Chromaticity Gradient Mapping for Interactive Control of Color Contrast in Images and Video

Ruyu Yan
ruyu.yan@princeton.edu
Princeton University
Princeton, New Jersey, USA

Jiatian Sun
jiatians@cs.cornell.edu
Cornell University
Ithaca, New York, USA

Abe Davis
abedavis@cornell.edu
Cornell University
Ithaca, New York, USA

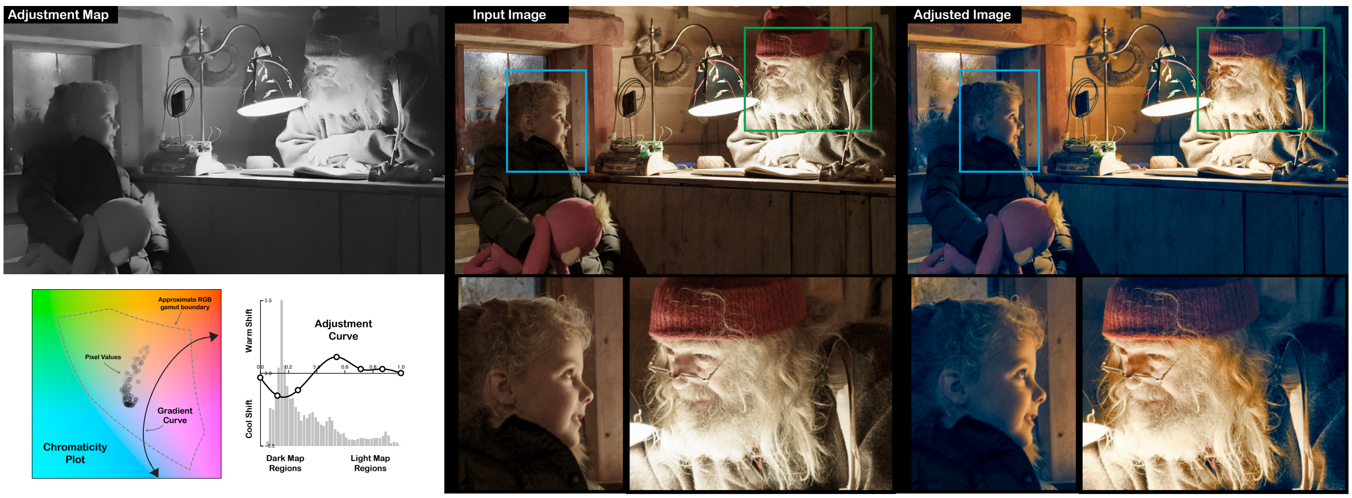

Figure 1: Contrast Enhancement with Chromaticity: We adjust per-pixel chromaticity along a gradient curve according to a user-controlled adjustment curve. This method let us enhance the perceived dynamic range and detail at a selected bandwidth without adjusting the illuminance of an image. Image credit: signatureedits.com.

ABSTRACT

We present a novel perceptually-motivated interactive tool for using color contrast to enhance details represented in the lightness channel of images and video. Our method lets users adjust the perceived contrast of different details by manipulating local chromaticity while preserving the original lightness of individual pixels. Inspired by the use of similar chromaticity mappings in painting, our tool effectively offers contrast along a user-selected gradient of chromaticities as additional bandwidth for representing and enhancing different details in an image. We provide an interface for our tool that closely resembles the familiar design of tonal contrast curve controls that are available in most professional image editing software. We show that our tool is effective for enhancing the perceived contrast of details without altering lightness in an image and present many examples of effects that can be achieved with our method on both images and video.

UIST '24, October 13–16, 2024, Pittsburgh, PA, USA

ACM ISBN 979-8-4007-0628-8/24/10
https://doi.org/10.1145/3654777.3676340

CCS CONCEPTS

• **Human-centered computing** → **User interface programming**; • **Computing methodologies** → **Graphics systems and interfaces**; **Image processing**.

ACM Reference Format:
Ruyu Yan, Jiatian Sun, and Abe Davis. 2024. Chromaticity Gradient Mapping for Interactive Control of Color Contrast in Images and Video. In *The 37th Annual ACM Symposium on User Interface Software and Technology (UIST '24), October 13–16, 2024, Pittsburgh, PA, USA.* ACM, New York, NY, USA, 16 pages. https://doi.org/10.1145/3654777.3676340

1 INTRODUCTION

> *"I want to do with colour what they do in black and white..."*
>
> — Paul Cézanne[1]

Most digital images represent color as the sum of three channels (red, green, and blue) stored in three separate 8-bit integers per pixel. This creates a palette of just 256 or fewer possible brightness variations per representable color, which poses a challenge for content creators who are often tasked with simultaneously conveying detail in both bright and dark regions of an image. At the same time, an 8-bit color pixel can take on 2^{24} (over 16 million) distinct values, with most of that representational power dedicated

[1] Quoted by Maurice Denis in"Cézanne" from *Modern Art and Modernism (1982)* [24]

Figure 2: Color Contrast for Detail Enhancement in Art: On the left is the painting *Sunrise* by Claude Monet. On the right, we see that visualizing the painting in lightness renders many of its details nearly invisible. In particular, the details of the sun and its reflection on the water are conveyed almost entirely through variations in chromaticity.

to variations of hue and saturation. Our work focuses on helping content creators repurpose some of that power toward conveying tonal detail (i.e., differences in brightness) in an image. Our strategy for this is inspired by clever uses of color contrast in the arts and data visualization. The idea is to map variations of image lightness to adjustments along some specified gradient of chromaticities. We call this technique *chromaticity gradient mapping* (CGM). In this paper, we examine how artists have employed similar mappings of chromaticity throughout art history, and what makes such mappings difficult to explore with current digital tools. We then derive a method and interface for defining and controlling such mappings in digital images. We also present three small user studies: a perceptual study comparing the perceived contrast of adjustments produced using our tool with those produced by traditional tonal contrast adjustments, a user interaction study observing how artists approach our tool compared to their familiar software, and a comparison study evaluating our interface against traditional adjustment curves applied to the chromaticity channels of an image. Finally, we demonstrate a wide range of adjustments achieved with our interface that are difficult or impossible to achieve with other existing tools.

1.1 Vocabulary

Our work brings together ideas from many disciplines: art history, information visualization, color science, and psychophysics, as well as computer graphics and signal processing. The interdisciplinary nature of this endeavor calls for establishing clear vocabulary from the outset, especially as several important terms hold conflicting meanings across different related fields. For this text, most of our terminology will be based on the International Electrotechnical Commission (CIE) standards [18]. In particular, *chromaticity* will refer to the properties of color orthogonal to luminance. *Lightness* will refer to a measure of luminance, noting that the same term may refer to color saturation in some literature on painting. We adopt the meaning of *tone* used in photography, which refers to luminance across different spatial scales, noting that this term also has different meanings in literature on painting. The term "saturation" means different things in imaging and color science, so we will use the term "clamping" as a substitute for the imaging definition, which describes the effect of an integrating sensor reaching its maximum representable value (e.g., "pixel saturation" in a camera sensor).

We refer to the more general practice of mapping lightness variations to changes in chromaticity as chromaticity mapping, and to the specific technique we propose for doing this as chromaticity gradient mapping. Note that while we introduce the term "chromaticity mapping", the practice is discussed at length using a variety of subjective terms in writings on art and art history (Section 3).

Finally, we note that much of the theoretical machinery we use to quantify image detail in this paper comes from comparatively recent work on signal processing. We provide an overview and references for this theory later, noting for exposition that "detail" will refer to measures of contrast over different spatial and intensity scales (e.g., separating "fine-scale detail" from "coarse structure").

2 RELATED WORK

2.1 Color Perception & Color Spaces

Much of our understanding of human color perception builds on Maxwell's color matching experiments of the 1850s [41], which helped quantify the human color sensitivity spectra. These spectra form the basis of a surprisingly predictive linear model for human perception of chromaticity. While subsequent research has explored a wide range of questions related to human color perception, our work is most heavily influenced by that on designing perceptually-motivated color spaces, which serve as an intermediary between the photometric analysis of light and predictive models of human perception. Most of our analysis will build on two color spaces, in particular: **sRGB** [31], which defines a photometric standard for interpreting the pixels of most digital (RGB) images, and **CIE** $L^*a^*b^*$ space [18] (sometimes "CIELAB", or simply "Lab"), which uses a *lightness* (L^*) channel to measure luminance and two additional channels (a^*b^*) that measure chromaticity using a parameterization that incorporates the opponent color model of human vision [30]. CIELAB is also designed to have L^2 Euclidean norms that approximate the perceptual distance between colors, making it a useful representation for reasoning about color contrast. Note that chromaticity, as measured by the a*b* dimensions in CIELAB, is distinct from the related concepts of *saturation* (sometimes called "chroma") and *hue* [32, 49]; at a very high level, one can think of a^*b^* and hue/saturation as different parameterizations of the same variations, but this analogy is not a perfect one.

2.2 Information Visualization

This work draws inspiration from design principles for palette and color map selection in data visualization. The challenge of using a limited palette to visualize differences in data is closely related to that of using a limited dynamic range to visualize image detail. In both cases, information encoded in variations of a signal—whether that signal is an image or more abstract data—must be mapped to a limited set of visual channels. Also in both cases, the relative importance of variations in the input may depend on the goals of a subjective user, and the evaluation of variations in the output is a product of human perception.

The question of how to use color for information visualization has been explored from a variety of directions (see [29, 47, 53] for a survey). Much of the work in this area deals with understanding how color interacts with other visual channels for encoding data. Two channels are considered separable if selective attention to one does not interfere with the other, and integral if it does [37, 48]. Color itself can be further separated into multiple visual channels (as we do with lightness and chromaticity in this work) with empirical studies confirming that these channels are at least partially separable [14, 37, 39]. Our work leverages this separability to map details represented in the lightness channel to variations represented in chromaticity.

One paper with a surprising connection to our own comes from Rogowitz and Kalvin [45], who propose a quick method for evaluating perceptual color maps by applying them to grayscale images of human faces and asking users to rate the rendered images on naturalness. They find a strong correlation between perceived naturalness and the degree to which the applied color map is monotonic in lightness. Our work maps lightness to luminance-neutral adjustments along a gradient of chromaticities, which effectively takes a more general class of spatially varying multi-channel color mappings and constrains it to guarantee the monotonic luminance property identified in their findings.

2.3 Detail Manipulation & Tone Mapping

Parts of the theory underlying our method come from previous research on HDR tone mapping, which focuses on mapping high-precision image data to lower-precision display formats. Global mappings (those that treat each pixel independent of its neighboring pixels) tend to result in a loss of detail, as ranges of distinct input lightness values will inevitably map to the same output value. Efforts to find mappings that balance the preservation of fine-scale image details with global contrast led to various strategies for quantifying detail across different scales of an image [6, 8, 21–23, 35]. Work in this space is also closely related to methods for fast edge-aware smoothing [3, 17, 22, 25]. Our algorithm builds on that of Paris et al [42] on Local Laplacian Filtering, which we repurpose for controlling chromaticity. Our implementation is based on the fast approximation introduced by Aubry et al [7].

In Section 5.3 we also show how our method can be applied to compensate for tonal contrast lost in HDR tone mapping. This idea was previously explored by [50], which uses changes in color saturation to compensate for lost tonal contrast. We discuss the limits of using color saturation as a proxy for color contrast later in the paper. Previous research has also focused on the role of color and human perception in HDR tone mapping, including [36] which models the Purkinje effect for tone mapping low-light photographs, and [40] on color correction. Some prior works have also explored interactive methods for tone mapping of lightness [20, 38].

2.4 Recoloring

Several comparatively recent works deal with colorization [28], style transfer [27, 54], or exemplar-based color grading [11], which can achieve similar effects to our approach given a suitable exemplar image. These approaches are focused on automatically inferring color changes based on a particular objective function rather than offering and parameterizing a creative space of variations for users to explore. More similar in spirit to our work is research on creative systems for color manipulation, e.g., [15, 16, 46, 52]. Our work focuses more specifically on the connection between chromaticity adjustments and tonal detail than any previous work, and is, to our knowledge, the first to explicitly support mapping of detail represented in the lightness channel of an image to contrast represented in chromaticity.

3 CHROMATICITY MAPPING IN ART

Artists have used chromaticity as an alternative visual channel to convey variations of luminance for centuries. Just as pixels have only a limited number of bits, a painter's palette can only represent a limited number of human-distinguishable colors, and a review of art throughout history will find chromaticity used in various creative ways to counter this limitation. The strategy is well-documented by art historians—we review a selection of illustrative examples here and refer to [10, 26, 51] for more comprehensive histories.

3.1 Chiaroscuro Paintings

The use of chromaticity mapping in art evolved over time. Earlier examples generally used the technique to approximate real perceptual effects, making its use akin to an optical illusion. This is particularly noticeable in Baroque chiaroscuro paintings, which are characterized by heavy contrast of light and shadow (see Figure 3). Chromaticity Mapping can be seen in the work of artists like Caravaggio and Rembrandt, often used to emphasize contrast between light and dark regions of candlelit scenes. The resulting colors—white highlights that fade into warmer, richer colors with distance from the candle—appear quite natural. However, this pattern of color makes little photometric sense: light does not naturally change its wavelength with distance from the source. Here, the artist is mimicking the effect of a much brighter subject on perceived chromaticity due to partial clamping.

(a) Reference Painting (b) Image Editing with Our Tool

Figure 3: Chromaticity Mapping in Chiaroscuro Painting: In the chiaroscuro painting *Vegetable and Fish Night Market in The Hague* by Petrus van Schendel in (a), we see a shift in chromaticity for conveying different luminance in the scene. Our method can be used to explore the application of similar effects to images, as shown in (b).

Figure 4: Substituting variations of chromaticity for variations of lightness: (right) *Mardi Gras* by Cezanne, painted in 1888. The shading of Pierrot's outfit shows an example of substituting cooler chromaticities for darker luminance in the scene. The middle image shows a similar effect achieved by coupling the chromaticity adjustments of our method with tonal contrast lost in an HDR tone mapping operator when it is applied to the input image shown on the left—doing mathematically exactly as Cezanne described in the quote that opens this paper. Image credit: signatureedits.com.

3.2 "Accidental Colors", Synthetic Blue, & Impressionist Painting

In 1743, mathematician Georges-Louis Leclerc, Comte de Buffon, a contemporary of many Baroque painters, described the blue and green appearance of shadows cast near dawn and dusk as "accidental colors" [26], noting them for their exceptional beauty. However, most blue pigments of the time were derived from the precious stone Lapis Lazuli, making blue paint rare and expensive, which may account for why most examples of chromaticity mapping from this time focused on warm hues. The 1820s brought about the discovery of a much cheaper process for synthesizing blue pigments, opening up a much larger space of chromaticities for artists to experiment with [43]. This coincided with the rise of Impressionism, and with it some of the most recognizable uses of chromaticity mapping in art. Many works, like Vincent van Gogh's *Starry Night*, use strong contrast in chromaticity as a replacement for variation in luminance. The perhaps most extreme example of this is Claude Monet's painting *Sunrise*, which manages to convey the brightness of a rising sun almost exclusively through chromaticity. Figure 2 shows that visualizing the lightness channel of Monet's painting renders the sun and its reflection near-invisible, demonstrating Monet's masterful grasp of human color perception. This is all the more remarkable considering that the work was created more than half a century before the first perceptual color space (a precursor to CIELAB) was defined [1].

Between the subtle use of chromaticity mapping in chiaroscuro and its more extreme application by Monet, we find the work of other impressionists. For example, Cezanne's depiction of Pierrot's outfit in the 1888 painting *Mardi Gras* (Figure 4, right) shows cooler chromaticities substituted for shading. Cezanne's own description of chromaticity mapping in his work serves as a clear and succinct statement of the motivation for this paper: "*I want to do with colour what they do in black and white*" [24]. To interpret it from an image-editing perspective, we want to enable the user to intuitively:

- Choose a color gradient in the 2D chromaticity space.
- Control the intensity and level of detail for contrast enhancement.

In the following sections, we will explain what makes this task difficult–often impossible, even–with existing digital editing tools. Then we will build on these observations to motivate the design of our own method and interface.

4 CONTRAST ADJUSTMENT

Most digital tools for manipulating contrast operate independently on one visual channel, meaning that input values of that channel are mapped directly to output values of the same channel. The reason for this is simple: the concept of contrast describes how values in an image differ, and considering multiple channels leads to exponential increase in the number of ways two pixel values can be different. This complicates the design of tools for manipulating contrast in chromaticity, which is fundamentally a two dimensional domain. In this section we review how contrast is controlled in the more traditional intra-channel setting. We then examine the limitations of applying existing approaches to chromaticity (e.g., using $a*$ and $b*$ adjustment curves), and establish a useful connection with comparatively recent developments in HDR tone mapping and style transfer. Finally, we review popular image and video editing tools that artists commonly use to manipulate color contrast.

4.1 Tonal Contrast & Tone Adjustment Curves

Tonal contrast refers to differences in lightness across different regions of an image. Regardless of how we choose to quantify it, absolute tonal contrast can always be scaled by simply scaling the value of each pixel in an image, which is analogous to adjusting the exposure and black level in a camera. However, when we limit pixel values to a finite range (e.g., the brightest and darkest values that a screen can display) this places an upper bound on how much contrast we can create in an image: as soon as two pixel values reach the upper and lower ends of our range, any further adjustment is guaranteed to bring at least two of our pixels closer together in value. Adjustment curves provide an elegant and compact way to represent global lightness adjustments subject to this limitation. A tone adjustment curve $f_L : \mathbf{L}_{in} \mapsto \mathbf{L}_{out}$ maps the 1D domain of normalized lightness values $L \in \{0, 1\}$ to itself. If we visualize this mapping as a curve defined by the points $[x = \mathbf{L}_{in}, y = \mathbf{L}_{out}]\ \forall L \in \{0, 1\}$, we can derive several important properties by looking at its shape (see Figure 5). In particular, continuity ensures that adjacent input lightness values will map to adjacent output values, while a monotonically increasing curve is guaranteed to preserve the ordering of lightness values and therefore the direction of edge gradients in the image (analogous to the predictor of naturalness found in [45]). Visually, the slope of our curve at each input brightness tells us how our adjustment changes contrast around certain tones in our image. A high or low slope at the start of our curve will increase or decrease contrast in dark regions, but maintaining monotonicity will ensure that the opposite happens in bright regions.

While tone (lightness) curves are by far the most common type of adjustment curve, the same approach can be applied to other

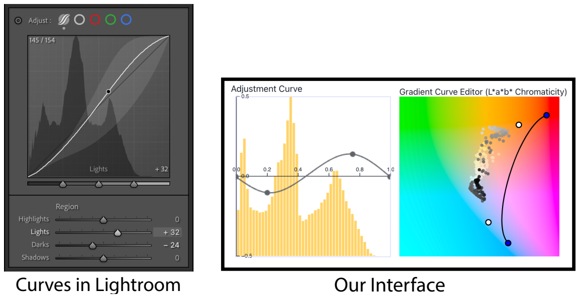

Figure 5: Adjustment Curve Interface: Similar to the *Curves* tool in Lightroom, our chromaticity adjustment curve interface shows a histogram of the filtered lightness channel of the image behind the curve. It lets the user adjust the shift distance along their selected chromaticity gradient, controlled by the cubic Bézier spline in the gradient curve editor, at different lightness values.

visual channels to control other types of contrast as well. However, understanding and visualizing adjustment curves for other channels can be more difficult than it is for lightness.

4.2 Histograms, Adjustment Maps, & The Span of an Adjustment Curves

It can be informative to interpret the adjustment curve for a target visual channel as a representation of our image in some alternative linear basis. Each position along the independent (input) axis of our curve corresponds to the set of pixels in our input that share a given value. Our curve itself then provides the coefficients for a simple linear combination of these sets, which we can think of as basis vectors that span all of the edits we can make with the provided tool. This offers a concrete way to reason about what edits are difficult or even impossible to make with a given combination of curve tools.

The basis vectors of an adjustment curve for a given visual channel are the level sets of that channel in our input. This means that the span of our adjustment curve depends on the number and distribution of unique values in our input. In other words, the expressive range of our curve tool depends on the variations present in our input. We can visualize this in two ways. First, we can visualize the dimensionality of our basis by looking at the sparsity of a histogram built on input values. This histogram is often displayed behind the adjustment curve during editing (e.g., the visualizations used in Photoshop and our own tool are displayed in Figure 5). Second, we can visualize the details spanned by our basis by simply displaying a spatial map of our input channel values. We call this spatial map an *adjustment map* for reasons that will be made clearer later in the paper. Note that the importance of visualizing the adjustment map for a given curve depends on how easy it is to deduce the target channel from simply looking at the image. This makes it less necessary for a tone adjustment curve than for curves applied to other visual channels.

Analyzing the linear basis of a given adjustment curve has another advantage, which is that we can compute the optimal fit of one curve to adjustments made in another curve. This can be thought of as projecting an edit from one tool into the closest edit one can make in another, which we can use to reason about the uniqueness of edits made in a given tool (see Section 4.4.2).

4.3 Global Contrast vs Detail

Coupling the adjustment of different pixels that share a common input value can be very limiting in scenes with detail distributed across a large bandwidth of spatial frequencies. To see this, consider the simple 1D image signal shown in Figure 6, which is constructed by adding *detail* components at different spatial frequencies to a broad-spectrum *edge* component. We can always increase the global contrast of a signal by applying a common S-shaped adjustment curve centered at its median value.[2] However, observe what this does to the signal shown in the top right of Figure 6. If we consider the left and right halves of the adjusted signal separately, the contrast of each region has been lowered significantly. If we relate this back to our original input components, the original detail information has been suppressed in favor of edge contrast. We can avoid this by replacing contrast enhancement with detail enhancement, which approximates applying contrast adjustments to each of the components in our signal individually before they are combined (see Figure 6 bottom right). Detail enhancement has been explored at length in work on edge-preserving filters [17, 21, 42]. Independently manipulating detail at different spatial scales is straightforward in the frequency domain [13], but suffers from ringing artifacts caused by the Gibbs phenomenon, often referred to as "halos". Local laplacian filtering [42] is a current state-of-the-art technique for dealing with such artifacts. The basic idea is to optimize separate contrast adjustments for each location and scale in an image before combining these local adjustments with Laplacian pyramid blending [12]. A fast approximation of this approach introduced by [7] is the basis for several of the adjustments available in Adobe Lightroom [4], (e.g., the "Presence" control sliders).

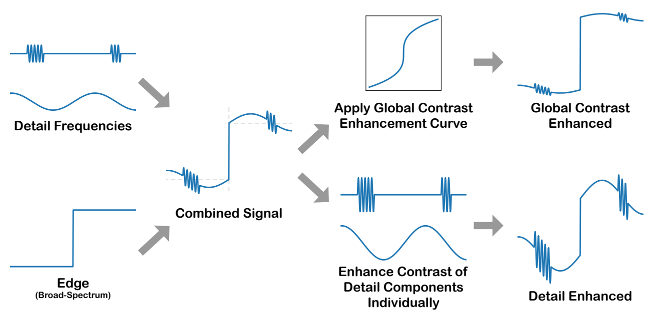

Figure 6: Contrast Enhancement Curves vs Detail Enhancement: A signal can be represented as the combination of detail frequencies and broad-spectrum edges. There is a tradeoff in contrast enhancement for these two components. Applying a global contrast enhancement curve would exaggerate the edges, but suppress the detail; enhancing the detail would create ringing artifacts sometimes referred as "halos".

[2]This is true up to clamping (i.e., repeated application of such a curve will converge to a binary image).

Some of the early examples of chromaticity mapping we see in the arts merely use chromaticity to enhance global contrast (e.g., chiaroscuro with a single light source). However, most of the more recent and extreme examples use chromaticity to selectively replace lightness variation at certain scales. To accomplish this, we need to drive chromaticity adjustments with a selectively filtered version of lightness.

4.4 Tools for Color Contrast Adjustment

Now that we have established how contrast is represented and controlled in lightness, we can use this as a lens to examine existing tools for editing color contrast. Image editing software like Adobe Photoshop [5] and Adobe Lightroom [4] offer a variety of tools for this purpose. Based on our own research and interviews with a small group of 6 design students and professionals, we identified what appear to be the most commonly used tools available today, which we examine here.

4.4.1 Saturation Adjustments. By far the most common strategy for manipulating color contrast is to adjust the saturation of an image—i.e., scale the S coordinate of each color's representation in HSV or HSL space [32, 49]. This strategy is quick and simple but also quite limited. Saturation adjustments map input saturation to new chromaticities, which we can visualize as a gradient over CIE a*b* color space. Figure 7 shows the gradient of saturation adjustments, which flows outward radially from the gray region of a*b* space. From this we can observe two problems. First, since the gradient of adjustments is orthogonal to hue, it can only adjust existing color contrast; it is impossible, for example, to add blueish shadows to a yellow flower. Second, the gradient diverges strongly around low-saturation values, which leaves adjustments poorly conditioned in this region of chromaticity space. This has the effect of amplifying chromatic noise in colors with low saturation (zoom in to the right of Figure 7 for an example) and leaves saturation adjustments undefined for grayscale images.

4.4.2 Channel Curves. The second most popular strategy we encountered involved using channel curves, which offer independent adjustment curves for each channel of an image in some chosen color space (typically RGB or $L^*a^*b^*$). Channel curves are capable of representing a much larger space of adjustments than a simple saturation slider, but for a chosen color space, the output values of each channel are limited to the functions of that channel. This limits different channel curves in different ways. Because $L^*a^*b^*$ space factors lightness and chromaticity into different channels, $L^*a^*b^*$ curves offer no way to map lightness variations to changes in chromaticity.

Figure 8 shows the optimal fit of $L^*a^*b^*$ curve adjustments to the example result from our interface designed to mimic the Cezanne painting in Figure 4. For RGB curves, changes to chromaticity are coupled with changes to lightness, with any single user adjustment guaranteed to impact both at the same time. Channel curves in general are also incapable of adapting adjustments to detail at different scales in an image. Figure 9 demonstrates this by comparing an image adjusted with our CGM tool to the closest adjustment one can achieve using RGB curves.

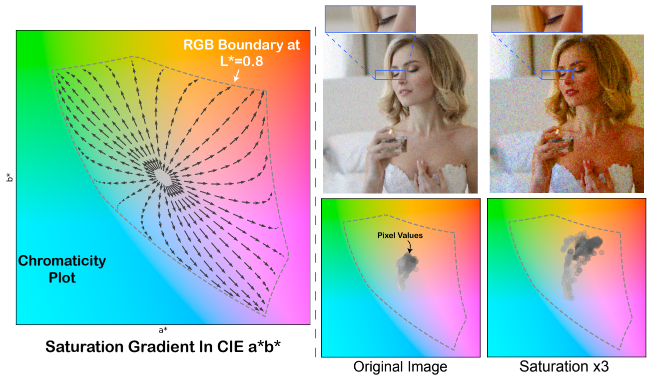

Figure 7: Problems With Saturation Adjustment: As increasing saturation maps chromaticities to other chromaticities, we can visualize the gradient of saturation in the perceptual chromaticity space defined by CIE a*b* (left). Increasing saturation approximately scales colors radially away from central gray. The direction of this scaling is poorly-conditioned at low-saturation pixels, and similar chromaticities map to similar chromaticities, which, making the operation extremely susceptible to noise (zoom in to see the difference). On the other hand, more saturated pixels with similar hue will scale in the same direction, with no guarantee that overall color contrast will increase. In fact, converging gradients in the plot on the left suggest that, in some cases, increasing saturation can actually reduce the perceived contrast between colors, even before reaching a saturation limit. Image credit: signatureedits.com.

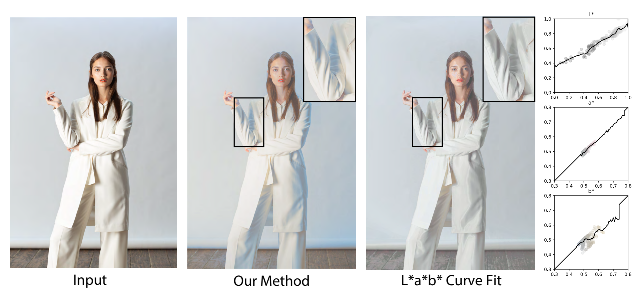

Figure 8: Fitting $L^*a^*b^*$ curves to adjustments made with the tone mapping operator of our tool: We fit $L^*a^*b^*$ curves to approximate edits made with our technique. The example shows that in the a^* and b^* channels, our method creates a mapping from one input value to multiple possible output values, which can not be modeled by a curve. The $L^*a^*b^*$ curves cannot produce the warm-cool contrast enhancement highlighted in the outputs. Image credit: signatureedits.com.

4.4.3 Selective Color Adjustments. Some of the more advanced Photoshop and Lightroom users we spoke to mentioned using a Selective Color Adjustment tool in Photoshop. The precise implementation of this tool is not public, but it offers selective control over color in specific regions of an image based on a selected initial

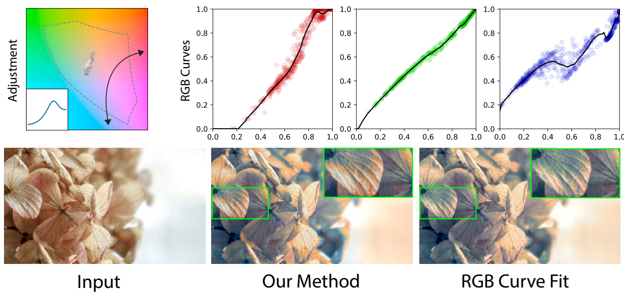

Figure 9: Channel Curves Comparison: We can fit a channel curve to approximate edits made with our technique as closely as possible by giving it a unique control point for every observed pixel value. This shows us two things; first, how unusual the shape of a channel curve needs to be to achieve a similar effect, and second, the limit of how close the channel curve can get. Here we see an example where, even with very detailed per-channel curves, it is impossible to fit an adjustment used to enhance structure contrast with our technique. Note that the color and local details of the highlighted regions do not match. Image credit: signatureedits.com

color trait. One of the more impressive demonstrations we found of an artist using existing tools for color contrast adjustment comes from a video posted by a professional photographer tasked with adding warm-cool contrast to an image (link: [33, 34]). Her strategy combines channel curves with the selective color adjustment tool and blending of multiple layers. The description she offers is detailed, running 8 minutes and 12 seconds long, which does not include the time she spent exploring to find her solution. Tellingly, the video concludes with her stating "*I hope I fulfilled this request properly. It was actually quite hard to do. It's a little bit harder than it looks to figure that out.*"

4.4.4 The Color Grading Tool. Advanced photographers and colorists in the film industry often uses the Color Grading tool for color correction and stylistic color manipulation. It is available in popular software such as Adobe Premiere, DaVinci Resolve, and Lightroom. The precise implementation of this tool is not publicly known, but it appears to selectively adjust the hue and saturation of three different tonal ranges in an image: "shadows", "midtones", and "highlights". Users often use it to create contrast in hue between tonal regions, enhancing perceived contrast in lightness and depths. For example, the cinematic "Teal and Orange" look is typically achieved by pushing the shadows to the blue and the highlights to the orange. The color grading tool only allows coarse and discrete chromaticity gradient mapping, and the control of contrast is implicit, relying on users' selection of hue and saturation.

4.4.5 The Gradient Map Tool. Another tool used primarily by illustrators and digital artists is the Gradient Map tool (e.g., in Adobe Photoshop), which lets users map lightness values directly to a chosen gradient of colors. This tool is typically used by digital artists to color grayscale artworks, as it ignores existing chromaticity values in the image. However, it can also be used for color grading by blending mapped copies of an image over top of the original image.

Notably, none of the tools mentioned in this section appear to offer selective control of color contrast for details at different spatial scales. To achieve frequency separation, sophisticated image editors sometimes use layers and blending tools in Adobe Photoshop. The low-frequency image ("color and tone") is obtained using Gaussian blur, while the high-frequency counter-part ("texture") is obtained by a subtractive blending of the input image with the Gaussian-blurred version. This technique is commonly used for portrait retouching, allowing users to adjust colors without affecting image details. While combining this method with other tools mentioned above can achieve color contrast editing at specific detail levels, the process is complex and involves significant manual effort. Our method enables chromaticity mapping operations in a way that makes targeting different spatial scales simple and efficient to explore.

5 CHROMATICITY GRADIENT MAPPING

At a high level, our approach factors control over chromaticity mapping into the specification of three things, as illustrated in Figure 10:

- A *gradient curve* ℓ, which controls a gradient field ∇_ℓ defined over chromaticity. Users control the gradient curve by adjusting the control points of a cubic Bézier spline in a^*b^* space.
- A *chromaticity adjustment curve* f_c, which maps filtered lightness values to chromaticity displacements along the specified gradient curve. Users control this curve the same way they would control tone adjustment curves, except the identity is a constant (horizontal) curve instead of a linear (diagonal) one, as outputs correspond to displacements instead of absolute values.
- An *adjustment map*, $\mathbf{M}$, which stores a map of selectively filtered lightness values that are provided as input to the adjustment curve.

We designed the gradient ∇_ℓ based on the notion of correlated color temperature, which measures a color's value relative to the Planckian Locus, a curve that describes the spectra emitted by blackbody radiation at different temperatures [9]. If the gradient curve ℓ is set to the Planckian locus, then our method approximately maps contrast adjustments to shifts of correlated color temperature. Our method generalizes this idea by letting users select their own alternative to the Planckian locus, and we can think of ∇_ℓ as a generalization of color maps used in visualization to 2D color gradients.

5.1 Chromaticity Adjustment Curves & Gradient Curves

Where a traditional adjustment curve f_L maps the lightness $L\mathbf{p}$ of each pixel $\mathbf{p}$ to some output lightness $f_L(L\mathbf{p})$, our chromaticity adjustment curve f_c maps $\mathbf{M}(\mathbf{p})$, the value of $\mathbf{p}$ stored in our adjustment map, to a change in chromaticity Δ_ℓ along the gradient ∇_ℓ. Here, we define ∇_ℓ to run parallel to ℓ (see Figure 10), so adjustments to f_c move the chromaticity of each pixel along a curve $\ell_\mathbf{p}$

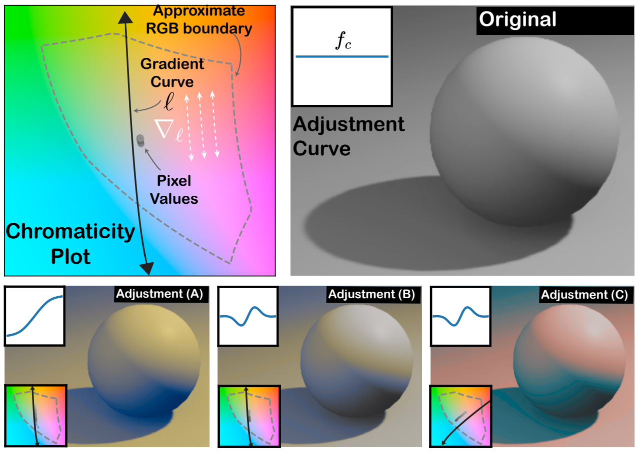

Figure 10: Contrast Enhancement with Lightness vs Color: We compute chromaticity adjustment by a user-controlled adjustment curve f_c, which maps lightness value Lp to amount of change ∇_ℓ along the gradient curve ℓ. (A) and (B) show the same gradient curve with different adjustment curves, while examples (B) and (C) have the same adjustment curve but use different gradients.

that passes through its chromaticity at a constant distance from ℓ. More formally, $\ell_\mathbf{p}$ is the level set of signed distance from ℓ defined by:

$$\ell_\mathbf{p}(t) = \ell(t) + ||\mathbf{p}_{ab} - \ell(t)||\hat{\ell_n}(t) \quad (1)$$

where $t \in [0, 1]$ is the progress along ℓ and $\hat{\ell_n}$ is the unit normal vector at $\ell(t)$ on the same side of ℓ as $\mathbf{p}_{ab}$, the coordinates of $\mathbf{p}$ in chromaticity space. Figure 10 shows 3 example adjustments applied to a simple scene of a shaded gray sphere where the adjustment map $\mathbf{M}$ in each case has been set the lightness values of the input image. Adjustments (A) and (B) show the same gradient curve with different adjustment curves, while examples (B) and (C) have the same adjustment curve but use different gradients. Here we observe that, unlike saturation, we can easily create color contrast in desaturated and grayscale images without boosting noise. This is because the direction of adjustments here is determined by our user gradient, which has minimal divergence for gradient curves with limited curvature.

5.2 Selective Detail Adjustments

Our approach to controlling the detail of contrast adjustments works by setting the values of $\mathbf{M}$ to a filtered version of our input image's lightness channel. Our filtering strategy builds on the local Laplacian filter [42], which works by blending the outputs of different adjustment curves tuned to each local neighborhood of an image. For completeness, we describe the original method and how we modify it here, noting that these details are primarily important for reproducibility and understanding the meaning of filter parameters that differ slightly from previous work. Note that we will also release an implementation of our version upon publication of the paper.

5.2.1 Local Laplacian Filtering. The original formulation of the local laplacian used a family of adjustment curves defined piecewise by an edge parameter β to control the contrast of edges, a detail parameter α to control the contrast of detail, and a threshold parameter σ to control the spatial scale that distinguishes between edges and detail (notation here differs from the original paper to better parallel our own implementation). The subsequent work of [7] introduced an approximation inspired by the bilateral grid [17] that tunes adjustments to intensity ranges instead of locations to yield a significant speedup. They also considered a more general class of local adjustments $r_g : i \mapsto r_g(i)$ that remap intensities i based on their value relative to some local average g, such that $r_g(i) = i - (i - g)f(i - g)$, where f is a continuous function. This includes curves that are smooth and map the local average to itself. Our filter falls into this family. We derive it by first defining parameters that play analogous roles to β, α and σ. In our case, the final adjustment is the sum of a linear component and a Gaussian one. The linear component, which we weigh with an edge parameter β, treats all spatial frequencies the same, effectively controlling broad-spectrum edges. We weigh the Gaussian component with parameter α and control the spatial bandwidth of details it affects by adjusting the standard deviation parameter σ. To simplify our equation, we write it in terms of $x : i - g$:

$$r_g(i) = g + x\left(\beta + \alpha e^{\frac{-x^2}{2\sigma^2}}\right) \quad (2)$$

This remapping function has the advantage of being differentiable everywhere, unlike the original piecewise formulation which has a discontinuity at the boundary between edge and detail. This is particularly relevant for our application, as chromaticity gradient adjustments can be thought of as something added to the original signal, so filtering edges or detail entirely from $\mathbf{M}$ is more likely than in typical tone mapping applications. We note that upon implementing the video version of our method we found a very similar remapping function used in an example application provided with the Halide source code repo [2]; they provide no explanation for the changes but we assume they are based on similar observations.

Equation 2 lets us balance the relative weight of edges and a controllable bandwidth of detail in $\mathbf{M}$, but we found the parameters β and α difficult to control in practice. To understand this, consider writing the matrix of color adjustments $\Delta\ell_\mathbf{p}$ at every pixel as a matrix-vector product $\mathbf{W}\mathbf{f_c}$ of f_c with some matrix $\mathbf{W}$. Here, every column of $\mathbf{W}$ represents a binary mask of pixels with shared values in $\mathbf{M}$. Note that $\mathbf{W}$ is low-rank, which is part of what makes adjustment curves simple to control. Scaling our adjustment map $\mathbf{M}$ in this case will simply rearrange the columns of $\mathbf{W}$, which does not change its column space. In other words, scaling $\mathbf{M}$ is equivalent to shifting the values of f_c. This redundancy is difficult to navigate, as absolute choices matter less than the relative values of β, α, and f_c. To address this, we instead let users control parameters A and θ, where

$$\alpha = A\theta \quad (3)$$
$$\beta = A(1 - max(\theta, 0)) \quad (4)$$

Here, A can be thought of as scaling the histogram on $\mathbf{M}$, while θ selects for global structure or local detail. Restricting A to positive values ensures that the derivative of our remapping function will

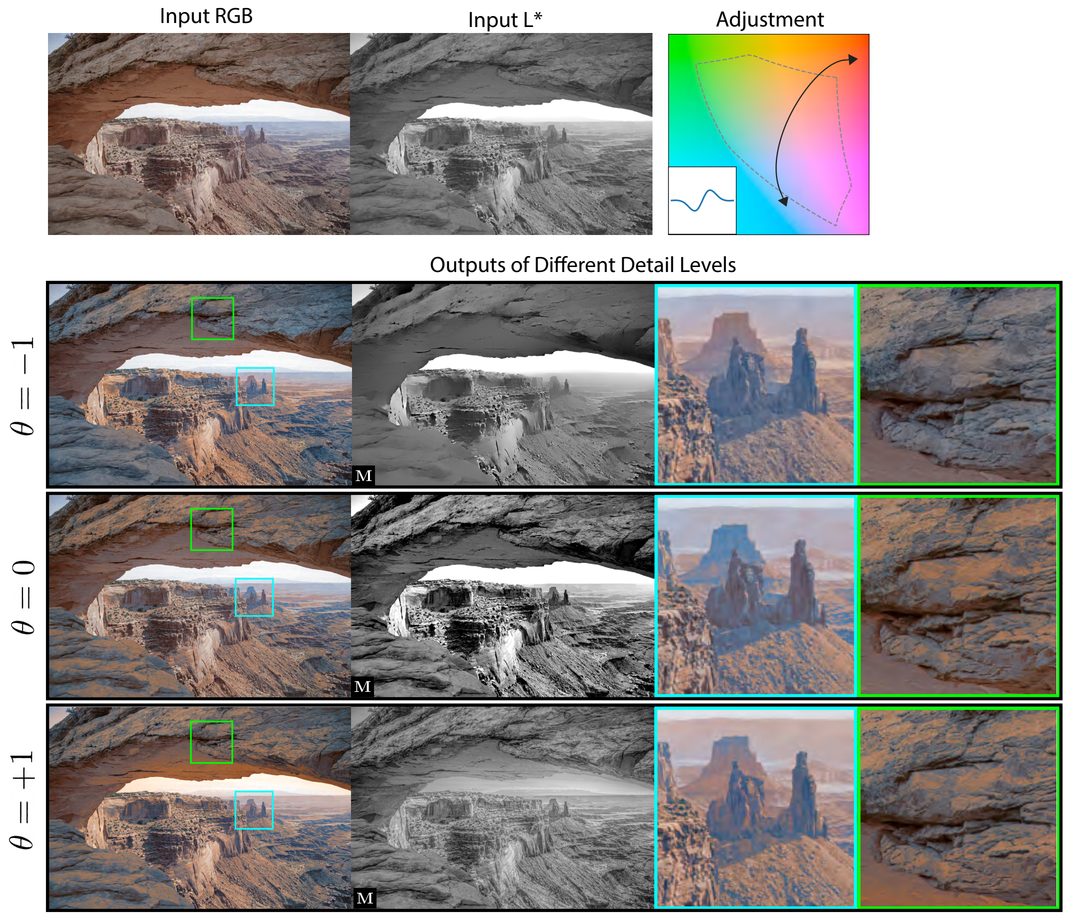

Figure 11: Detail Manipulation in Chromaticity: By filtering the adjustment map M, we control which frequencies of the input image drive the chromaticity adjustment. The figure shows three adjustments with different θ values. Each adjustment's corresponding M is displayed in grayscale. The top adjustment ($\theta = -1$) uses the chosen color gradient to enhance edge contrast, the middle ($\theta = 0$) enhances a mix of edge and detail contrast, and the bottom ($\theta = +1$) enhances detail contrast. Here, enhancing structure adds color contrast to the shadows of large structures. As the enhancement shifts toward detail, the warm-cool contrast is used instead to enhance the texture of rocks. Our method allows users to explore these alternative adjustments by focusing contrast on details at different scales. Image credit: signatureedits.com.

remain non-negative to avoid gradient reversals. We let θ range from -1, which corresponds to filtering out detail, to 1, which corresponds to filtering out edge structure. This range would be extreme in normal applications of the local Laplacian, but in our case, the map $\mathbf{M}$ represents how our image will change, so we can think of $\theta = -1$ as selectively adjusting edge structure, and $\theta = +1$ as selectively adjusting fine detail. Most expressiveness comes from controlling θ and σ, which have the most impact on the column space of $\mathbf{W}$. However, this same property also means the distribution of values in $\mathbf{M}$ is very sensitive to these variables. To keep a given adjustment curve meaningful across different sets of detail parameters we center our curve around the median value of $\mathbf{M}$ by default (users can override this if they choose). Figure 11 compares three adjustments made with the same curve and gradient at three different values of θ.

5.3 Coupled HDR Tone Mapping

Our filtering strategy provides convenient symmetry with tone mapping algorithms, which we can use to automatically generate chromaticity adjustments that compensate for contrast lost during tone mapping. The simplest way to do this is to use the difference between an image before and after tone mapping as our adjustment map $\mathbf{M}$. If our tone mapping operator is also a local Laplacian, we can represent this difference in terms of our filter parameters, allowing for further interactive adjustment by the user. Our interactive application offers a check box that activates this functionally. When selected, the threshold parameter σ is coupled with whatever choice is used for tone mapping lightness values, and the edge parameters β are coupled as well such that their sum is constant. This means that as the user compresses the dynamic range of the image, lost

Original Tone Mapped Tone Mapped With Chroma Compensation

Figure 12: Coupled HDR Tone Mapping: We show the original image, the detail-enhancing tone-mapped image, and the tone-mapped image with compensational global contrast enhancement through chromaticity mapping.

lightness contrast will automatically increase the contrast added to chromaticity (see Figure 12).

Even more generally, since our adjustment is orthogonal to the tone compression operation used for HDR tone mapping, it offers a flexible way to interactively compensate for visual detail lost in that process.

5.4 Video

Our method is designed to ensure that discontinuities within an edit can only happen where lightness discontinuities already exist in an image. This already does most of the work in preventing artifacts when extending the method to video. The one additional constraint we add is that all control points used to specify the curves of an adjustment should evolve smoothly in the time dimension. To accomplish this we let users specify adjustment keyframes, which we interpolate smoothly in time. Example video results are best seen through the videos in our supplemental material.

6 RESULTS

We implemented three versions of our method: one in Python for detailed analysis, another as a multi-pass shader in WebGL (Figure 13) to show that it can run at interactive rates in a browser, and a third in C++ using Halide [44] to process video results more quickly.

So long as the adjustment curve and gradient are smooth, our adjustments vary fairly smoothly with changes to an input image, as does interpolating between smooth curves and gradients over time. As such, our method works on video without modification. Our separate implementation in Halide is purely meant to speed up processing in this case.

Figure 14 shows a gallery of adjustments created with our technique, and more comprehensive results can be found in our supplemental gallery and videos.

7 USER STUDIES

To validate the color contrast enhancement ability of chromaticity gradient mapping (CGM), we need to answer two questions:

(1) *Can adjustments in chromaticity (CGM) achieve similar perceptual contrast enhancement as lightness adjustments?*
(2) *Do artists find it easier to control color contrast via CGM compared to that of existing tools?*

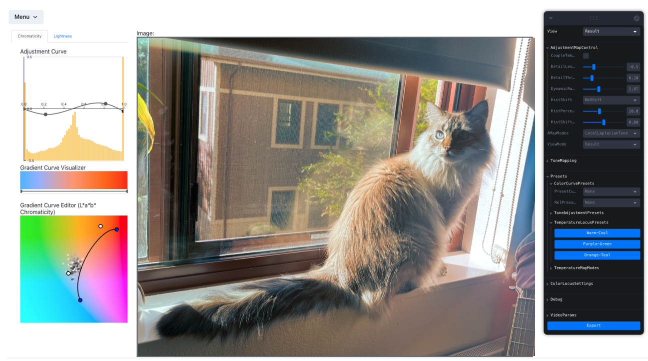

Figure 13: User Interface: At the upper left corner, we have two tabs for adjustment curves, one for chromaticity and one for lightness. In the lower-left corner, we have an editor for the gradient curve, which allows the user to define the chromaticity gradient using Bezier splines. The color of the gradient is visualized in the horizontal bar above. On the right, we have a control panel for Local Laplacian tone mapping parameters and presets of curve adjustments. The image is rendered in the center, and the user can toggle between the input image, adjustment map, and the output image.

Although our method is inspired by observations from the arts, we are not aware of any formal study on how chromaticity mapping affects the perception of contrast in images. To answer the first question, we conducted an online crowd-sourcing user study (Section 7.1) to address this question. For the second question, we recognize that the ideal subjects should possess advanced knowledge in both color science and existing image editing tools. However, constrained by resource limitations, we conducted a small-scale interview (Section 7.2) to evaluate how visual art practitioners compare our tool against other commercial software. We also conducted another experiment (Section 7.3) to directly compare CGM against $L^*a^*b^*$ curves (Section 4.4.2), the latest publicly available tool closely aligned with the user interface and functionality of CGM.

Figure 14: Result Gallery: Each pair of images shows an example of applying chromaticity gradient mapping with our web interface. The image on the left is the input, and the image on the right is the output. The two thumbnails at the lower left corner of the input image mark the gradient curve and chromaticity distribution (larger) and the adjustment curve (smaller). Please see our web gallery in the supplementary material for more visualizations of the adjustments. Images (a) to (e) credit: signatureedits.com.

7.1 Perceptual Study

To prove that perceptual contrast enhancement can be achieved by CGM, we designed a perceptual study that investigates the correlation between the scale of chromaticity adjustments and the perceptual contrast of images. If the chromaticity adjustment scale exhibits a positive correlation with users' perceived contrast, it would provide supporting evidence that increasing variations in chromaticity through CGM can indeed lead to enhanced perceptual contrast. It would validate our theoretical hypothesis that CGM can be effectively utilized to enhance the perceived contrast of images.

7.1.1 Study Design. 100 participants took part in our user study via *Connect*, an online crowd-sourcing platform with the compensation of 1 dollar for a five-minute survey. In the survey, we asked users to measure the amount of perceived contrast in an image edited by CGM in the scale of tone/lightness. For each task, users were shown a CGM or tonal edited image (reference) on the left side of the screen, and the original image (input) on the right (see examples in Figure 15). Users were then instructed to adjust the *tonal* contrast of the input image by controlling a slider placed under the input image. They were tasked with matching the contrast of the right image (tone adjusted) with the left one (chromaticity adjusted) as closely as possible, providing a continuous approximate measure of the perceived contrast in the chromaticity-adjusted image. We also asked participants to rate the "naturalness" of each image on a Likert scale with the explanation that "A natural image should not look artificially over-edited."

To ensure the quality of the study, we take additional steps to eliminate bias and supervise the quality of responses. The bias was reduced by randomizing the order and subset of images shown to each participant, randomizing the tonal contrast of the input images. The quality of user responses was tested by additionally

Figure 15: Perceptual Study Examples: We show 5 out of the 15 CGM samples that we used in the perceptual study. Each sample consists of an input image and a reference image edited with our method. We ask the participants to match the appearance of the reference image by adjusting the tonal contrast of the input image, and then rate the naturalness of both images. We report the mean user-selected contrast, mean naturalness (out of scale 1 to 5), and the change in chromaticity distribution along the first principal component. Images (a), (b), (c), (e) credit: signatureedits.com.

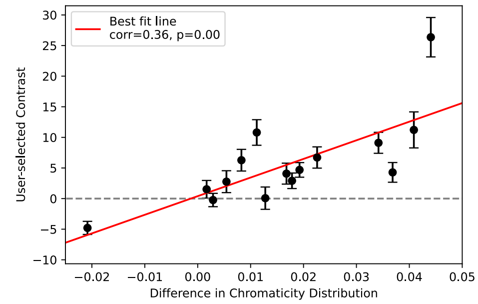

Figure 16: Correlation between Chromaticity Distribution and Contrast: Our visualization demonstrates a positive correlation between the perceived contrast enhancement and the standard deviation of chromaticity distribution along its first principal component. We use the standard error of the mean for the error bar.

tasking them to measure tonal contrast for reference images with known ground truth, those that have no adjustment or only tonal contrast adjustment. These quality control tasks provided a baseline for evaluating whether a participant can accurately gauge the tonal adjustments needed to match the appearance of two images exactly. To summarize, each participant was presented with 8 CGM samples, 2 tonal samples, and 2 unadjusted samples.

7.1.2 Results. To ensure the accuracy of our results, we excluded responses from participants who selected contrast values that deviated from the ground truth on more than 2 tonal or unadjusted samples by more than 10%. This left us with 89 validated responses.

We first evaluate our hypothesis that contrast adjustment with CGM will lead to a corresponding change in perceived tonal contrast. For each of the 15 CGM samples, we conducted Wilcoxon signed-rank test with the alternative hypothesis that the user will increase the tonal contrast of the input image in order to match the reference. We obtained statistically significant results (p-value < 0.05) for 11 out of the 15 CGM samples, with 10 showing that an increase in color contrast was perceived as increased tonal contrast, and 1 showing that decreased color contrast was perceived as decreased tonal contrast. To further investigate the outlier of decreased contrast, we computed the change in variance of chromaticities along their first principal component in a*b* space. By computing the Pearson correlation coefficient, we found that the change in chromaticity distribution positively correlates with user-perceived tonal contrast adjustment (p-value = 0.01), as shown in Figure 16.

We then study how the edits of our method affect the naturalness of the image. On the absolute scale of naturalness, we find that 11 out of the 15 samples are rated to be above the neutral level of naturalness, and 2 out of the 15 samples are rated to be unnatural. By computing the Pearson correlation coefficient, we find that both the absolute naturalness of reference images (p-value = 0.02) and the relative naturalness (p-value = 0.01) offset by the tonal-adjusted result negatively correlate with the user-selected contrast value. In other words, smaller contrast adjustments are perceived as similarly natural to corresponding tone adjustments, but as the amount of contrast added increases the naturalness of using color contrast appears to decrease.

7.1.3 Discussion. The statistical results show that by mapping lightness to changes in chromaticity, we can achieve the perceptual effect of adjusting the tonal contrast of an image without changing its lightness. We further investigated the 4 statistically non-significant color samples where our adjustment did not result in a conclusive perceived change in contrast. We found that these samples were different edits of two images (one shown in Figure 15 (d)). We suspect that this may be due to the fact that this scene has a high dynamic range with details in bright and dark regions of the image. Increasing color contrast without changing lightness can increase the visibility of these details. To achieve a similar increase in visibility with a standard tonal contrast adjustment one has to decrease overall tonal contrast. It is conceivable that some of the users may have been looking at the visibility of these details when making their decision, which could result in responding with lower contrast.

We observed a trade-off between perceived contrast and the naturalness of the image when applying our method. However, except for the reference image in Figure 15 (e), the participants do not find the edited images artificially over-edited. We argue that it is an encouraging result, and that creative exploration of even more extreme edits–like those of Monet and Cezanne–is sometimes desirable.

7.2 User Interaction and Feedback

To understand how artists approach our tool and potentially integrate it into their workflow, we conducted interviews with a small group of 4 participants over Zoom. The participants are advanced in visual arts (either through academic training or professional experience) and are recruited through personal contacts.

7.2.1 Experiment Setting. We first gave the participants a walk-through of our web interface. Then, we let them freely explore our tool and save a result that they felt satisfied with. Next, we asked them to try to duplicate their adjustment using the software they were most familiar with. We did not set a time limit, and the participants roughly spent 15 minutes individually on both of the tools. Finally, we collected their feedback through a survey with Likert-scale and open-ended questions. The images generated during the experiments as well as the survey responses are included in our supplementary material.

7.2.2 Free Exploration Results. The participants generally enjoyed playing with the tool, and they all selected neutral to positive responses on a Likert scale in response to whether they would use or recommend our tool for creative production. However, some users noted that human clients tend not to like dramatic color adjustments in photos with people, so the tool may be more useful for adjusting images with content that is less sensitive to color scheme change, or for expressing strong emotions through drastic edits. The interface itself poses a learning curve, and the participants suggested improvements by adding to our visualization of adjustment parameters, simplifying controls, and providing a set of meaningful presets as starting points.

7.2.3 Duplication Task Results. Three participants chose Adobe Photoshop and one chose PhotoScape. Participants spent most of the time adjusting color balance and channel curves, and some used hue/saturation adjustment, color grading, and color mixer. The most challenging part for users was to precisely apply a color gradient across different illuminance levels. The participants could usually adjust the midtones and highlights to their desired color scheme, but struggled more with adjustments to shadows. This makes sense, as chromaticity is very sensitive to small changes in individual channels at those lightness levels.

One of our most interesting observations was that, despite having explicitly explained to participants that our method does not change the lightness or tonal contrast of pixels, two of the participants made direct adjustments to the brightness/contrast controls and the tonal curve during their attempt at reproduction. This provides some support for our motivating idea of mapping tonal adjustments to chromaticity gradients as a way to expand the bandwidth available for contrast and detail manipulation. All participants responded positively that our method enables functionality that is hard to achieve with existing tools.

7.2.4 Discussion. Our analysis strongly points to a lack of support in existing tools for chromaticity mapping. In theory, chromaticity mapping without selective detail support can be done with very careful manipulation of channel adjustment curves, but these manipulations must be done in a precise balance that is difficult to navigate with three separate curves. By contrast, our tool makes it much easier to, for example, create gradients of cold color in the shadows, which is often desired but less intuitive to achieve with other tools. Moreover, the observation that users were inclined to increase tonal contrast in their attempts to visually match the output of our method supports the motivation of our work.

7.3 Comparison Study

We conducted a comparative experiment to evaluate the effectiveness of CGM in manipulating contrast at specific details against $L^*a^*b^*$ curves. In this study, we asked participants to complete image editing tasks by alternating between a simplified version of CGM and the $L^*a^*b^*$ curves.

We recruited 16 participants through personal contacts and Slack posts in the academic department. The participants have an equal gender split, ages 22 to 33. None of them reported color vision deficiency. Among these participants, two individuals were designated pilot users, helping us improve the robustness of the study's implementation.

7.3.1 Study Design. The study was performed asynchronously through a website. It started by presenting participants with a tutorial about the $L^*a^*b^*$ color space and the concepts of global and local contrast enhancement. The participants can experiment with both the $L^*a^*b^*$ curves and CGM with the interactive widgets. After familiarizing themselves with both tools, the participants were prompted with several image editing tasks, which asked them to enhance specific details in images with both tools. Following each task, participants are presented with a brief survey of five-point Likert-scale questions, each focusing on user satisfaction and ease of use for the respective tool, as well as user preference between the two tools. A post-study survey is presented at the end to evaluate participants' overall impression of the tools and gain qualitative feedback.

7.3.2 Results. We first aggregated the survey responses for each task. We performed Wilcoxon Signed-Ranks test to compare the distributions of Likert-scale questions on individual tools, as shown in Figure 17. Notably, the result shows that the participants find our method more intuitive to use ($p = 0.01$) and easier to enhance specific details in images ($p = 0.04$). We also evaluated the response on Likert-scale questions directly asking for a preference between CGM and $L^*a^*b^*$ curves with the Wilcoxon signed-rank test. As shown in Figure 18, the users report stronger preference on CGM for the edit result ($p = 0.03$), exploring creative decisions ($p = 0.01$), and effectiveness on detail enhancement ($p = 0.02$). The post-study survey does not show statistically significant results on users'

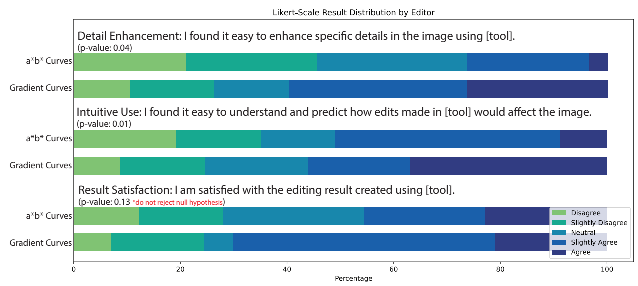

Figure 17: Evaluation Score Distributions For Each Method We asked the users to rate CGM and a*b* Curves individually on their effectiveness in detail enhancement, intuitiveness of use, and result satisfaction on a five-point Likert-scale. We ran Wilcoxon Signed-Ranks test to compare the distribution, and found that CGM is rated higher on effectiveness in detail enhancement and intuitiveness of use. Survey prompts and p-values reported on the right of the figure.

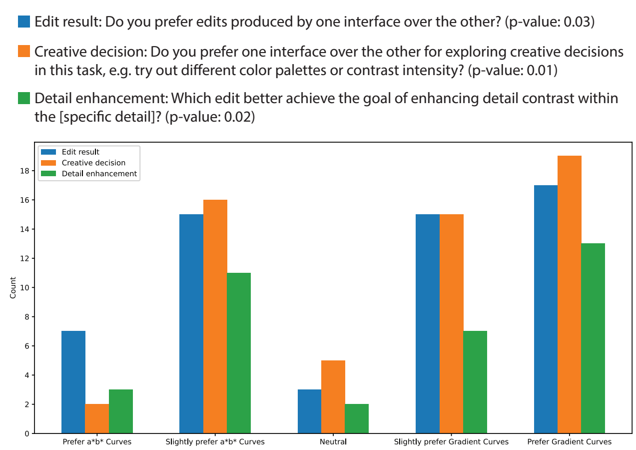

Figure 18: CGM vs $L^*a^*b^*$ Preference Scores We evaluate users' responses to the Likert-scale questions for preference between CGM and $L^*a^*b^*$ using Wilcoxon rank sum test. Results show that users overall prefer CGM. Survey prompts and p-values reported on the right of the figure.

ratings of satisfaction, ease of use, and control precision for both tools.

7.3.3 Discussion. The experiment result confirmed our theoretical analysis in Section 4.4.2 that CGM is a better method for contrast enhancement in chromaticity channels. When completing individual image editing tasks, users tend to prefer CGM for enhancing details and exploring creative decisions in the choice of color gradients. However, to our surprise, users did not show a significant preference for either tool in their post-study evaluation. We hypothesize that average users are not familiar with the $L^*a^*b^*$ color space, which makes both tools difficult to grasp in a short online study.

Please see more information regarding our user studies in the supplemental material, including illustrations of the study websites, survey forms, and more analysis of user responses. We note that the results of our user interaction study and comparison study are prone to participant response bias [19], even though we employed randomization methods to normalize users' expectations of the tested methods.

8 DISCUSSION

8.1 Limitations and Future Work

While the current design of our tool is promising, initial feedback from users points to areas for improvement and future work. For example, one somewhat counter-intuitive property of chromaticity gradient curves is that adjustments made with a particular gradient curve will not necessarily result in colors that lie on that curve. This is an intentional result of designing our adjustments to modify images in a smooth and continuous way. However, it does suggest that there may be a better way to convey the meaning of a given chromaticity curve to users, and feedback from some initial user interviews does echo this suggestion.

Our method occasionally produces a color banding effect when an adjustment curve with multiple inflection points is applied within a lightness range covered by a single image gradient. This issue is uncommon but can occur when the user optimizes for foreground content in scenes with high-contrast background edges blurred by shallow depth of field. For example, see the blue contour of the leaf in Figure 14 (e). It can be addressed by using masks to separate foreground and background adjustments, but is not implemented in our system.

8.2 Conclusion

Our work offers an algorithmic way to map tonal and detail adjustments to changes in chromaticity, as well as an interface for controlling this mapping.

Separating our representation of adjustments into one curve that controls the gradient in a^*b^* space and another that controls the magnitude of adjustments helps to interactions relatively simple despite the increased dimensionality of chromaticity relative to lightness. We believe that photo editing software could gain a lot from incorporating this type of control for chromaticity mapping, and that the algorithms underlying this tool may also be useful in other applications, such as visualizing information that varies across detail scales in images.

Finally, the original inspiration for this work grew from observing a surprising connection between the technique of impressionist painters and the design of color maps for information visualization. Understanding this connection and how to derive from it the design of our interactive tool required bringing together ideas from many different disciplines. One of our most surprising findings was how each of these disciplines independently deals with similar concepts quite differently—an observation reflected in the often contradictory vocabulary used across different fields. However, we believe that the perspective offered by each discipline carries its own unique and useful insights, and we hope in addition to the tool we have created, that this paper will help connect those insights for others and lead to similar efforts in the future.

REFERENCES

[1] 1948. Minutes of the Thirty-First Meeting of the Board of Directors of the Optical Society of America, Incorporated. *J. Opt. Soc. Am.* 38, 7 (Jul 1948), 651–651. https://doi.org/10.1364/JOSA.38.000651

[2] Andrew Adams. 2022. "Halide/apps/local_laplacian at Main Halide/Halide". "https://github.com/halide/Halide/tree/main/apps/local_laplacian"

[3] Andrew Adams, Jongmin Baek, and Myers Abraham Davis. 2010. Fast High-Dimensional Filtering Using the Permutohedral Lattice. *Computer Graphics Forum* 29 (2010).

[4] Adobe. 2021. Adobe Lightroom.

[5] Adobe. 2021. Adobe Photoshop.

[6] Ahmet Oğuz Akyüz, Kerem Hadimli, Merve Aydinlilar, and Christian Bloch. 2013. Style-Based Tone Mapping for HDR Images. In *SIGGRAPH Asia 2013 Technical Briefs* (Hong Kong, Hong Kong) *(SA '13)*. Association for Computing Machinery, New York, NY, USA, Article 23, 4 pages. https://doi.org/10.1145/2542355.2542384

[7] Mathieu Aubry, Sylvain Paris, Samuel W. Hasinoff, Jan Kautz, and Frédo Durand. 2014. Fast Local Laplacian Filters: Theory and Applications. *ACM Trans. Graph.* 33, 5, Article 167 (sep 2014), 14 pages. https://doi.org/10.1145/2629645

[8] Soonmin Bae, Sylvain Paris, and Frédo Durand. 2006. Two-Scale Tone Management for Photographic Look. *ACM Trans. Graph.* 25, 3 (jul 2006), 637–645. https://doi.org/10.1145/1141911.1141935

[9] L. E. Barbrow. 1964. International Lighting Vocabulary. *Journal of the SMPTE* 73, 4 (1964), 331–332. https://doi.org/10.5594/J07262

[10] M. Baxandall. 1995. *Shadows and Enlightenment.* Yale University Press. https://books.google.com/books?id=_gDAQgAACAAJ

[11] Nicolas Bonneel, Kalyan Sunkavalli, Sylvain Paris, and Hanspeter Pfister. 2013. Example-Based Video Color Grading.

[12] Peter J. Burt and Edward H. Adelson. 1983. A Multiresolution Spline with Application to Image Mosaics. *ACM Trans. Graph.* 2, 4 (oct 1983), 217–236. https://doi.org/10.1145/245.247

[13] Peter J. Burt and Edward H. Adelson. 1987. *The Laplacian Pyramid as a Compact Image Code.* Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 671–679.

[14] T. C. Callaghan. 1984. Dimensional interaction of hue and brightness in preattentive field segregation. *Percept Psychophys* 36, 1 (Jul 1984), 25–34.

[15] Cheng-Kang Ted Chao, Jason Klein, Jianchao Tan, Jose Echevarria, and Yotam Gingold. 2023. ColorfulCurves: Palette-Aware Lightness Control and Color Editing via Sparse Optimization. *ACM Transactions on Graphics (TOG)* 42, 4 (Aug. 2023).

[16] Cheng-Kang Ted Chao, Jason Klein, Jianchao Tan, Jose Echevarria, and Yotam Gingold. 2023. LoCoPalettes: Local Control for Palette-based Image Editing. *Computer Graphics Forum (CGF)* (June 2023). Special issue for Eurographics Symposium on Rendering (EGSR).

[17] Jiawen Chen, Sylvain Paris, and Frédo Durand. 2007. Real-Time Edge-Aware Image Processing with the Bilateral Grid. *ACM Trans. Graph.* 26, 3 (jul 2007), 103–es. https://doi.org/10.1145/1276377.1276506

[18] CIE Colorimetry. 2004. Report No: CIE Pub No 15. *Vienna: CIE Central Bureau* (2004).

[19] Nicola Dell, Vidya Vaidyanathan, Indrani Medhi, Edward Cutrell, and William Thies. 2012. "Yours is better!": participant response bias in HCI. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Austin, Texas, USA) *(CHI '12)*. Association for Computing Machinery, New York, NY, USA, 1321–1330. https://doi.org/10.1145/2207676.2208589

[20] Frédo Durand and Julie Dorsey. 2000. Interactive Tone Mapping. In *Proceedings of the Eurographics Workshop on Rendering Techniques 2000.* Springer-Verlag, Berlin, Heidelberg, 219–230.

[21] Frédo Durand and Julie Dorsey. 2002. Fast Bilateral Filtering for the Display of High-Dynamic-Range Images. *ACM Trans. Graph.* 21, 3 (jul 2002), 257–266. https://doi.org/10.1145/566654.566574

[22] Zeev Farbman, Raanan Fattal, Dani Lischinski, and Richard Szeliski. 2008. Edge-Preserving Decompositions for Multi-Scale Tone and Detail Manipulation. In *ACM SIGGRAPH 2008 Papers* (Los Angeles, California) *(SIGGRAPH '08)*. Association for Computing Machinery, New York, NY, USA, Article 67, 10 pages. https://doi.org/10.1145/1399504.1360666

[23] Raanan Fattal, Dani Lischinski, and Michael Werman. 2002. Gradient Domain High Dynamic Range Compression. *ACM Trans. Graph.* 21, 3 (jul 2002), 249–256. https://doi.org/10.1145/566654.566573

[24] Charles Harrison Francis Frascina. 1983. *Modern Art And Modernism: A Critical Anthology (Icon Editions).* Westview Press. https://libgen.fun/book/index.php?md5=dccfcabc635604b7cdeb7a87861d5569

[25] Eduardo S. L. Gastal and Manuel M. Oliveira. 2012. Adaptive Manifolds for Real-Time High-Dimensional Filtering. *ACM Trans. Graph.* 31, 4, Article 33 (jul 2012), 13 pages. https://doi.org/10.1145/2185520.2185529

[26] Comte de Buffon" "Georges-Louis Leclerc. 1743. "Dissertation sur les couleurs accidentelles". *"Memoires de l'Academie Royale des Sciences"* (1743).

[27] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. 2014. Generative adversarial nets. *Advances in neural information processing systems* 27 (2014), 2672–2680.

[28] Mingming He, Dongdong Chen, Jing Liao, Pedro V Sander, and Lu Yuan. 2018. Deep exemplar-based colorization. *ACM Transactions on Graphics (TOG)* 37, 4 (2018), 47.

[29] C.G. Healey. 1996. Choosing effective colours for data visualization. In *Proceedings of Seventh Annual IEEE Visualization '96.* 263–270. https://doi.org/10.1109/VISUAL.1996.568118

[30] Ewald Hering. 1920. *Grundzüge der Lehre vom Lichtsinn.* J. Springer.

[31] International Electrotechnical Commission. 1999. IEC 61966-2-1:1999 - Multimedia Systems and Equipment - Colour Measurement and Management - Part 2-1: Colour Management - Default RGB Colour Space - sRGB. , 51 pages.

[32] George H. Joblove and Donald Greenberg. 1978. Color Spaces for Computer Graphics. In *Proceedings of the 5th Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH '78).* Association for Computing Machinery, New York, NY, USA, 20–25. https://doi.org/10.1145/800248.807362

[33] Kayleigh June. 2021. Kayleigh June Photography. https://www.kayleighjune.com/

[34] Kayleigh Kayleigh June. 2016. How to create a warm/cool contrast colour grade [photoshop tutorial]. https://www.youtube.com/watch?v=GoaMJobmLg8

[35] Min H. Kim and Jan Kautz. 2008. Consistent Tone Reproduction. In *Proceedings of the Tenth IASTED International Conference on Computer Graphics and Imaging* (Innsbruck, Austria) *(CGIM '08).* ACTA Press, USA, 152–159.

[36] Adam G. Kirk and James F. O'Brien. 2011. Perceptually Based Tone Mapping for Low-Light Conditions. In *ACM SIGGRAPH 2011 Papers* (Vancouver, British Columbia, Canada) *(SIGGRAPH '11).* Association for Computing Machinery, New York, NY, USA, Article 42, 10 pages. https://doi.org/10.1145/1964921.1964937

[37] L. 2010. Filtering and condensation tasks with integral and separable dimensions. https://api.semanticscholar.org/CorpusID:195748206

[38] Dani Lischinski, Zeev Farbman, Matt Uyttendaele, and Richard Szeliski. 2006. Interactive Local Adjustment of Tonal Values. *ACM Trans. Graph.* 25, 3 (jul 2006), 646–653. https://doi.org/10.1145/1141911.1141936

[39] Yang Liu and Jeffrey Heer. 2018. Somewhere Over the Rainbow: An Empirical Assessment of Quantitative Colormaps. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems* (Montreal QC, Canada) *(CHI '18).* Association for Computing Machinery, New York, NY, USA, 1–12. https://doi.org/10.1145/3173574.3174172

[40] R. Mantiuk, R. Mantiuk, A. Tomaszewska, and W. Heidrich. 2009. Color correction for tone mapping. *Computer Graphics Forum* 28, 2 (2009), 193–202. https://doi.org/10.1111/j.1467-8659.2009.01358.x

[41] James Clerk Maxwell. 1857. The diagram of colors. *Trans. R. Soc. Edinburgh* 21 (1857), 275–298.

[42] Sylvain Paris, Samuel W. Hasinoff, and Jan Kautz. 2015. Local Laplacian Filters: Edge-Aware Image Processing with a Laplacian Pyramid. *Commun. ACM* 58, 3 (feb 2015), 81–91. https://doi.org/10.1145/2723694

[43] P. Perrin and M. Kisiel. 2018. *Colours of Impressionism: Masterpieces from the Musée d'Orsay.* National Gallery Singapore. https://books.google.com/books?id=8f_2DwAAQBAJ

[44] Jonathan Ragan-Kelley, Connelly Barnes, Andrew Adams, Sylvain Paris, Frédo Durand, and Saman Amarasinghe. 2013. Halide: A Language and Compiler for Optimizing Parallelism, Locality, and Recomputation in Image Processing Pipelines. In *Proceedings of the 34th ACM SIGPLAN Conference on Programming Language Design and Implementation* (Seattle, Washington, USA) *(PLDI '13).* ACM, New York, NY, USA, 519–530. https://doi.org/10.1145/2491956.2462176

[45] B.E. Rogowitz and A.D. Kalvin. 2001. The "Which Blair project": a quick visual method for evaluating perceptual color maps. In *Proceedings Visualization, 2001. VIS '01.* 183–556. https://doi.org/10.1109/VISUAL.2001.964510

[46] Maria Shugrina, Amlan Kar, Karan Singh, and Sanja Fidler. 2018. Color Sails: Discrete-Continuous Palettes for Deep Color Exploration. *CoRR* abs/1806.02918 (2018). arXiv:1806.02918 http://arxiv.org/abs/1806.02918

[47] Samuel Silva, Beatriz Sousa Santos, and Joaquim Madeira. 2011. Using color in visualization: A survey. *Computers & Graphics* 35, 2 (2011), 320–333. https://doi.org/10.1016/j.cag.2010.11.015 Virtual Reality in Brazil Visual Computing in Biology and Medicine Semantic 3D media and content Cultural Heritage.

[48] Stephen Smart and Danielle Albers Szafir. 2019. Measuring the Separability of Shape, Size, and Color in Scatterplots. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems* (Glasgow, Scotland Uk) *(CHI '19).* Association for Computing Machinery, New York, NY, USA, 1–14. https://doi.org/10.1145/3290605.3300899

[49] Alvy Ray Smith. 1978. Color Gamut Transform Pairs. In *Proceedings of the 5th Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH '78).* Association for Computing Machinery, New York, NY, USA, 12–19. https://doi.org/10.1145/800248.807361

[50] Kaleigh Smith, Grzegorz Krawczyk, Karol Myszkowski, and Hans-Peter Seidel. 2006. Beyond Tone Mapping: Enhanced Depiction of Tone Mapped HDR Images. *Computer Graphics Forum* (2006). https://doi.org/10.1111/j.1467-8659.2006.00962.x

[51] Paul Smith. 2021. *Painting, Science, and the Perception of Coloured Shadows: 'The Most Beautiful Blue'.* Routledge. http://gen.lib.rus.ec/book/index.php?md5=9044293E50A1F91CED504647AA195F48

[52] Jianchao Tan, Jose Echevarria, and Yotam Gingold. 2018. Efficient palette-based decomposition and recoloring of images via RGBXY-space geometry. *ACM Transactions on Graphics (TOG)* 37, 6, Article 262 (Dec. 2018), 10 pages. https://doi.org/10.1145/3272127.3275054

[53] Liang Zhou and Charles D. Hansen. 2016. A Survey of Colormaps in Visualization. *IEEE Transactions on Visualization and Computer Graphics* 22, 8 (2016), 2051–2069. https://doi.org/10.1109/TVCG.2015.2489649

[54] Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A Efros. 2017. Unpaired Image-to-Image Translation using Cycle-Consistent Adversarial Networks. In *Computer Vision (ICCV), 2017 IEEE International Conference on*.