make trade-offs among saturation, brightness, and hue [ICC 22].
- ICC color profiles have three rendering intents:
 - Absolute Colorimetric Intent: Preserves the exact color values, including the white point of the source.
 - Relative Colorimetric Intent: Preserves the color values relative to the destination white point.
 - Perceptual Intent: Preserves the overall visual appearance and relationships between colors, even if the exact values change.
- The choice between these methods ultimately depends on the goal; there is no silver bullet.



When mapping a color to a smaller color volume, you must always make trade-offs among saturation, brightness, and hue. The characteristics of gamut mapping depend on the artistic decision or artistic goal.

For example, ICC color profiles have these three rendering intents:

Absolute Colorimetric Intent, which preserves the exact color values.

Relative Colorimetric Intent, which preserves the color value relationship to the white point.

Perceptual Intent, which preserves the overall visual appearance and relationships, but the exact values may change.


Gamut Mapping by Desaturation

- Simple gamut mapping
- Just apply “desaturation” matrix

$$\begin{pmatrix} R \\ G \\ B \end{pmatrix} = \begin{pmatrix} 1 - 2x & x & x \\ x & 1 - 2x & x \\ x & x & 1 - 2x \end{pmatrix} \begin{pmatrix} R \\ G \\ B \end{pmatrix}$$

The simplest way to perform gamut mapping is to apply desaturation.


Gamut Mapping by Desaturation (AgX)



For example, AgX tone mapping utilizes a custom inset color space to achieve a desaturation effect.

Gamut Mapping with an Intermediate Color Space

- Some gamut mapping methods use an intermediate color space.
 - XYZ, LAB, ICtCp, Jzazbz, etc.
- Color-preserving operations are easier to perform in these spaces.




Another gamut mapping method uses an intermediate color space, such as CIE XYZ or LAB, etc. In such a color space, color-preserving operations are much easier than in pure RGB space.

Gamut Mapping with an Intermediate Color Space (ACES 2.0)



For example, ACES 2.0 tone mapping utilizes the JMh color space to perform color calculations.




Here we compare the gamut mapping results produced by two tone mapping methods and a no gamut mapping picture. Each method reflects unique design choices and exhibits distinct characteristics.

Commonly Used Tone Mappings in Games

- Scale and Clamp
- Reinhard [Reinhard et al. 02]
- Filmic Curve (Hable's [Hable 10] and AMD VDR [Lottes 16])
- ACES 1.3 [Academy 21] & 2.0 [Academy 25]
- AgX [Sobotka 22]
- Khronos PBR Neutral Tone Mapper [Lalish 24]

According to articles and presentations, the following tone mapping methods are commonly used in games nowadays.

Scale and Clamp



- Commonly used when "HDR rendering" first appeared in games
 - A primitive form of tone mapping [McTaggart et al. 06] [Carucci 06]
- Often used in conjunction with auto exposure

$$y = \text{clamp}(x * \text{Exposure}, 0, 1)$$

The first tone mapping method is the simple scale and clamp. It was commonly used in the early days of HDR rendering in games.


The diagram illustrates how values change with varying exposure levels.

Note:

Various techniques were developed as part of exposure adjustment.

In recent years, exposure adjustment has been separated from tone mapping.

Reinhard



- Per-channel
- First introduced in 2002
- Inspired by photographic film techniques
- Designed to balance high-luminance roll-off and the white point (w)

$$y = \frac{x * \left(1 + \frac{x}{w^2}\right)}{1 + x}$$


Reinhard is a popular tone mapping method.

This method is inspired by photographic film techniques and is designed to balance high-luminance roll-off and the white point.

Note:

$w = 4$

Filmic Curve (Hable's)



- Per-channel
- Popular in early 2010 games
- Originally designed for “Uncharted 2”
- Smooth curve with polynomial formulation
- Hard to control
- SDR output only

$$y = \frac{x * (A * x + C * B) + D * E}{x * (A * x + B) + D * F} - \frac{E}{F}$$

POLYPHONY™
DIGITAL

Filmic curve introduced by John Hable is a popular tone mapping function used in early 2010s games.

This curve uses a polynomial formulation to produce a smooth S-shaped function.

As a result, it was hard to control the shape of the curve by adjusting the parameters.

Note:

A = 0.15

B = 0.50

C = 0.10


D = 0.20

E = 0.02

F = 0.30

W = 11.2

Filmic Curve (AMD's VDR)



- Per-channel
- Designed to have variable peak luminance
- Supports VDR (Variable Dynamic Range) output
- Hard to control

$$y = \frac{x^a}{b * x^{ad} + c}$$

POLYPHONY™
DIGITAL

There is another filmic curve, introduced by Timothy Lottes at GDC 2016.

This curve is designed to provide controllability of the peak luminance, which corresponds to the name VDR: variable dynamic range.

Note:

MidIn = 0.18

MidOut = 0.18

HdrMax = 64.0

Contrast = 1.3

Shoulder = 0.995

Parameters a, b, c, and d are derived from the values above.

AgX (1/2)



- Color volume mapping
- A mix of three techniques:
 - "Rotated" primary color space conversion to reduce hue twist
 - Normalized log2 space conversion for dynamic range
 - Flexible S-shaped function with a linear mid-tone

 POLYPHONY™
DIGITAL

AgX is a tone mapper that has recently become popular.

AgX includes three techniques:

- A rotated primary color space conversion for some gamut compression.
- Normalized log2 conversion for dynamic range compression.
- An S-shaped function for the final touch.

AgX (2/2)

- Used in Blender, Godot Engine, Three.js, etc.
- The popular implementation uses a pre-baked LUT, which is generated by Eary Chow's generator* [Chow 25], and supports SDR output only.
- An HDR extension is a work in progress [Pestaluky 25].


*: Eary Chow's AgX is based on Troy Sobotka's SB2383 [Sobotka 23].  POLYPHONY™
DIGITAL

This tone mapping is used in popular software such as Blender, Godot Engine, and Three.js.

The popular implementation uses a pre-baked look-up table generated by Eary Chow's generator.

An HDR extension is a work in progress.

Khronos PBR Neutral Tone Mapper



- Per-channel mapping with white highlights
- Designed for e-commerce use
 - Simple to implement, good performance
 - Faithful reproduction of color while minimizing HDR highlight artifacts
- Supports only Rec. 709 and SDR output



Khronos PBR Neutral is a tone-mapping operator for e-commerce HDR product views. It keeps mid-tone hues accurate while reducing highlight artifacts, giving faithful, consistent colors.

Note:

The official implementation supports only Rec. 709 and SDR output.

ACES 1.3



- Color volume mapping
- Industry-standard color mapping by the Academy Color Encoding System (ACES)
- A per-channel (RGB) approach with a fitted curve [Hill 16] [Narkowicz 16] is very common in games.

ACES 1.3 is a tone mapping function standardized by the Academy Color Encoding System.

Originally, ACES is a color volume mapping.

Some games use a per-channel approach with fitted S-shaped curve extracted from ACES.

ACES 2.0



- Newer version of the ACES tone mapping
- Changes from 1.3: [Academy 25]
 - Improved consistency of the tone scale
 - Hue skew minimized
 - Harsh clipping improved
 - Easy to use, easy to customize

ACES 2.0 is the latest version of the ACES tone mapping function. This tone mapper uses the JMh color space to perform color calculations for a perceptually natural result. Consistency of the tone scale, harsh clipping, and hue skew are improved compared to the previous ACES tone mapping function.

ACES Small Details

- Implementations are available in CTL (the Color Transform Language).
- An optional configuration is needed to support different peak luminance levels.
 - ACES 1.3: ODT (Output Device Transform)
 - ACES 2.0: Output Transform
- Maps 0.18 mid-gray to 10 to 15 nits depending on the ODT, which is almost 1/1.8 times darker than other tone mappers.
- The 10-nit mid-gray comes from the movie standard [Pines 10].



Implementations of both ACES 1.3 and 2.0 are available in CTL, the Color Transform Language.

To support different peak luminance levels, an optional configuration is needed.

ACES tone mapping maps 0.18 mid-gray to 10 to 15 nits, depending on the output device performance.

As other tone mapping functions map 0.18 mid-gray to 18 nits on an SDR display, the result from ACES tone mapping usually looks darker.

Note:

ACES 1.3 maps 0.18 to 10 nits.

ACES 2.0 mapping changes according to the display device peak luminance.

Tone Mapping Methods by Approach

Global Tone Mapping

- Scale and Clamp
- Reinhard
- Filmic
- ACES 1.3 & 2.0
- AgX
- Khronos PBR Neutral Tone Mapper
- GT Tone Mapping
- GT7 Tone Mapping

Localized Tone Mapping

- Bilateral Grid [Patry 21]
- GI-Based local tone mapping [McAuley 18]

Here is a brief list of tone mapping operators.
Our tone mappings are both global tone mapping operators.

Tone Mapping Methods by Approach

Per-channel Mapping

- Scale and Clamp
- Reinhard
- Filmic (Hable's, AMD)
- GT Tone Mapping

Color Volume Mapping

- ACES 1.3 & 2.0
- AgX
- Khronos PBR Neutral Tone Mapper
- GT7 Tone Mapping



Here is a brief list of tone mapping methods for per-channel mapping and color volume mapping.

We used a per-channel mapping method for GT Tone Mapping, but we moved to color volume mapping in GT7 Tone Mapping. We will talk about the reason in the next section.

Qualitative Comparison of Tone Mapping Methods for Games

Tone Mapping	Approach	Supports Variable HDR Peaks	Ease of Parameter Adjustment
Reinhard	Per-Channel	No	Easy
Filmic (Hable)	Per-Channel	No	Difficult
Filmic (AMD)	Per-Channel	Yes	Difficult
AgX	CVM	No	Fixed (De facto LUT* used)
PBR Neutral	CVM	No	Fixed
ACES 1.3	CVM	Additional Work Needed	Additional Work Needed
ACES 2.0	CVM	Additional Work Needed	Additional Work Needed
GT Tone Mapping	Per-Channel	Yes	Easy
GT7 Tone Mapping	CVM	Yes	Easy

- CVM = Color Volume Mapping
- De facto LUT: As of June 2025, Blender, Godot Engine, and Three.js use AgX LUTs derived from Eary Chow's implementation.

Tone Mapping: Implementation of "GT7 Tone Mapping"

Requirements for Tone Mapping in Gran Turismo

- The game requires a realistic representation of actual vehicles.
- Lighting conditions are always changing.
 - Real and fictional courses
 - Weather and time-of-day changes
- The tone mapping method in Gran Turismo should:
 - Produce images that look natural to the human eye
 - Be consistent under diverse lighting conditions

“GT Tone Mapping” for GT Sport (2016)

“GT7 Tone Mapping” for Gran Turismo 7 (2022)



Gran Turismo demands a highly realistic portrayal of real-world vehicles.

Because lighting conditions change constantly, we needed a tone-mapping method that consistently produces images that look natural to the human eye across diverse scenarios.


To meet this need, we developed two tone-mapping versions: one for Gran Turismo Sport and another for Gran Turismo 7.

Trials before GT Sport


- In 2015-2016, we began working on HDR display support.
- Initially adopted the AMD Filmic curve [Lottes 16] because it supports both SDR and HDR.
- Conducted many trials and identified key points.
 - Linear mid-tone is crucial for realism.
 - High and low luminance areas need an S-shape function.
 - Color clipping is not a problem if the colors are in proper exposure.




We initially adopted AMD's Filmic curve for tone mapping tests. After extensive testing, we identified several key requirements: A linear response within the correctly exposed range is essential for realism, while very bright and very dark regions should follow an S-shaped curve. If exposure is set properly, color clipping is not an issue.



This car paint is out of the gamut of Rec. 709, but color clipping is not a big problem if the picture is in proper exposure.



Here is a sample of our game's rendering output using a simple scale and clamp tone mapping.



Scale and clamp has no S-shaped function, so in a dark environment, it loses contrast in the picture and causes artifacts around the high-brightness area.
This is why an S-shaped function is needed.

Scale and clamp is one of the answers if...

Scale and clamp is enough if dynamic range compression is not required.

- If you are using an ideal display device
- If the picture contains few highlights
- If the picture is complex enough to hide clamp artifacts



For example, this picture is processed using a simple scale and clamp method, but the image has enough complexity to hide clamp artifacts.

Why Design a New Curve?

- The existing tone mapping curves at that time were polynomials.
 - It was difficult to adjust parameters independently.
- To satisfy all of the following requirements at the same time:
 - Perfectly linear mid-tones
 - Director's requests for highlights and shadows
 - Maintaining compatibility with both SDR and HDR
- We decided to design a new curve:
 - "GT Tone Mapping"



The previous tone-mapping curves were polynomial and changed dramatically with small parameter tweaks, making it hard to find one setting that satisfied all needs.

We therefore created a new curve called GT Tone Mapping.


GT Tone Mapping

Introduced in Gran Turismo Sport (2017) [Uchimura and Suzuki 18]

- Per-channel tone mapping method
- Easy parameter adjustment
- Supports both SDR and HDR displays with flexible peak luminance
 - This idea is inspired by several presentations [Lottes 16] [Kawase 16].

GT Tone Mapping is a per-channel tone mapping method that supports HDR displays with various peak brightness levels. The details were explained in the SIGGRAPH ASIA 2018 talk, but today we will give a brief overview as well.

GT Tone Mapping Curve




<https://www.desmos.com/calculator/ribmus8yt4>

POLYPHONY™
DIGITAL

Here is the curve itself. This curve is made up of three sections.


Enhanced Contrast for the Dark Areas



POLYPHONY™
DIGITAL

The toe region.


Linear Response in the Mid-tones



POLYPHONY™
DIGITAL

Linear middle segment.


Smooth Roll-off in the Highlights



POLYPHONY™
DIGITAL

And a smooth shoulder for highlights.

Variable HDR Peak Luminance Support



POLYPHONY™
DIGITAL

Adjusting the maximum-luminance parameter lets the same tone-mapping curve seamlessly handle both SDR and HDR, automatically reshaping itself as the peak luminance changes.

Characteristics of GT Tone Mapping

- The curve is applied independently to the R, G, and B channels.
- There is no gamut compression effect in the low- and mid-luminance regions.
- As the curve flattens in the high-luminance range, the hue twists and shifts toward white.



POLYPHONY™
DIGITAL

GT Tone Mapping applies its curve to the R, G, and B channels individually.

Low- and mid-luminance areas receive no gamut compression.

Hues in the highlights drift toward white.

GT Tone Mapping: Pros/Cons

Pros

- Controllable
- Perfectly linear mid-tones
- Adjustable peak luminance for HDR displays
- Consistent appearance across SDR and HDR

Cons

- Hue twists
- Too desaturated at high brightness

Here are the pros and cons of GT Tone Mapping.

GT Tone Mapping is sufficiently controllable.

Linear mid-tones provide realism to the picture.

Consistent appearance across SDR and various HDR displays is achieved by adjustable peak luminance.

However, color hue twists due to per-channel tone mapping.

Also, the colors become too desaturated at high brightness.

Motivation for New Tone Mapping

Introduced time-of-day and weather changes in Gran Turismo 7.

- Skies from our new renderer
- Sharp reflections and highlights from ray tracing
- Vivid artificial lights in night scenes

Not satisfied with the GT Tone Mapping results.

- A new tone mapping method is needed.
- “GT7 Tone Mapping”



To showcase GT7’s advanced lighting and rendering, we introduced a new tone mapping method that avoids desaturation and hue shifts under high brightness.

Avoiding Hue Twisting

- Hue twist can be avoided by using a luma-chroma separation
 - Yxy, Yuv, ICtCp, etc.
- We want both:
 - High-luminance, hue-twist prevented look
 - “Camera-like” feeling of a hue-twisted look
- Answer: Blend both.



Preventing hue shifts in highlights is essential for the new tone mapping.



Separating luma and chroma eliminates most twisting.

But, at the same time, a slight hue twist adds a natural, camera-like feeling.

We blend both approaches.

Uniform Color Space as Luma-Chroma Separation

- Separates luma and chroma to process each component independently.
- Ensures a more perceptually consistent distribution of differences, enabling reliable mathematical operations.
 - This property is suitable for matching mathematical operations with perceptual experience.
 - Inspired by the Frostbite presentation [Fry 18]



$$\text{Var}\left(\frac{\Delta E_{\text{Perception}}}{\Delta E_{\text{UCS}}}\right) \ll \text{Var}\left(\frac{\Delta E_{\text{Perception}}}{\Delta E_{\text{RGB}}}\right)$$

To achieve more natural adjustments to color and luminance, we perform most processing in a uniform color space, UCS. A UCS typically separates color into luma and chroma components that are evenly spaced perceptually.

Uniform Color Spaces (excerpt)

Year	Common Name	HDR Support
1964	CIEUVW	No
1976	CIELUV	No
1976	CIELAB	No
1998	IPT	No
2006	CAM02-UCS	Yes
2016	ICtCp	Yes
2017	CAM16-UCS	Yes
2017	Jzazbz	Yes
2020	Oklab	No

There are many uniform color spaces in the world.

UCS with HDR

ICtCp [Pytlarz and Atkins 23]

- Based on the PQ (Perceptual Quantizer) curve research
- Efficient computation
- Standardized in 2016
- Used in Frostbite [Fry 18]

Jzazbz [Safdar et al. 17]

- More uniform than ICtCp
- More circular MacAdam ellipses
- Not standardized

We've chosen two UCSs to test:

ICtCp, proposed by Dolby Research, and it is very fast to compute.

Jzazbz, which is more perceptually uniform.

ICtCp

```
// -----  
// ICtCp conversion.  
// Reference: ITU-T T.302 (https://www.itu.int/rec/T-REC-T.302/en)  
// -----  
void  
rgbToICtCp(const float* rgb, float* ictCp) // Input: linear Rec.2020  
{  
    float l = (rgb[0] * 1688.0f + rgb[1] * 2146.0f + rgb[2] * 262.0f) / 4096.0f;  
    float m = (rgb[0] * 683.0f + rgb[1] * 2951.0f + rgb[2] * 462.0f) / 4096.0f;  
    float s = (rgb[0] * 99.0f + rgb[1] * 309.0f + rgb[2] * 3688.0f) / 4096.0f;  
  
    float lPQ = inverseEotfst2084(l);  
    float mPQ = inverseEotfst2084(m);  
    float sPQ = inverseEotfst2084(s);  
  
    ictCp[0] = (2048.0f * lPQ + 2048.0f * mPQ) / 4096.0f;  
    ictCp[1] = (6610.0f * lPQ - 13613.0f * mPQ + 7003.0f * sPQ) / 4096.0f;  
    ictCp[2] = (17933.0f * lPQ - 17390.0f * mPQ - 543.0f * sPQ) / 4096.0f;  
}  
  
void  
ictCpToRgb(const float* ictCp, float* rgb) // Output: linear Rec.2020  
{  
    float l = ictCp[0] + 0.00860904f * ictCp[1] + 0.11103f * ictCp[2];  
    float m = ictCp[0] - 0.00860904f * ictCp[1] - 0.11103f * ictCp[2];  
    float s = ictCp[0] + 0.560031f * ictCp[1] - 0.320627f * ictCp[2];  
  
    float llin = eotfst2084(l);  
    float mlin = eotfst2084(m);  
    float slin = eotfst2084(s);  
  
    rgb[0] = std::max(3.43661f * llin - 2.50645f * mlin + 0.0698454f * slin, 0.0f);  
    rgb[1] = std::max(-0.79133f * llin + 1.9836f * mlin - 0.192271f * slin, 0.0f);  
    rgb[2] = std::max(-0.0259499f * llin - 0.0989137f * mlin + 1.12486f * slin, 0.0f);  
}
```

- Simple matrix operations
- Based on the PQ EOTF
- Designed with hardware implementation in mind

Here is the reference implementation of the ICtCp conversion.
Only two matrix operations and a single PQ EOTF are applied.

Jzazbz

```
// -----  
// Jzazbz conversion.  
// Reference:  
// Muhammad Safdar, Guihua Cui, Youn Jin Kim, and Ming Rommier Luo,  
// "Perceptually uniform color space for image signals including high dynamic  
// range and wide gamut," Opt. Express 25, 15131-15151 (2017)  
// Note: Coefficients adjusted for linear Rec.2020  
// -----  
#define JZAZBZ_EXPONENT_SCALE_FACTOR 1.7f // Scale factor for exponent  
  
void  
rgbToJzazbz(const float* rgb, float* jab) // Input: linear Rec.2020  
{  
    float l = rgb[0] * 0.530004f + rgb[1] * 0.355704f + rgb[2] * 0.086090f;  
    float m = rgb[0] * 0.289388f + rgb[1] * 0.525395f + rgb[2] * 0.157481f;  
    float s = rgb[0] * 0.091098f + rgb[1] * 0.147588f + rgb[2] * 0.734234f;  
  
    float lPQ = inverseEotfSt2084(l, JZAZBZ_EXPONENT_SCALE_FACTOR);  
    float mPQ = inverseEotfSt2084(m, JZAZBZ_EXPONENT_SCALE_FACTOR);  
    float sPQ = inverseEotfSt2084(s, JZAZBZ_EXPONENT_SCALE_FACTOR);  
  
    float iz = 0.5f * lPQ + 0.5f * mPQ;  
  
    jab[0] = (0.44f * iz) / (1.0f - 0.56f * iz) - 1.6295499532821566e-11f;  
    jab[1] = 3.524000f * lPQ - 4.066708f * mPQ + 0.542708f * sPQ;  
    jab[2] = 0.199076f * lPQ + 1.096799f * mPQ - 1.295875f * sPQ;  
}  
  
void  
jzazbzToRgb(const float* jab, float* rgb) // Output: linear Rec.2020  
{  
    float jz = jab[0] + 1.6295499532821566e-11f;  
    float iz = jz / (0.44f + 0.56f * jz);  
    float a = jab[1];  
    float b = jab[2];  
  
    float l = iz + a * 1.386090432715393e-1f + b * 5.884731615611869e-2f;  
    float m = iz + a * -1.386090432715393e-1f + b * -5.884731615611869e-2f;  
    float s = iz + a * -9.601924202631895e-2f + b * -8.118918969560390e-1f;  
  
    float llin = eotfSt2084(l, JZAZBZ_EXPONENT_SCALE_FACTOR);  
    float mlin = eotfSt2084(m, JZAZBZ_EXPONENT_SCALE_FACTOR);  
    float slin = eotfSt2084(s, JZAZBZ_EXPONENT_SCALE_FACTOR);  
  
    rgb[0] = llin * 2.990669f + mlin * -2.049742f + slin * 0.088977f;  
    rgb[1] = llin * -1.634525f + mlin * 3.145627f + slin * -0.483037f;  
    rgb[2] = llin * -0.042505f + mlin * -0.377983f + slin * 1.448019f;  
}
```

- Uses numerically optimized parameters for a uniform distribution
- Avoids hue twists in high-chroma blues
- Achieves a uniform luma distribution via a slightly modified PQ-like curve

Here is the Jzazbz implementation.

Jzazbz uses numerically optimized parameters for a uniform distribution.

GT7 Tone Mapping Strategy

1. Prepare the “twisted” RGB values by processing each channel separately.
2. Convert the RGB values to UCS (Uniform Color Space) values.
3. Compute the “untwisted” RGB values for the highlights.
4. Blend the two results to obtain the final RGB values.



Here is the basic strategy of GT7 Tone Mapping.

First, we prepare the “hue-twisted” RGB values by processing each channel separately.

Second, we convert the RGB values to UCS values.

Third, we calculate the “untwisted” RGB values for the highlights.

Finally, we blend the two results to obtain the final RGB values.

Step 1: Prepare the “Twisted” RGB Values

- Initialize GT Tone Mapping V2 curve with a target peak luminance (peakL)

$$\text{curve}(x) = \text{GTToneMappingV2Curve}_{\text{peakL}}(x)$$

- Apply tone mapping to each RGB channel using the curve


$$R_1 = \text{curve}(R)$$

$$G_1 = \text{curve}(G)$$

$$B_1 = \text{curve}(B)$$

Step 1: Apply per-channel tone mapping using a slightly modified GT Tone Mapping curve to prepare “twisted” RGB values.

GT Tone Mapping V2



- We slightly modified the GT Tone Mapping
- To guarantee convergence

* This curve function itself may be convenient.

<https://www.desmos.com/calculator/bfe382ab4c>

POLYPHONY™
DIGITAL

The original GT Tone Mapping curve did not converge to the peak luminance, resulting in an ill-defined output range.

The new curve, by contrast, is designed with a clearly bounded and well-defined value range.

Note:

The shoulder has a sharp clipping point, but in our use case, this was not a significant issue.

Since TVs also perform some degree of tone mapping, we considered it acceptable to clip the shoulder of the curve.

Step 2: Convert the RGB Values to UCS Values

- Calculate the UCS coordinates for both the original RGB and the twisted RGB

$$\begin{pmatrix} L_u \\ a_u \\ b_u \end{pmatrix} = \text{RGBtoUCS} \begin{pmatrix} R \\ G \\ B \end{pmatrix}$$

$$\begin{pmatrix} L_t \\ a_t \\ b_t \end{pmatrix} = \text{RGBtoUCS} \begin{pmatrix} R_1 \\ G_1 \\ B_1 \end{pmatrix}$$

In step 2, calculate the UCS coordinates for both the original RGB and the twisted RGB.

Step 3: Compute the “Untwisted” RGB Values

- Whiten the RGB values in UCS (make (R_2, G_2, B_2))
 - Calculate the scale to whiten the color according to the luminance in UCS
 - Scale a, b (chromaticity) by them
 - Use “twisted” luminance for smoother highlight

$$scale = 1 - \text{smoothstep}\left(f_s, f_e, \frac{L_u}{\text{UCS}(peakL)}\right)$$

$$\begin{pmatrix} R_2 \\ G_2 \\ B_2 \end{pmatrix} = \text{UCStoRGB} \begin{pmatrix} L_t \\ a_u * scale \\ b_u * scale \end{pmatrix}$$

Next, calculate a whitening scale from the UCS luminance and the peak luminance parameter.

Then, compute the untwisted, desaturated color in UCS coordinates.

Step 4: Blend the Two Results

- Blend twisted and untwisted RGB to obtain the final result.
 - We tried blending in the UCS, but blending RGB looks better to us.
 - We used linear Rec. 2020 space as our RGB color space.
- The ratio lets you control the final appearance of the image.

$$\begin{pmatrix} R' \\ G' \\ B' \end{pmatrix} = \text{lerp} \left(\begin{pmatrix} R_1 \\ G_1 \\ B_1 \end{pmatrix}, \begin{pmatrix} R_2 \\ R_2 \\ R_2 \end{pmatrix}, ratio \right)$$

Finally, we blend the twisted and untwisted RGB values to produce the final image.

The blend is performed in RGB space, although blending in a uniform color space also works.

A sample implementation is provided at the end of these slides.

Sample Implementation

```
float  
SmoothCurve(float x, float a, float b)  
{  
    return 1.0f - smoothStep(x, a, b);  
}  
  
// "GT Tone Mapping" curve with convergent shoulder.  
// -----  
struct GToneMappingCurvev2  
{  
    float peakIntensity_;  
    float alpha_;  
    float midPoint_;  
    float linearSection_;  
    float toeStrength_;  
    float kA_, kB_, KC_;  
  
    void initializeCurve(float monitorIntensity,  
                        float alpha,  
                        float grayPoint,  
                        float linearSection,  
                        float toeStrength)  
    {  
        peakIntensity_ = monitorIntensity;  
        alpha_ = alpha;  
        midPoint_ = grayPoint;  
        linearSection_ = linearSection;  
        toeStrength_ = toeStrength;  
  
        // Pre-compute constants for the shoulder region.  
        float k = (linearSection_ - 1.0f) / (alpha_ - 1.0f);  
        kA_ = peakIntensity_ * linearSection_ + peakIntensity_ * k;  
        kB_ = peakIntensity_ * k * std::exp(linearSection_ / k);  
        KC_ = -1.0f / (k * peakIntensity_);  
    }  
  
    float evaluateCurve(float x) const  
    {  
        if (x < 0.0f)  
        {  
            return 0.0f;  
        }  
  
        float weightLinear = smoothStep(x, 0.0f, midPoint_);  
        float weightToe = 1.0f - weightLinear;  
  
        // Shoulder mapping for highlights.  
        float shoulder = kA_ + kB_ * std::exp(x * KC_);  
  
        if (x < linearSection_ + peakIntensity_)  
        {  
            float toeMapped = midPoint_ * std::pow(x / midPoint_, toeStrength_);  
            return weightToe * toeMapped + weightLinear * x;  
        }  
        else  
        {  
            return shoulder;  
        }  
    }  
};
```

- First, prepare the GT Tone Mapping v2 curve as an S-shaped function.

From here, I will show you some sample implementations.
Since I know that the code text is too small to read, we have
included the code at the end of the presentation slides.

Sample Implementation

```
float framebufferLuminanceTarget_;
float framebufferLuminanceTargetUcs_; // Target luminance in UCS space
GTToneMappingCurveV2 curve_;

float blendRatio_;
float fadeStart_;
float fadeEnd_;

// Initializes the tone mapping curve and related parameters based on the target display luminance.
// This method should not be called directly. Use initializeAsHDR() or initializeAsSDR() instead.
void initializeParameters(float physicalTargetLuminance)
{
    framebufferLuminanceTarget_ = physicalValueToFrameBufferValue(physicalTargetLuminance);

    // Initialize the curve (slightly different parameters from GT Sport).
    curve_.initializeCurve(framebufferLuminanceTarget_, 0.25f, 0.538f, 0.444f, 1.280f);

    // Default parameters.
    blendRatio_ = 0.6f;
    fadeStart_ = 0.98f;
    fadeEnd_ = 1.16f;

    float ucs[3];
    float rgb[3] = { framebufferLuminanceTarget_,
                    framebufferLuminanceTarget_,
                    framebufferLuminanceTarget_ };
    rgbToUcs(rgb, ucs);
    framebufferLuminanceTargetUcs_ =
        ucs[0]; // Use the first UCS component (I or Jz) as luminance
}
```

- Prepare the parameters.
- These are our sample parameters.

Next, we precompute some fixed parameters.

Sample Implementation

```
// Input: linear Rec.2020 RGB (frame buffer values)
// Output: tone-mapped RGB (frame buffer values);
// - In SDR mode: mapped to [0, 1], ready for sRGB OETF
// - In HDR mode: mapped to [0, framebufferLuminanceTarget_], ready for PQ inverse-EOTF
// Note: framebufferLuminanceTarget_ represents the display's target peak luminance converted to a frame buffer value.
// The returned values are suitable for applying the appropriate OETF to generate final output signal.
void applyToneMapping(const float* rgb, float* out) const
{
    // Convert to UCS to separate luminance and chroma.
    float ucs[3];
    rgbToUcs(rgb, ucs);

    // Per-channel tone mapping ("skewed" color).
    float skewedRgb[3] = { curve_.evaluateCurve(rgb[0]),
                          curve_.evaluateCurve(rgb[1]),
                          curve_.evaluateCurve(rgb[2]) };

    float skewedUcs[3];
    rgbToUcs(skewedRgb, skewedUcs);

    float chromaScale =
        chromaCurve[ucs[0] / framebufferLuminanceTargetUcs_, fadeStart_, fadeEnd_];

    const float scaledUcs[3] = { skewedUcs[0], // Luminance from skewed color
                               ucs[1] * chromaScale, // Scaled chroma components
                               ucs[2] * chromaScale };

    // Convert back to RGB.
    float scaledRgb[3];
    ucsToRgb(scaledUcs, scaledRgb);

    // Final blend between per-channel and UCS-scaled results.
    for (int i = 0; i < 3; ++i)
    {
        float blended = (1.0f - blendRatio_) * skewedRgb[i] + blendRatio_ * scaledRgb[i];
        // When using SDR, apply the correction factor.
        // When using HDR, sdrCorrectionFactor_ is 1.0f, so it has no effect.
        out[i] = sdrCorrectionFactor_ * std::min(blended, framebufferLuminanceTarget_);
    }
}
```

- Perform the calculation.
- Use any UCS you prefer.
- Blending in a UCS works as well.
- SDR paper white
 - We assume a paper white = 250 nits
 - In sRGB, 1.0 = 100 nits; we assume 2.5x brightness.
 - We multiply by 1 / 2.5 after SDR tone mapping to adjust.

The final computation is as follows.

Image Comparison

- Here we compare several tone mapping approaches:
 - ACES 1.3 (mapping 0.18 to 18 nits)
 - ACES 2.0 (mapping 0.18 to 18 nits)
 - AgX (from <https://github.com/EaryChow/AgX>)
 - Khronos PBR Neutral Tone Mapper (from <https://github.com/EaryChow/AgX>)
 - GT Tone Mapping
 - GT7 Tone Mapping
- While the results are in SDR, the findings apply similarly to HDR.
- We do not want to label them as “Good” or “Not Good.”
- Each tone mapping method reflects different design priorities.



From this slide, we will compare the results.

These results are all in SDR. However, HDR output leads to similar overall conclusions.


Note:

ACES 1 and 2 map linear 0.18 to 10 nits on an SDR display, so they usually look darker than the other tone mappers.

For this comparison, we scaled them to match the exposure.



First sample is a white car running on an urban highway in Tokyo.



The sky gradation here is our design target for GT7 Tone Mapping.



From left to right:

ACES 1.3,

ACES 2.0,

AgX,

Khronos PBR Neutral,

GT Tone Mapping,

and GT7 Tone Mapping, our recent method.

When the exposure is set correctly, the tone mapping results show few noticeable differences.



When the exposure is increased, the differences become more apparent.



This image shows the results at +2 steps of exposure.

GT7 Tone Mapping demonstrates its characteristics in the blue-sky gradation.

ACES 1.3, ACES 2.0, AgX, and GT Tone Mapping produce desaturated gradations.

The Khronos PBR Neutral result preserves saturation but loses luminance gradation.

GT7 Tone Mapping achieves a balance between desaturation and luminance gradation.




The next sample is a scene featuring an expensive supercar with vivid orange paint.

This car paint is outside the Rec. 709 gamut, so color compression plays a key role here.



GT7 Tone Mapping is designed to preserve these vivid colors.



This is a reference for the actual car paint color.



As you can see, all tone mapping methods have their own characteristics for color compression.
The car paint color produced by the GT7 Tone Mapping result looks good to us.



As the exposure increases, the differences become more noticeable.



Since this car paint is a vivid orange, some tone mappers make it appear pinkish.



Here is a comparison between the real car paint and the in-game rendering again.
Reproducing the actual car paint color is a high priority for Gran Turismo.



The last sample is a sunset scene with brake lamps.
This scene contains a large dynamic range.




It is important to preserve the color and luminance of the brake lamps to ensure that users can play comfortably during night scenes.




In GT Tone Mapping, the brake lamp color shifts to orange. However, in GT7 Tone Mapping, the brake lamp color is preserved.

Hue Twist Flight

GT Tone Mapping



GT7 Tone Mapping



POLYPHONY™
DIGITAL

This diagram shows the color flight paths of GT Tone Mapping and GT7 Tone Mapping.

Each color starts from the border of the Rec. 2020 gamut, and the exposure is swept up.


In GT Tone Mapping, most of the colors converge to cyan, magenta, and yellow due to per-channel tone mapping.

GT7 Tone Mapping avoids hue twisting in the mid-tones and introduces some twist in the highlights.


It also utilizes the display color volume effectively.

Key Differences

GT Tone Mapping



GT7 Tone Mapping



POLYPHONY™
DIGITAL

As you can see in this diagram, the highlight colors in GT Tone Mapping shrink to cyan, magenta, and yellow.

On the other hand, GT7 Tone Mapping utilizes the color gamut more evenly and effectively.

Thanks to the uniform color space, the results of GT7 Tone Mapping are more perceptually uniform.

Hue Twist Comparison

GT Tone Mapping



GT7 Tone Mapping



POLYPHONY™
DIGITAL

Thanks to the uniform color space, GT7 Tone Mapping avoids unwanted hue twisting.

In the diagram, highlight colors in GT Tone Mapping collapse toward cyan, magenta, and yellow.

GT7 Tone Mapping, by contrast, distributes the color more evenly across the gamut.

Look-up Table

- Bake a 3D LUT to accelerate tone mapping
- Apply nonlinear ($\text{pow}(4)$) indexing to reduce banding artifacts

We bake the GT7 Tone Mapping results into a 3D look-up table for faster operation.

We apply a power-of-4 nonlinearity to avoid banding artifacts.

Impact of Sampling Nonlinearity

Resolution 32, pow(x, 4)



Resolution 64, linear



POLYPHONY™
DIGITAL

This sample shows the difference between nonlinear and linear look-up tables.

The sample on the right uses twice the resolution of the table, but banding is still noticeable.

Interpolation is turned off to emphasize the visual differences in this sample.

Extending Tone Mapping with Perception-Aware Physically Based Post-Processing

So far, we've treated tone mapping as a pixel-to-pixel operation, adjusting luminance and color from local physical values. Now we broaden the view to how it interacts with perception over space and time.

Perception-Aware Physically Based Post-Processing

- **Glare (blooming) effects**
 - Expanding the perceptual output range of tone mapping, in the **spatial domain**
- **Auto-exposure and WB (white balance)**
 - Optimizing luminance and color distribution for tone mapping, in the **temporal domain**




As we mentioned at the beginning, we treat glare, auto-exposure, and auto-white-balance as extensions of tone mapping.

Even after physically based rendering was introduced to Gran Turismo, they were still tuned manually for each lighting context, with no shared baseline, causing inconsistency and repeated revisions.

This is why we integrated their control with tone mapping, designing explicit, dynamic links.


Spatial Extension | Glare Effect



POLYPHONY™
DIGITAL

These are fully saturated pixels; all seven look the same.

Spatial Extension | Glare Effect



POLYPHONY™
DIGITAL

Once the glare effect is applied, the difference in luminance becomes visible.

When a pixel is fully saturated, it carries no more information by itself.

By diffusing intensity outward from saturated areas, glare effects add perceptual structure absent in the pixel data.



This is especially powerful when exceptional high-luminance pixels are scattered against a dark background, such as in night scenes.

It has become a very common effect in games today.



The same goes for defocus and motion blur.

Although the physical luminance is reduced due to spatial spreading, the perceptual luminance remains extremely high — because of the visual context present.




In GT, defocus and motion blur are physically based and compatible with HDR.

They made us realize that glare effects — traditionally crafted manually by artists for each scene, often at significant cost — might also be integrated as a physically based perceptual extension.

Glare Effect | Fraunhofer Diffraction

- **Approximation using Fraunhofer diffraction**
 - Depends on the **aperture size, shape, and wavelength** of light
 - Wavelength-integrated PSF for each F-number
 - Mapped to scene luminance via aperture size



POLYPHONY™
DIGITAL

Although perceptual glare is subjective, we used a physically grounded Fraunhofer diffraction model, whose pattern varies with aperture size and thus exposure.

The aperture size is interpreted as a camera exposure setting and mapped to scene luminance.




With this design, the glare effect appears clearly in bright scenes and stays subdued in dark ones.



Additionally, on HDR, glare appears narrower than on SDR since we treat the peak luminance gap as an exposure difference in F-number. Higher peak luminance reduces the need for perceptual exaggeration.

Glare Effect | Implementation



However, the implementation is typical itself, the Fraunhofer diffraction is just approximated by the sum of multiple Gaussian blurs.

Note:

First step: Extract high-luminance pixels from the frame buffer and repeatedly apply Gaussian blur and down-sample to obtain images for each frequency band.

Second step: Weight and composite them to approximate Fraunhofer diffraction.




I manually calibrated 24 parameters for each of the ten F-numbers to match the physically based ground truth — yes, all by hand. It took some doing.



I manually calibrated 24 parameters for each of the ten F-numbers to match the physically based ground truth — yes, all by hand. It took some doing.

Temporal Extension | Auto-Exposure and WB

- Auto-exposure and white balance are forms of **temporal tone mapping**.



Next, automatic exposure and white balance are also common in real-time graphics, and we treat them as dynamic range extension mechanisms over time.

Auto-Exposure and WB | Requirements

- **Time of day: [around 16 EV]**
 - Noon to midnight
 - Extremely rapid changes both at sunrise and sunset
- **Weather: [around 6 EV]**
 - Clear sunny to severe storm
 - Sun suddenly behind clouds
- **Local environment: [up to 10 EV]**
 - Tunnel, underpass, forest, indoor, in-car, etc.
 - Fast cars need a fast response

Fully real-time, continuous, and **recordable**



Gran Turismo 7 introduced dynamic changes in time and weather, dramatically expanding the dynamic range of continuous scenes — creating multi-dimensional contexts for adaptation.

Note: The experience needed to be realistic, yet easy to drive, while ensuring that replays reproduce the same impression as live play.

Auto-Exposure and WB | Requirements



Gran Turismo is a game with many cameras.

There are eight on-board cameras, both front and rear, that can be switched freely during a race.

During replays, the camera is constantly switched every 4 to 5 seconds. In addition, the car being watched can be freely changed by the user.

Therefore, no continuity in camera position or orientation can be expected.

Auto-Exposure and WB | Requirements

- Context matters
 - *“The sun is setting — it’s getting dark.”*
 - *“It looks like the rain is about to stop.”*
 - *“The weather is getting worse.”*
 - *“It’s too dark for daytime — a storm is coming.”*



There is one more important point.

Say, in a 24-hour race, changes in the scene are not mere spectacle — they carry real consequences.

Whether the change is sudden or gradual, subtle or sweeping, it signals a shift in risk.

Such cues should not be veiled by auto-exposure.

Auto-Exposure and WB | Typical Algorithm in Real-time

Measurement & Evaluation

- Measure Scene Luminance
 - Downscale
 - Create Histogram
 - Rejection
- Evaluate Target Exposure

Adaptation

- Exponential Smoothing

This slide explains a typical auto-exposure pipeline.

First, in 'Measurement and Evaluation,' a target exposure is calculated based on the scene's average luminance.

The image is downsampled, and a histogram is generated to reject outlier pixels, such as the sun.

Then, a target exposure is determined that brings the obtained average luminance to the target luminance.

Next, in 'Adaptation,' the exposure transitions smoothly over time. Exponential smoothing is used to create a natural curve that is fast at first and then slows down.

However, this camera-like model alone makes every frame perfectly, properly exposed — even when it shouldn't be.



Life doesn't always want to be properly exposed.

In real life, the night stays dark, and a snowy field or beach stays blinding.

Auto-Exposure and WB | Typical Contextual Control

- **Method**
 - Exposure limitation (absolute clamp)
 - Exposure compensation (relative adjustment)
- **Scope**
 - Global (e.g., time-of-day exposure LUT)
 - Local (per-region via bounding volumes)
- **Key Contexts**
 - Time of day / Weather (e.g., dawn, dusk, overcast)
 - Local environment (e.g., tunnel, forest, in-car)
 - Artistic direction (e.g., calm of shade, tension of a storm)



Most games introduce hand-authored overrides for contextual control.

Exposure limitation and exposure compensation are often used both globally and locally — via bounding volumes like a Post Process Volume in Unreal Engine.

It's flexible and powerful but often tends to be messy; we ourselves once depended on it.


Auto-Exposure and WB | Our Approach

Contexts that require exposure compensation are **simply delays in adaptation** or cases of **reaching physiological limits**.




This time, our core approach was to simulate the physiological response of the human eye to handle these contexts naturally, rather than relying on pre-authored maps.

Adaptation | Recap



Now, let's recap the adaptation mechanism of the human eye.

Adaptation | Recap




There is a significant difference in operating time — the change in receptor sensitivity is remarkably slow.

Adaptation | Reconstruct

Short-Term Adaptation

Pupil and neural gain

- Range: **3-5 EV**
- Adapts in a **few seconds**



Long-Term Adaptation

Receptor sensitivity

- Range: **12 EV**
- Adapts in **tens of minutes**



POLYPHONY™
DIGITAL

We reconstructed adaptation into two stages of the auto-exposure mechanism in our implementation.

Short-term adaptation is modeled after the pupil and neural gain mechanisms.

By limiting its range, it offers a faster onset to handle sudden changes, such as entering a tunnel.

Long-term adaptation is modeled after changes in receptor sensitivity.


It progresses so gradually that it goes largely unnoticed, yet it covers a very broad range.

Note:

The pupil's adjustment range is approximately 3-4 EV, but it is expanded somewhat through the local adaptation of neural gain control.

The effective range of receptors is also adjusted, accounting for the dynamic range of luminance as it appears in-game, though the real effective range is said to reach 10 log units for rods and cones combined.


Schematic Diagram of Dual Adaptations



POLYPHONY™
DIGITAL

Long-term adaptation limits both the range of short-term adaptation and the entire adaptation range.

Schematic Diagram of Dual Adaptations



POLYPHONY™
DIGITAL

Note:

During the day, short-term adaptation pushes darker but always hits the physiological limit set by long-term adaptation — keeping the scene mildly glaring.

Schematic Diagram of Dual Adaptations




POLYPHONY™
DIGITAL

Note:

As evening approaches, short-term adaptation fully engages, allowing the glare to fade.


Schematic Diagram of Dual Adaptations



POLYPHONY™
DIGITAL

Note:
It gets darker and darker.

Schematic Diagram of Dual Adaptations




POLYPHONY™
DIGITAL

Note:

At dusk, delays in long-term adaptation hinder the progress of short-term adaptation.

Schematic Diagram of Dual Adaptations




POLYPHONY™
DIGITAL

Note:

At dusk, delays in long-term adaptation hinder the progress of short-term adaptation.

Schematic Diagram of Dual Adaptations




POLYPHONY™
DIGITAL

Note:

As night falls, short-term adaption is still limited by the physiological constraints of long-term adaption — and the view remains dark.

Schematic Diagram of Dual Adaptations



POLYPHONY™
DIGITAL

Note:

Eventually, dawn arrives. Short-term adaption is still limited by delays in long-term adaption,

Schematic Diagram of Dual Adaptations




POLYPHONY™
DIGITAL

Note:

but because light adaptation is faster than dark adaptation and white balance is different,

Schematic Diagram of Dual Adaptations




POLYPHONY™
DIGITAL

Note:

the perception is not symmetrical with dusk — even at the same solar elevation.

Schematic Diagram of Dual Adaptations



POLYPHONY™
DIGITAL

Note:

the perception is not symmetrical with dusk — even at the same solar elevation.

Implementation Detail | Measuring References

- Reference from frame buffer alone was insufficient.


	Frame Buffer
Spatial	View-local (ITL)
Angular	Directional, within FOV
Value	Average by Histogram
	

So long-term adaptation requires much higher stability.

Luminance measurement from the frame buffer is highly accurate and flexible, but its scope is inherently limited to the in-game camera frustum.

Implementation Detail | Measuring References

- We use multiple references for stable photometry.

	Frame Buffer	Light Probes	Sky Dome	Sunlight	Sky Visibility
Spatial	View-local (TTL)	Position-local	Position-independent	Position-independent	Position-local
Angular	Directional, within FOV	Full-sphere	Hemispherical	Directional	Hemispherical (SH)
Value	Average by Histogram	Integrated	Integrated	Illuminance	Integrated
					



To capture a broader context, we also reference the luminance and color from interpolated light probes, sky-dome luminance and color, and direct sunlight illuminance and color — just as photographers need an exposure meter to know scene-wide information.


We use fixed weights to mix them for each exposure and white-balance separately.

Note:

Sky visibility is derived from spherical harmonics–based ambient occlusion, which estimates the visibility of both the sky-dome and sunlight at each local position.




This was very helpful for the tuning of auto white balance, avoiding overfitting and ensuring accurate color representation of the cars.




This visualizes the temporal changes in our long-term and short-term adaptation, along with the resulting exposure from short-term adaptation.

Scene luminance variation is also shown.


Auto-Exposure Trend Diagram




Auto-Exposure Trend Diagram





Auto-Exposure Trend Diagram



Auto-Exposure Trend Diagram









The exposure and white balance of these images was determined solely by these mechanisms; manual adjustments by artists are no longer needed. Believe it or not, there is only one parameter set throughout the game.

One World, Consistent

- Unified lights, posts, assets
- Visual Cohesion
- Structured diversity



Once these post effects behaved consistently under an explainable model, our standard quietly shifted — from beauty tailored to specific conditions to robustness for system-wide diversity.

Excessive fine-tuning was avoided, assets stabilized, conflicts eased, and the game's look stayed coherent — more than simply time-saving.

Auto White Balance

- Using the same mechanism as auto-exposure
 - Control color temperature and delta instead of luminance
 - Gradually adapt the scene's color temperature to D65
- Measurement of scene color
 - Using the same multiple references as auto-exposure
 - Adding a virtual light for references to stabilize measurement
- Adaptation logic
 - Adapted using similar adaptation logic to avoid abrupt changes
 - Limits based on the AWB specifications of actual cameras



We use the same mechanism as auto-exposure, but instead of controlling luminance, we control the scene's color temperature and delta.

The goal is to slowly adapt the scene's color tone to match D65, which is the standard daylight white.

To measure the color of the scene, we use multiple references—just like in auto-exposure.

Since a measurement result becomes unstable at night, we add a faint virtual light to improve the results.

The adaptation logic is similar to auto-exposure.

It prevents sudden changes in color tone.

Also, we follow limits based on real camera AWB specifications.

Note:

To blend the measurement results, each measured RGB is converted to the Yxy color space and blended in that space.


During blending, luminance-weighted averaging is applied (excluding the virtual light source), after which the result is blended with the ideal light source.

We tried Rec. 2020, Yxy, and XYZ as blending color spaces, and chose Yxy based on parameter tuning and result quality.

Using a uniform color space is a potential future improvement.

The blended Yxy values are then converted to color temperature and color delta, which are used for the subsequent adaptation process.

Overview of Tone Mapping and Post Effects Pipeline



This is the overview of the tone mapping and post-effects pipeline.

Note:

Post-effects not relevant to this talk (e.g., TAA) have been excluded.

Summary

Physically Based Rendering $\neq$ Perceptual Realism

- An accurate simulation does not guarantee correct perception.
 - This is especially important for human-eye-oriented content like Gran Turismo.
 - Photorealistic images have always relied on human vision.



Let me now summarize the key ideas from this presentation.

First and most importantly:

Physically based rendering doesn't always match what people actually see.

Even with an accurate simulation, the final image depends on the display.

This matters a lot in games like Gran Turismo, where we try to match human vision.

But it's also true for other content like live-action films, where photorealism has always depended on how we see light.

Physically Based Rendering $\neq$ Perceptual Realism

- An accurate simulation does not guarantee correct perception.
 - This is especially important for human-eye-oriented content like Gran Turismo.
 - Photorealistic images have always relied on human vision.
- What the user perceives depends only on the display output.

It is important to remember that the user never sees the raw rendering results.
What they see is only what comes out of the display.

Physically Based Rendering ≠ Perceptual Realism

- An accurate simulation does not guarantee correct perception.
 - This is especially important for human-eye-oriented content like Gran Turismo.
 - Photorealistic images have always relied on human vision.
- What the user perceives depends only on the display output.
- Tone mapping connects rendering results to the visual experience.
 - As display technology evolves, tone mapping remains essential for preserving visual intent across diverse hardware.



That is why tone mapping is essential.
It bridges the gap between rendering results and the actual light that reaches the viewer's eye.

Evolution of Displays

- HDR displays are becoming a solution.
 - They can achieve peak brightness over 1,000 nits and support a wide color gamut.
 - This may reduce or eliminate the need for tone mapping.
- SDR displays are still common.
 - Most media today still target sRGB or Rec. 709.
- Tone mapping remains essential.
 - Preserving contrast and hue conveys meaning.
 - It ensures artistic intent across diverse devices.



And, of course, one of the biggest factors affecting perception is the display itself.

Let's now look at how display technology affects tone mapping.

HDR displays are a possible solution.

With peak brightness above one thousand nits and a wide color gamut, they can present rendering results more faithfully.

In the future, tone mapping might even become unnecessary in some cases.

However, SDR displays are still very common.

Most of the images we see on the web, in print, or even in streaming content are still based on sRGB or Rec. 709 standards.


That's why tone mapping is still important.

As long as we preserve relative contrast and hue, the viewer can still understand the visual meaning.

Tone mapping is the key to delivering the intended artistic expression, even when the same image is shown across many different kinds of displays.

It allows us to maintain visual integrity, even in a very diverse hardware environment.

Overall Structure



This slide shows the structure of our physically based tone mapping again.

At the core is what we call "GT7 Tone Mapping," which is based on color volume mapping.

In addition, several post effects are controlled based on physical and psychophysical principles.

Together, we refer to this entire system as "Physically Based Tone Mapping".

Physically Based Tone Mapping in Gran Turismo 7

- Core component: "GT7 Tone Mapping"
 - Nonlinear curve: Sigmoid shape matching human perception
 - Linear region matters: Smooth transition between highlights and shadows
 - Gamut mapping: Necessary for out-of-gamut colors
- Consistent pipeline
 - Integrated with exposure, glare, and display characteristics



As we have seen, tone mapping remains essential today because of the wide variety of displays.

Previously, we introduced the details of our "GT7 Tone Mapping." Here, I would like to briefly summarize its core elements.

We use a sigmoid-shaped nonlinear curve to match human brightness perception, with a linear region that ensures smooth highlight–shadow transitions.

Gamut mapping is also essential for out-of-gamut colors on the target display.

However, tone mapping does not work in isolation.

It is part of a consistent pipeline that integrates tightly with our exposure control, glare rendering, and display characteristics.

These components work together to produce stable, perceptually grounded output across a wide range of viewing environments. Altogether, we refer to this as our "Physically Based Tone Mapping."

Future Work

- Modeling
 - Modeling the differences between camera sensors and the human visual system
 - High-precision modeling of glare, automatic exposure, and white balance adjustment, and integration into tone mapping
 - Modeling the mesopic vision (e.g., the Purkinje shift)
 - Modeling diversity in color perception

Let us briefly talk about future work.

Modeling remains a challenge.

We need to better match camera sensors and human vision, especially for glare, exposure, and white balance.

Also, mesopic vision and differences in color perception should be considered.

Future Work

- Modeling
 - Modeling the differences between camera sensors and the human visual system
 - High-precision modeling of glare, automatic exposure, and white balance adjustment, and integration into tone mapping
 - Modeling the mesopic vision (e.g., the Purkinje shift)
 - Modeling diversity in color perception
- Parameter adjustment
 - Automatic adjustment across a wide range of display devices
 - Real-time adjustment in VR environments using eye tracking

Next, parameter adjustment must support a wide range of displays.

In VR, real-time tuning based on eye tracking is especially important.

Future Work

- **Modeling**
 - Modeling the differences between camera sensors and the human visual system
 - High-precision modeling of glare, automatic exposure, and white balance adjustment, and integration into tone mapping
 - Modeling the mesopic vision (e.g., the Purkinje shift)
 - Modeling diversity in color perception
- **Parameter adjustment**
 - Automatic adjustment across a wide range of display devices
 - Real-time adjustment in VR environments using eye tracking
- **Quantitative link**
 - Establishing a more quantitative link between physical measurements and subjective visual evaluation

Finally, the quantitative link.

It's still hard to connect physical data with how people subjectively see and feel images.

Conclusion for Today

- Why physically based tone mapping works:
 - Accounting for physical, psychophysical, and display characteristics
 - Enabling stable realism in Gran Turismo's human-eye-oriented rendering
 - Isolating and solving visual issues in other use cases



To finish, here's what we learned:

Tone mapping that includes both physical and perceptual effects works well for achieving human-eye realism in Gran Turismo.

We believe this approach could also be useful beyond Gran Turismo.

Even in games that don't aim for realism, using high-quality tone mapping may help to find and fix visual problems.

Conclusion for Today

- Why physically based tone mapping works:
 - Accounting for physical, psychophysical, and display characteristics
 - Enabling stable realism in Gran Turismo's human-eye-oriented rendering
 - Isolating and solving visual issues in other use cases
- Tone mapping remains a foundation in games.
 - Needs to adapt to evolving display technologies
 - Examples: new console hardware, smartphones, HDR, Mini LED, etc.
 - Must support diverse and rapidly changing content



In that sense, tone mapping is not just a finishing touch. It is a foundation.

As displays and content types evolve, tone mapping must adapt while preserving artistic intent.

That is why we believe physically grounded tone mapping will continue to be essential in the future.

Conclusion for Today

- Why physically based tone mapping works:
 - Accounting for physical, psychophysical, and display characteristics
 - Enabling stable realism in Gran Turismo's human-eye-oriented rendering
 - Isolating and solving visual issues in other use cases
- Tone mapping remains a foundation in games.
 - Needs to adapt to evolving display technologies
 - Examples: new console hardware, smartphones, HDR, Mini LED, etc.
 - Must support diverse and rapidly changing content

See the [Physically Based Shading](#) site for the full version of these slides, including many bonus slides!

The presentation will end here — but for the full version of these slides, including many bonus slides, please visit the [Physically Based Shading](#) site.

Acknowledgements

- Polyphony Digital
 - Kazunori Yamauchi
 - Shuichi Takano
- Additional Thanks
 - Stephen Hill
 - Stephen McAuley
 - Troy James Sobotka



Thank you for your attention.

Bibliography

- [Academy 21] ACES 1.3 is available, <https://acescentral.com/aces-1-3-is-available/>
- [Academy 25] Academy Color Encoding System, <https://github.com/ampas/aces>
- [Blackwell 1946] Contrast Thresholds of the Human Eye, Journal of the Optical Society of America 36, 624-643, 1946
- [Burkhardt 94] Light Adaptation and Photopigment Bleaching in Cone Photoreceptors in situ in the Retina of the Turtle, Journal of Neuroscience 14, 1091–1105, 1994
- [Burley 12] Physically Based Shading at Disney, Practical Physically-Based Shading in Film and Game Production, ACM SIGGRAPH 2012 Courses 10, 1–7, 2012
- [Carucci 06] HDR Meets Black & White 2, GDC 2006
- [Chow 25] AgX, <https://github.com/EarlyChow/AgX>
- [Fry 17] High Dynamic Range Color Grading and Display in Frostbite, GDC 2017
- [Hable 10] Filmic Tonemapping Operators, <http://filmicworlds.com/blog/filmic-tonemapping-operators/>

Bibliography

- [HGiG 19] For Better HDR Gaming, <https://www.hgig.org/>
- [Hill 16] ACES.hls, <https://github.com/TheRealMJP/BakingLab/blob/master/BakingLab/ACES.hls>
- [Hirai and Takano 22] Real-Time Ray Tracing Techniques and Mesoscale Material Representation in Gran Turismo 7 [in Japanese], CEDEC+KYUSHU 2022
- [Hurter and Driffield 1890] Photo-Chemical Investigations and a New Method of Determination of the Sensitiveness of Photographic Plates, Journal of the Society of Chemical Industry 9, 455–469, 1890
- [ICC 22] Specification ICC.1:2022 (Profile version 4.4.0.0), International Color Consortium
- [Ignite 25] XDisplayHdrModelInfo, <https://learn.microsoft.com/en-us/gaming/gdk/docs/reference/system/xdisplay/structs/xdisplayhdrmodelinfo>
- [Jakob et al. 22] Mitsuba 3, <https://www.mitsuba-renderer.org/>, 2022
- [Kawase 16] Theory and Practice for HDR Support in Silicon Studio [in Japanese], CEDEC 2016
- [Kunkel and Reinhard 2010] A Reassessment of the Simultaneous Dynamic Range of the Human Visual System Proceedings of the 7th Symposium on Applied Perception in Graphics and Visualization, 17–24, 2010.

Bibliography

- [Lalish 24] Khronos PBR Neutral Tone Mapper, <https://modelviewer.dev/examples/tone-mapping>
- [Lottes 16] Advanced Techniques and Optimization of VDR Color Pipelines, GDC 2016
- [McAuley 18] The Challenges of Rendering an Open World in Far Cry 5 , SIGGRAPH 2018 Course: Advances in Real-Time Rendering in Games, 2018
- [McTaggart et al. 06] High Dynamic Range Rendering in Valve's Source Engine, ACM SIGGRAPH 2006 Courses, 7-es, 2006
- [Narkowicz 16] ACES Filmic Tone Curve, <https://knarkowicz.wordpress.com/2016/01/06/aces-filmic-tone-mapping-curve/>
- [Patry 21] Real-Time Samurai Cinema, SIGGRAPH 2021 Course: Advances in Real-Time Rendering in Games, 2021
- [Pestaluky 25] An Adjustable Tonemapping Curve for Variable / Extended Dynamic Range, <https://allenwp.com/blog/2025/05/29/allenwp-tonemapping-curve/>
- [Pines 10] From Scene to Screen, SIGGRAPH 2010 Course: Color Enhancement and Rendering in Film and Game Production, 2010

Bibliography

- [Pytlarz and Atkins 23] Perceptual Color Volume, Dolby Professional White Paper, Version 7.1
- [Reinhard et al. 02] Photographic Tone Reproduction for Digital Images, ACM Transactions on Graphics 21, 267–276, 2002
- [Safdar et al. 17] Perceptually uniform color space for image signals including high dynamic range and wide gamut, Optics Express 25, 15131-15151, 2017
- [Sobotka 22] AgX, <https://github.com/sobotka/AgX>
- [Sobotka 23] SB2383-Configuration-Generation, <https://github.com/sobotka/SB2383-Configuration-Generation>
- [Stockman et al. 06] Human cone light adaptation: From behavioral measurements to molecular mechanisms, Journal of Vision 6, 5, 2006
- [Suzuki and Yasutomi 23] Realistic Real-time Sky Dome Rendering in Gran Turismo 7, GDC 2023
- [Synopsys 25] Optical Scattering Measurements & Equipment, <https://www.synopsys.com/optical-solutions/scattering-measurements.html>

Bibliography

- [Uchimura 23] Reproducing the Material of That Car!? BRDF Fitting Using a Differentiable Renderer [in Japanese], CEDEC 2023
- [Uchimura and Suzuki 18] Practical HDR and Wide Color Techniques in Gran Turismo SPORT, SIGGRAPH Asia 2018 Courses 15, 1–4, 2018
- [Watson and Yellott 12] A unified formula for light-adapted pupil size, Journal of Vision 12, 12, 2012
- [X-Rite 25] MA-T12: Fully Characterize Complex Material Finishes, <https://www.xrite.com/categories/portable-spectrophotometers/ma-family/ma-t12>

