Color Field: Developing Professional Vision by Visualizing the Effects of Color Filters

Matthew T. Beaudouin-Lafon
University of California, San Diego
La Jolla, California, USA
mbeaudouinlafon@ucsd.edu

Jane L. E
University of California, San Diego
La Jolla, California, USA
je@ucsd.edu

Haijun Xia
University of California, San Diego
La Jolla, California, USA
haijunxia@ucsd.edu

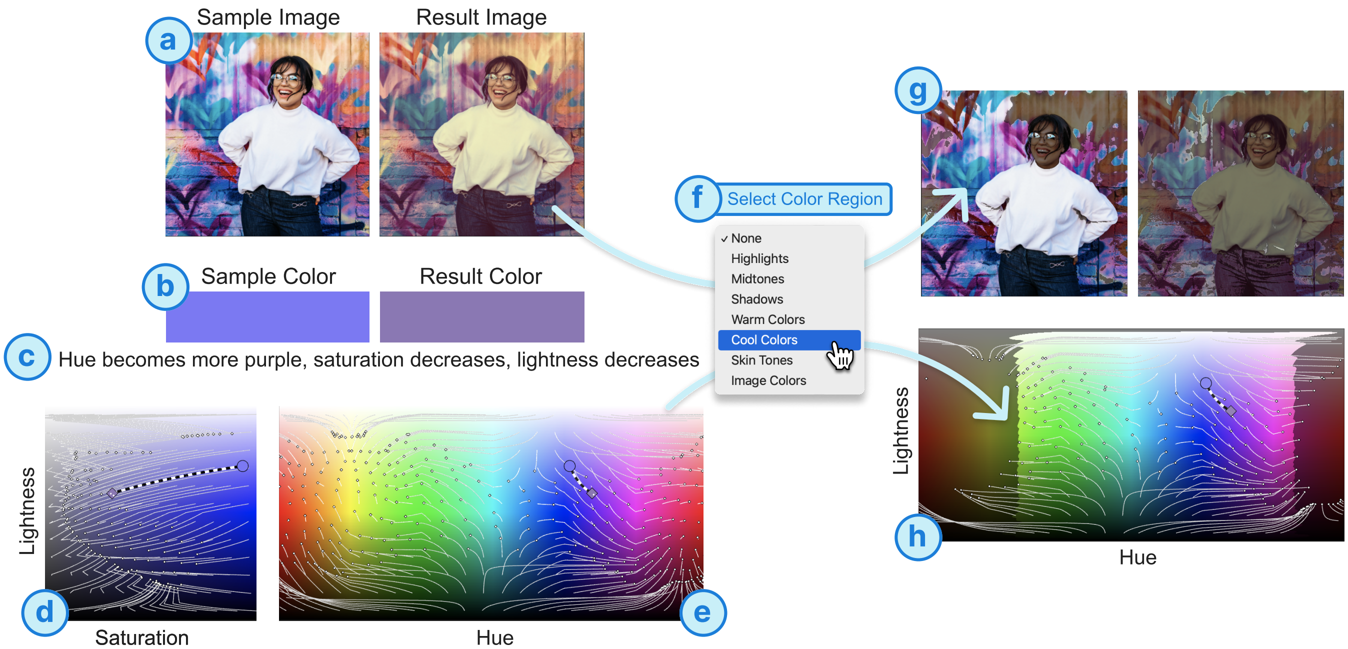

Figure 1: Color Field visualizes the effect of color filters with a *sample image* before and *result image* after applying the filter (a), a *sample and result color* before and after the filter (b), a *natural language description* of the filter's effect on the sample color in terms of hue, saturation and lightness (c), a *vector field* showing how the filter affects saturation and lightness at a fixed hue (d), and a second *vector field* showing how the filter affects hues and lightness at a fixed saturation (e). The sample color determines the fixed hue and saturation of views (d) and (e). When a user selects the "Cool Colors" region (f), the images (g) and color field (h) highlight cool colors and de-emphasize warm colors. Note that the color filter makes the image warmer (a), which is shown in the vector field (e) as gravitating away from blue and towards red, orange and yellow (h).

ABSTRACT

Color filters are ubiquitous across visual digital media due to their transformative effect. However, it can be difficult to understand how a color filter will affect an image, especially for novices. In order to become experts, we argue that novices need to develop Goodwin's notion of Professional Vision [29]. Then, they can "see" and interpret their work in terms of their domain knowledge like experts. Using the theory of Professional Vision, we present two design objectives for systems that aim to help users develop expertise. These goals were used to develop Color Field, an interactive visualization of color filters as a vector field over the Hue-Saturation-Lightness color space. We conducted an exploratory user study in which five color grading novices and four experts were asked to analyze color filters. We found that Color Field enabled multiple strategies to make sense of filters (e.g. reviewing the overall shape of the vector field) and discuss them (e.g. using spatial language). We conclude with other applications of Color Field and future work to leverages Professional Vision in HCI.

UIST '23, October 29–November 01, 2023, San Francisco, CA, USA

ACM ISBN 979-8-4007-0132-0/23/10.
https://doi.org/10.1145/3586183.3606828

CCS CONCEPTS

• **Human-centered computing** → *HCI theory, concepts and models*; *Information visualization*; **Graphical user interfaces**.

KEYWORDS

visualization, creativity support tool, professional vision, support novice, color filter, color grading

ACM Reference Format:
Matthew T. Beaudouin-Lafon, Jane L. E, and Haijun Xia. 2023. Color Field: Developing Professional Vision by Visualizing the Effects of Color Filters. In *The 36th Annual ACM Symposium on User Interface Software and Technology (UIST '23), October 29–November 01, 2023, San Francisco, CA, USA.* ACM, New York, NY, USA, 16 pages. https://doi.org/10.1145/3586183.3606828

1 INTRODUCTION

Color filters are ubiquitous in digital art. They help graphic designers develop color palettes, digital painters make their colors more cohesive, and film makers set the tone of their movies. Color filters take many forms, including blending modes, gradient maps and Look Up Tables (LUTs). Since color is central to visual media, color filters are a powerful tool to reshape the look and feel of a piece. An artist's approach to manipulating color can even become part of their artistic identity. For example, movies from Hollywood directors like Wes Anderson can be recognized at a glance just from their distinctive use of color.

Despite its importance, manipulating color is remarkably difficult. To edit colors, experts create color filters that alter an input image to match the effect they seek to achieve. This process, known as color grading, involves manipulating a multitude of parameters with confounding effects. Moving one slider might need to be compensated by adjusting another slider in a careful balance. However, the complexity of these tools is made warranted by the precision and control they afford.

On the other end of the spectrum, pre-made filters let users explore a range of color options while avoiding the complexity of authorship altogether. Although pre-made filters let casual photographers achieve compelling results with little effort, they lack control and expressivity. To cross this gap, novices motivated to learn how to create their own color filters will use tools that are progressively more complex. For photography, they might start using basic sliders in their phone's camera app, move to Adobe Lightroom [34] and eventually use Adobe Photoshop [33]. Each successive application introduces features that afford finer control, but also expect familiarity with the intricacies of color grading. For example, the *lift* control in Da Vinci Resolve [19] is only useful to users who can identify the shadows in an image and decide what they should look like. Using progressively more complex software helps scaffold learning more complex tool, but does little to scaffold learning the mechanics of the medium itself.

Goodwin calls this affinity for the medium "Professional Vision": experts use their domain knowledge to "see" differently than novices [29]. For example, a color grading expert may analyze a filtered image and claim that "the midtones are too saturated" [7, 58]. To arrive at this conclusion, they must mentally separate the image's highlights, midtones and shadows based on pixel brightness and think about colors in terms of hue, saturation, and lightness. These features help structure the design space of color grading and guide actions within that space.

Without Professional Vision, novices can be overwhelmed by the complexity of the unstructured medium. For example, they may notice that a photo looks hazy without knowing what contribute to that effect; whereas an expert would identify that the darkest points of the image are relatively bright. Since color filters are functions, they imply a change in color rather than a static result. Therefore, they are often indirectly represented as their effect on images, or their parameters in the editor. As a result, novices can only build their Professional Vision by looking at many examples over time and inferring the effect of parameters from repeated use. Instead, filters would be easier to interpret if they were represented explicitly. Such a representation could then scaffold the acquisition of Professional Vision.

We present Color Field, an interactive visualization system for color filters. Since a color filter is a function mapping colors to colors, it can be represented as a vector field on a color picker, where each vector represents how an input color is affected by the filter (Figure 1). To understand how to help users develop Professional Vision, we synthesized Goodwin's work into two design objectives (Section 3.2) and applied them to the domain of color grading.

In an interview study with three color grading experts, we validated that Color Field represented the right information and identified how the system might help novices develop Professional Vision. We also conducted an exploratory user study with five novices and four experts to understand how Color Field could influence their analysis of color filters and support Professional Vision. Novices described how Color Field helped them analyze color filters in multiple ways. In fact, even after the visualization of the field was withheld, some novices simulated the field visualization in their heads to support their interpretation of the filters. Experts appreciated that the tool visualized the inner mechanism of filters, in comparison to current tools which only show the result of a filter. These findings provide some initial support that Color Field has the potential not only to support the development of Professional Vision for color filters, but also to complement experts' existing Professional Vision. This paper contributes:

- **design objectives** synthesized from Goodwin's work [29] to design tools for developing Professional Vision,
- a **case study** in applying these objectives to the domain of color filters based on insights from expert interviews,
- an **interactive visualization**, Color Field, that presents a novel interactive representation of color filters,
- an **exploratory user study** with novices and experts to see how Color Field could affect their analysis of color filters.

Although this paper explores the application of our design objectives in a single domain, we believe that they may be applied to other domains and foster the development of creativity support tools that help users develop Professional Vision and better understand the medium they work with.

2 RELATED WORK

We first summarize work related to understanding and editing color filters, then explore different perspectives on gaining expertise, and finally review tools with external representations that support the development of Professional Vision.

2.1 Understanding and Editing Color Filters

Researchers have explored multiple ways to edit colors in visual media ranging from previewing edits directly in the camera [3], enabling edit commands through high-level natural language [43], and automatically customizing edits based on learned personal preferences [40]. These tools reduce tedious manual editing, letting experts quickly make rough edits before refining the results. However, they require clear artistic goals, which novices often lack. Hence they are better suited for experts who know what they want than novices who are still learning.

Another direction has focused on simplifying color grading tasks by letting users select pre-made color filters. Popularized by Instagram [35], many mobile operating systems now come with similar pre-made filters that users can apply to their photo. Meanwhile, computer graphics researchers have produced simple but more expressive tools that provide meaningful control over the color of an image. For instance, researchers have leveraged the abundance of image examples for color style transfer. These systems extract a color palette from an example and apply it to the target image [11, 15, 56, 71]. This lets users apply the overall look of an image that they like onto their work with minimal effort, but with more stylistic flexibility than picking arbitrary pre-made color filters. Chang et al. provide more granular control by extracting and adjusting a color palette from the target image, in turn affecting the image's colors [14].

Pre-made filters primarily benefit people who prioritize performance and efficiency. They let users quickly test different edits to arrive at an acceptable result with minimal effort. However, selecting curated filters is much less expressive than working with color grading controls. As a result, these systems might not be sufficiently flexible for novices interested in honing their artistic preferences. Instead of simplifying color grading controls, Color Field aims to help users understand the task and become experts.

Color filters are often described natural language and examples on sample images. For example, this is how Adobe Photoshop's documentation describes blending modes, a type of color filter [21]. Similarly, when photographers sell color filters they created, they are often marketed using an evocatice description and examples with representative photographs. These methods indirectly suggest what a filter might do to a user's image, but are limited by the abstract nature of language and the creator's selection of images. Alternatively, in order to characterize how blending modes behave, Valentine applies blend modes to color pickers to show how color responds [70]. This strategy can help users understand the range of colors a filter will produce and suggest its general behavior, but since its representation is static, it can be difficult to pick out how individual colors are transformed. The crux of the challenge with visualizing color filters is that they are functions, which imply an action rather than an object. And yet, to our knowledge, color filters have not been directly visualized as such. Hence, instead of relying on imprecise descriptions or showing the result of a filter, Color Field represents how colors map to each other using a vector field.

2.2 Developing Expertise

We use Professional Vision to guide the design of tools that help users become experts. In particular, Goodwin demonstrates that experts use their domain knowledge to *highlight* regions of their subject as part of their discursive practices (Section 3.1). We describe several other lenses that researchers across cognitive science, psychology, and education have proposed to describe the development of expertise. Chase and Simon demonstrated that experts could only memorize chess board layouts better than novices when the layout was representative of a real game of chess [16]. For random layouts, experts performed as well as novices. Chase and Simon argued that experts could no longer "chunk" the unplayable boards based on prior experience and therefore relied on the same brute-force memorization as novices. Deliberate practice [26], learning by doing [57], and mastery learning [8, 9] are all frameworks suggesting that expertise can be achieved with focused practice and time.

Researchers have also directly studied the impact of "seeing" like an expert. For example, to the untrained eye, it is nearly impossible to determine the sex of a day old chick. However, experts can be trained to determine the sex of over one thousand chicks in an hour with almost perfect accuracy [47]. Researchers have also found that subtly directing a novice's gaze to follow the scan-path of an expert radiologist can significantly improve the novice's ability to identify abnormalities [4, 60]. The concepts behind Professional Vision have been studied in a variety of different contexts within Cognitive Science, and have been used in HCI to understand the practices of user groups. Professional Vision is therefore a promising lens to study expert practice grounded in human cognition.

2.3 Tools to Support Learning

Prior work demonstrates methods to support learning with a variety of interactions and tools which roughly falls into two categories: supporting learning the software and learning domain knowledge. Approaches to support software learning have ranged from leveraging application context to provide assistance [28, 31, 36], improving software tutorials [18, 38], and progressively disclosing features in the tool to ease a novice into complex software [59]. Other interactions have been designed to help users develop expertise through frequent use and familiarity [5, 42]. While designing systems for Professional Vision may help its users develop expertise in the software, our main goal is to help users learn domain knowledge independently of the tools they use.

Since tools can be more readily available as an educational resource than limited expert and instructor time, researchers have extensively studied how computation can support domain learning, such as through a wide range of intelligent tutoring systems [30, 41, 52] or tools to scaffold learning [49]. Many such tools provide the support that an expert might provide while minimizing the need for expert intervention, for example using customized, context-aware feedback. While some of these tools help users learn from experts, many provide scaffolding for peers to provide more effective, expert-like feedback [13, 37, 48, 51, 74]. We are interested in helping novices develop Professional Vision on their own using the visual scaffolding provided by Color Field.

2.4 Using External Representations to Support Professional Vision

Despite being a common cognitive phenomenon, Professional Vision has rarely been used explicitly in HCI. However, some prior work implicitly supports Professional Vision by leveraging expert visualizations and annotations to help novices learn a new domain. For example, Head et al. analyzed numerous textbooks and educational resources to synthesize patterns in how instructors augment math formulas and figures for the purpose of instruction [32]. In the context of photography, experts train their eye to identify potential subject and improve their work. For example, photographers annotate images with shapes and grid lines to focus attention on different regions of the image and explain photographic concepts to help the reader see as they are seeing [25]. Graphics and HCI researchers have proposed a number of algorithms and tools to provide these visualizations on a collection of images or directly in the camera [2, 22–24, 46, 50, 73]. To support visual design, GUIComp visualizes saliency maps for users to understand the hierarchy of their current design, along side some algorithmically computed visual aesthetics scores [45]. These tools, like our work, use external representations to support the development of Professional Vision.

External representations have widely been shown to help cognitive processes [29, 39, 64, 65, 67]. First, they reduce cognitive load by externalizing information that would otherwise be kept in mind. More interestingly, their spatial nature can help highlight relationships and facilitate reasoning [64]. For example, architects use sketches not only to illustrates ideas in their minds, but also to surface functional and spatial relationships between parts of the building and help structure their ideas [65]. Additionally, visuospatial representations facilitate certain types of inference, or what Tufte calls "graphic arithmetic" [66]. Color Field's design demonstrates these benefits: users can make inferences about the color filter based on the overall shape of the field, and use these observations to structure the exploration of the visualization.

Closest to our work is Chevalier et al.'s Histomages, which improves the familiar histogram, a visualization common to most desktop photo editors [17]. Since novices struggle to map the shape of the histogram to the content of the image, Histomage lets users directly manipulate the image with the histogram. In other words, selecting pixels in the histogram highlights them in the image, and vice versa. This bidirectional mapping helps novices interpret a visualization designed for experts. However, histograms, and therefore Histomages, only represents the current state of the image and does not represent color filters. With Color Field, we introduce a new visualization that represents the change induced by color filters, rather than just the outcome.

3 PROFESSIONAL VISION

Goodwin defines Professional Vision as "socially organized ways of seeing and understanding events" that are relevant to experts [29]. Although Goodwin studies this phenomenon as discursive practices, we focus on the cognitive implications of his findings; that experts "see" more than novices. As people become experts in a domain, they learn its theory and learn to apply it in their daily work. Therefore, when they look at a subject in their domain of expertise, their domain knowledge enhances specific characteristics of the subject. For example, a novice only sees a photo for what it shows, but a photographer's eye is intuitively drawn to regions deemed important to consider, such as the image's shadows or its highly saturated components. This specialized perspective helps

Figure 2: Examples of highlighting. The camera *highlights* over-exposed areas of the image with an animated zebra pattern. Furthermore, this caption draws the reader's attention to the zebra pattern instead of the camera's buttons.
Image from Unsplash by Ben Griffiths.

them quickly analyze and make decisions, for instance by choosing to lower the image's saturation.

Goodwin studied Professional Vision through three interrelated practices, summarized in Section 3.1. In Section 3.2, we synthesize them into two design objectives for systems that strive to help novices become experts.

3.1 Expert Practice

As a sociologist, Goodwin studied experts by observing how they act and communicate. He investigated three practices in particular: the use of *coding schemes* to understand the domain, *highlighting* to apply that knowledge, and the use of *material representations* to act on this knowledge. These practices are evidence that experts craft how they see their domain of expertise.

3.1.1 Coding Schemes. A *coding scheme* is a system of codes that organizes the explicit knowledge that experts have learned. It helps them make sense of the domain and communicate with other experts. The coding scheme includes theory and vocabulary that simplifies the domain into its principle components. By applying this theory to a specific subject, all of its irrelevant details can be ignored. For example, an image can be analyzed in many ways, but an expert color grader can focus on its lightness using codes like *shadows* and *highlights*. This coding scheme helps experts structure their thoughts about the medium and frame the actions they will carry out. Therefore, the coding scheme is more than jargon; it structures the domain and helps them think about it more clearly. However, experts don't just use the coding scheme in the abstract, it is applied to subjects on a case by case basis to make sense of them. To do this, experts must "see" the coding scheme in subjects.

3.1.2 Highlighting. Experts guide attention by *highlighting* regions. This helps them understand and explain parts of the their work using the coding scheme. For example, a color grading expert may use their mouse cursor to attract a collaborator's attention to the darker regions of the image, referring them as "the shadows of the image", and discuss if it should be made brighter. Goodwin argues that this practice makes manageable subjects that are otherwise too perceptually complex and full of detail.

Highlighting comes in many forms: a caption on a figure draws the readers attention to the regions that the author considers important, and an animated zebra pattern overlaid on the camera tells the user which areas of the image are overexposed (Figure 2). In the latter case, highlighting is supported by a tool, which leads to the third practice.

3.1.3 Material Representations. Goodwin observed that experts create and use tools, referred as *material representations*, to highlight and interpret their subjects. This is consistent with the extensive research in Cognitive Science that suggests that external representations help people think [1, 39, 67, 68]. Hence, the representations that experts create provide insight into the way they think and the tasks they want to achieve. For example, photographers use *waveform monitors* to represent the brightness of pixels in the image and help them objectively calibrate the brightness in their image. This shows that these experts sought alternate objective representations of the image to support their work, in this case to balance the brightness of the image.

In summary, **coding schemes** represent *what* experts know, **highlighting** demonstrates *where* that knowledge is relevant, and **material representations** illustrate *how* it is used. Goodwin's work suggests that the path to expertise involves developing Professional Vision. In the next section, we synthesize these concepts into two design objectives for tools that support the development of Professional Vision.

3.2 Design Objectives to Support the Development of Professional Vision

We propose the following design objectives based on Goodwin's work. They are meant to serve as goals for systems that make novices engage in Goodwin's expert practices.

3.2.1 Design Objectives. Tools that support the development of Professional Vision should:

DO1: Help users understand the coding scheme. Being cognizant of the domain's *coding scheme* is a pre-requisite to using *highlights* and *material representations*. Therefore, the *coding scheme* should not be treated as complexity to hide from novices, but as necessary background. Designers should pay more attention to *essential domain knowledge*, invariants associated with the domain, than *accidental domain knowledge*, which may originate from replaceable tools and applications. For example, the concepts of shadows and highlights are essential to the experience of a photograph, whether it is digital or chemically developed, but understanding compression artifacts is a bi-product of computer technology. Designers should prioritize teaching essential domain knowledge because it is fundamental to the medium and applicable across tools. Accidental domain knowledge may be necessary to use a tool, but it may become obsolete. If possible, designers should create tools that obviate the need for accidental domain knowledge.

DO2: Help users apply the coding scheme. To "see" like a professional, novices must learn how to use a *coding scheme* in context. One strategy is to scaffold *highlighting*. Creativity support tools can serve as *material representations* that computationally highlight regions until the user can do it themselves. This can help novices simultaneously understand when a theoretical concept is applicable and serve as a concrete example for that concept.

Additionally, material representations help people externalize cognition by supporting "graphic arithmetic", where people can make inferences based on the relative size or shapes of objects. [64, 66]. Designers should consider what tasks are essential to experts and construct external representations that facilitate them. For example, *waveform monitors* help colorists balance the brightness of the image by visually lining up the top and the bottom of the waveform, which complements their analysis of the image. Like with highlighting, these material representations should scaffold expert tasks until novices can complete them on their own [49].

3.2.2 Applying Design Objectives to Color Filters. These objectives aim are meant to guide the design of tools that support the development of Professional Vision. However, an expert's Professional Vision depends on their domain of expertise. We next describe how we applied these objectives to develop a tool that helps novices make sense of color filters.

In pursuit of **DO1**, we sought to *develop a representation that explains what color filters do in terms of the color grading coding scheme*. First, based on color grading textbooks [7, 58], we found that experts commonly refer to colors in terms of hue, saturation and lightness. These terms in turn define regions of the image: highlights, midtones and shadows respectively refer to bright, average, and dark areas of the screen (Figure 5). Warm colors refer to red, yellow and orange hues, while cool colors refer to blue, green and purple hues. Second, at its most basic level, a color filter is a function that changes one color into another. Hence, the first goal for the system is to represent color filters as a change in the context of hue, saturation and lightness.

In the pursuit of **DO2**, the system should help users *apply the coding scheme in the context of the user's specific filter*. This means that the system should use *highlighting* to explain filters in terms of hue, saturation and lightness, and in terms of color regions, such as *shadows* or *warm colors*. The system should also use the same coding scheme on images. This allows for a more holistic view of the filter's behavior on arbitrary images, rather than solely the contextualized impact of the filter on a specific example image.

3.3 Applying the Framework to Color Filtering

To understand what experts currently "see" with their Professional Vision (**DO1**) and how it influences their work (**DO2**), we conducted an interview with film and photography professionals who create color filters as part of their work. The goal of this study was to validate that what we learned from textbooks was ecologically valid (Section 3.2.2) and understand how they applied color grading knowledge. We also used this interview to receive feedback on how an initial prototype of Color Field could fit in their work.

Participants. We interviewed three professionals (P1–P3; 1 female, 2 male; ages 24–35) with 4+ years of experience color grading for professional and freelance projects. All three participants had experience in photography and two had experience color grading videos. Participants were compensated with a $30 gift card.

3.3.1 Method. Participants were interviewed for one hour using remote video conferencing. They were first asked about their color grading process using recent projects as examples. They were then asked to describe three color filters using before-and-after images. Finally, we presented Color Field to participants and allowed them to freely explore the tool.

3.3.2 Findings. We describe our findings on expert's workflows and initial responses to Color Field. We then discuss terminology used to describe color filters to interpret their coding schemes.

Color Grading Workflow. Because the participants are responsible for taking the photos and videos they edit, they have the opportunity to consider the color content of the image long before editing. For example, P3 planned the location of a photography shoot around the color palettes of a product. These production decisions influence the color filters they will choose in post-production. As a result, P3 notes that many of the filters they purchase are not directly useful to them because they shoot in different locations than the filter's creator. Instead, they purchase Lightroom presets to learn how other photographers achieve certain looks.

When color grading, P1 noted that finely tuning parameters was time consuming because adjustments destined for one color occasionally had undesirable cascading effects throughout the whole image. For example, they described that *"if I changed something and the building looked weird, I would need to adjust it after"*.

3.3.3 Color Grading Coding Scheme. During the description task, we observed that experts described color filters by jumping between levels of detail: from low detail and conceptual (eg. *"Grungy"*, *"Dreamy"*, *"Professional"*), to overall regions of the image (warm and cool tones; highlights, shadows, skin tones). They occasionally referred to specific items in the scene and their colors (eg. *"gold and red in her lips"*). This suggests that the coding scheme for color grading has as a hierarchy of details. Colors can be described generally or in detail as needed. Professional tools usually operate at the intermediate level by affecting color regions, although masks also enable more specific adjustments. This suggests that *highlighting* involves mapping between these levels of detail. For example, referring to the *"red in her lips"* as too saturated means that desaturating the entire image will adjust the lips appropriately. Since the pre-made filters are only presented using high level descriptions (eg. *"Vintage", "Cinematic"*), they do not directly teach their users how to highlight the image.

Color Field Design Probe. All participants found the prototype intellectually compelling and fun to explore. P3 wanted to use it to analyze filters they have made. P1 said: *"It's a lot more data than I think about just looking at a scene [...] it's cool to see what is happening behind the scenes"*. They also appreciated that the vector field reflected how colors change rather than just the result of the filter. P2 compared this with their usual workflow: *"Sometimes you see an overall effect, but you're not sure what's happening to the image; [Color Field] is very clear, when the tile is this big, [the color] is changing"*. As a whole, this validates that the vector field visualization is capturing the appropriate coding scheme. Overall, participants had

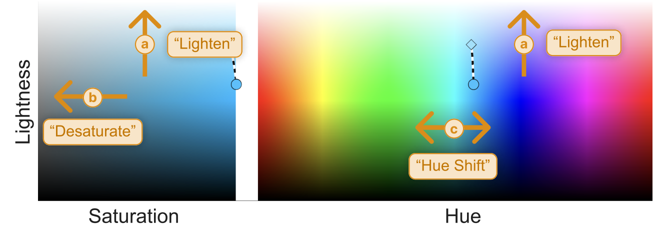

Figure 3: A key to interpret vectors in the Color Field: (a) an increase in lightness in both spectra, (b) a decrease in saturation, (c) adjusting hue—from this point in cyan, pointing left would shift towards green, yellow, etc. and moving right would shift towards blue, violet, etc.

mixed opinions on Color Field's potential for novices. On the one hand, they acknowledged the value of a tool that visually represents color grading terminology. On the other hand, they expressed concern that the system was too dense with information for novices. P3 noted that it could be very effective in the hands of an educator: *"I feel like professors should be using this in classes, you know how long it took me to understand the difference between vibrance and saturation? You need to find someone who explains it well with the right visuals"*. In response to the design probe, we shifted from a static tutorial to a guided tour of the Color Field's functionality to make new users learn the system more efficiently.

We found that the terminology used in the interview and in the description task consistent, which suggests that the task can help study how people think about color filters. This informed the design of our user study in Section 5. Furthermore, the results of the interview study validated our understanding of the expert's coding scheme and in turn that Color Field describes the right concepts.

4 SYSTEM OVERVIEW

Using the design objectives from Section 3.2, we present Color Field. The user interface has three sections (see Figure 1). In order to keep analysis grounded in the user's image, the top section displays an image before and after the application of a filter (**DO2**, Figure 1a). To highlight individual colors (**DO2**), the second section displays a *sample color* before and after applying the filter (Figure 1b). To help users interpret the filter's impact on the sample color (**DO2**), the system automatically generates a natural language description of the color filter's impact on the sample color in terms of hue, saturation and lightness (Figure 1c). Finally, to ground various concepts in terms of hue, saturation and lightness (**DO1**), the system represents the color filter's effect on all colors as a vector field over the Hue-Saturation-Lightness (HSL) color space (Figure 1d, e). The source of a vector represents a color before applying the filter, and the target of the vector shows the color after applying the filter. The sample color can be changed by clicking and dragging in the color field or in the image. This representation enables visual inference about the overall behavior of the color filter (**DO2**).

We next describe these sections of the interface in turn, explaining how they address the design principles.

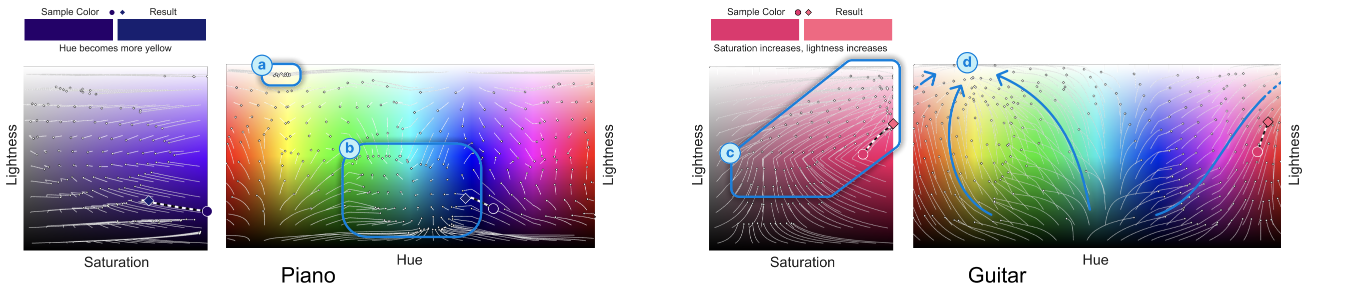

Figure 4: Vectors in Color Field draw the eye towards the colors that the filter produces in the Hue-Saturation-Lightness space (HSL). Here we show two filters, *Piano* (left) and *Guitar* (right). For Piano, the concentrated region at the top in the light yellows (a) shows that the brightest highlights are shifting towards yellow. There is a second region centered around a mid-lower cyan (b) showing that shadows are shifting towards blue. With Guitar, the Lightness-Saturation spectrum shows that magentas are getting lighter and more saturated (c), while the Lightness-Hue spectrum shows that all colors are "moving" towards light yellows, shifting hues and becoming brighter (d).

4.1 Color Filters as a Vector Field

Since experts concepts are meant to generalize across examples, a proper representation of those concepts should somehow describe all images. We make four observations:

- images are made of colors,
- color filters are functions that map one color into another,
- color pickers visualize all colors,
- vector fields represent vector functions.

Together, these observations suggest that color filters can be represented as a vector field over a color picker.

4.2 Choosing a Color Space

Color Field uses the Hue-Saturation-Lightness (HSL) color space to strike a balance between using a perceptually meaningful space, familiarity, and minimizing distortion by staying close to the computer's native representation of color (RBG).

HSL is a cylindrical model of color, where hue is the cyclical value. It is perceptually meaningful because each axis independently represents an intuitive aspect of color. Video color grading software often use Hue-Chroma-Lightness for controls, but as a double-cone model, it does not lend itself well to planar representations.

While the axes of HSL are perceptually meaningful, they are not perfectly independent from each other. Other color spaces like CIELab have more perceptually independent axes. However, since most computation occurs in RGB, visualizing those functions on a CIELab vector field would result in more artifacts. They are less common than HSL and therefore less familiar.

In summary, we selected the HSL color space as the best compromise to represent Color Field's coding scheme (**DO1**). Next, we start using the color space as a basis for describing color filters.

4.3 Interpreting the Vector Field

4.3.1 Motion in the Vector Field. Verbs like "lighten" and "saturate" can be interpreted as motion in HSL space. With Color Field they are respectively vertical motion (in both panels, see Figure 3a) and rightward motion in the Lightness-Saturation panel (Figure 3b). Horizontal motion in the Lightness-Hue panel represents a change in hue (Figure 3d). For example, a filter that makes the image "warmer" will shift hues towards oranges, reds and yellows. Therefore, Color Field satisfies **DO1** by embedding the concept of motion along these axes. Furthermore, general shifts towards a hue (eg. cyan) will appear as if the field is "gravitating" towards that hue (Figure 4). This serves as an easy to parse visual summary of the behavior of whole regions of the filter (**DO2**).

4.3.2 Curved Vector Paths. Color Field vectors trace out the path that a color would take as the filter is applied from 0% opacity, where the filter is not applied, to 100% opacity. To be consistent with other color grading software, the interface refers to opacity as "intensity". Since intensity is well defined for any filter, it is a universally meaningful control for color filters and is therefore worth representing. Additionally, this approach tends to reduce the occurrence of overlapping vectors and draws the eye as if vectors were flowing to their destinations. More importantly, the curves help guide the eye towards regions where the vector field converges (see Figure 4), making it easier to infer the filter's behavior and resulting in a stronger external representation (**DO2**).

The vector field is computed by uniformly sampling the color space, and the final colors are represented with with small white diamonds. By moving the scroll wheel up and down, the user can easily see the diamonds move along the vector paths on the field. This also serves as a user controlled animation of the color field, helping them interpret motion in the Color Field (**DO2**) [69].

4.3.3 Sample Color. The sample color's vector is styled more prominently. The source is shown as a circle that the user can drag around the color field. It can be used to pick another specific sample color, or continuously probe the field to focus on one area at a time. The filter's output for the sample color is shown as a large diamond whose position is automatically updated on the field. This lets users manually explore the space with specific colors, fulfilling **DO2**. The sample color can also be set from the image. Users can then examine what happens to the colors of an object in the image, such as the sky, and where it is in the color space. This fulfills **DO2** by embedding colors of interest from the image into a general representation.

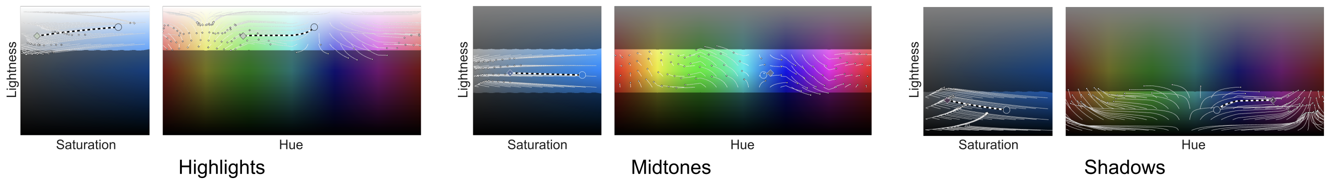

Figure 5: Visualizing the selection options of Highlights (left), Midtones (middle), and Shadows (right) within the Color Field. Since these regions are a function of lightness, they result in horizontal bands across the visualizations.

The effect of color filters is often shown with examples before and after applying a filter to an image. To our knowledge, color filters have not been illustrated in a single unified representation. By doing so on a perceptually meaningful color space, Color Field makes color filters explainable (**DO1**). We now show another way in which Professional Vision *highlights* can make the field more manageable while introducing additional color grading terminology.

4.4 Highlighting Color Regions

Color grading experts usually refer to sections of images in terms of their lightness (often also referred to as brightness). "Whites", "highlights", "mid tones", "shadows" and "blacks" are categories of pixel lightness. They are often reflected in color grading controls, where the user can affect the colors of these groups independently. Since these terms are a function of lightness, they can simply be characterized as horizontal bands in the Color Field (Figures 5, 6). This embedding serves **DO1** by visually explaining how the lower level concept of lightness relates to the higher level concepts of highlights, mid tones and shadows.

Color Field can also highlight "warm" and "cool" colors as a range of hues, and "skin tones". Together, these highlights meet **DO1** by embedding more expert terminology in the same representation.

Finally, the user can highlight the colors that appear in the image (Figure 7). This feature follows **DO2** by embedding the concrete image into the abstract representation, which simultaneously describes the image in terms of the color field (eg. "this image mostly contains oranges and blues"), and makes the color field more concrete through examples (eg. "fully saturated images are not that common").

4.5 Generating Natural Language Descriptions of Color Filters

To fulfill **DO2**, we aim to scaffold how users learn color filter vocabulary. Since vectors and selections are embedded in HSL, describing the color field can result in talking like an expert. For example, the field in Figure 6 is easily described with "the midtones get brighter and become more orange". It sounds like an expert description but was simply produced by following patterns of vectors in the field.

Since Color Field is designed for novices, it should spell out the language it wishes to elicit. Hence, Color Field generates a text description of the effect of the filter on the sample color. It simply detects the HSL dimensions in which the vector moves significantly and describes whether the color increase or decreases in each dimension. For example, with the filter used in Figure 6, selecting a purple will print "hue becomes more pink, saturation decreases" while selecting green prints "hue becomes more yellow, saturation increases, luminance increases". Hence, if the novice user becomes overwhelmed by the system, they can fall back on these natural language descriptions to suggest what they should think of the color field.

4.6 Implementation

Color Field was implemented in TypeScript as a web application. Images and the Color Field were laid out and rendered using ThreeJS to benefit from WebGL's rendering performance. Tweakpane was used for the side panel for additional user interface elements.

Color filters are stored as cube Look Up Tables (LUTs). Selections are also encoded using LUTs: the output color is white if it is selected, and black if it isn't. All LUTs have resolution between $25\times25\times25$ and $32\times32\times32$, which are common in the industry. Values that fall between indices of the LUT are interpolated from nearby indices. Combined with the conversion from RGB space to HSL, this contributes to artifacts around the edges of selections (in the image) and for colors near the corners of the RGB cube.

4.7 Summary of Design Objectives

Color Field aims to help novices learn the domain knowledge of color grading (**DO1**) and help them apply it to specific filters and images (**DO2**). The vector field achieves **DO1** with a visual representation of color filters that explains how hue, saturation and lightness relate, and shows how the color space relates to regions, such as "shadows", and colors in the image. Color Field achieves **DO2** by representing color filters as an interactive and visually inspectable vector field. It helps users apply the color grading coding scheme by visual highlighting salient regions of the color space, and it scaffolds the use of professional language with text descriptions.

5 USER STUDY

In order to understand how Color Field can fulfill our design objectives, we conducted an exploratory qualitative user study. It served as an initial exploration with rich data from a small number of participants, rather than a full pedagogical evaluation of Color Field. Although novices constituted the primary audience, the expert interview study suggested that they could provide insights based on their experience becoming experts.

Our user study addressed the following research questions:

- **RQ1**: How is Color Field used to *understand* the coding scheme? (addressing DO1)
- **RQ2**: How is Color Field used to *apply* the coding scheme? (addressing DO2)

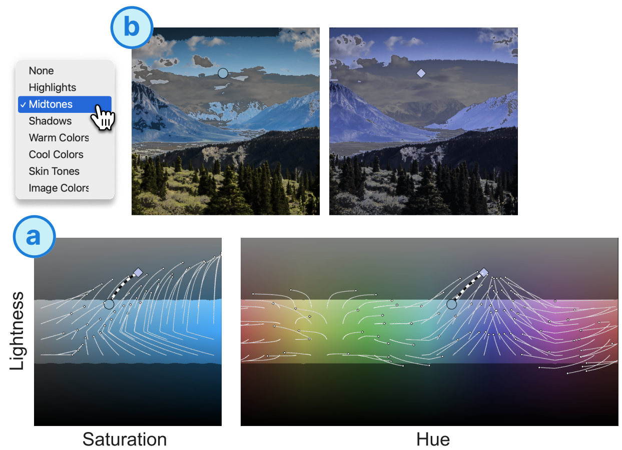

Figure 6: User selects the Midtones. Color Field highlights Midtones both in the field as a horizontal strip (a) and picks out the midtone pixels in the images (b). Note that the filtered image (right) shows fewer mid tones and that the vector field shows mid tones moving towards highlights.

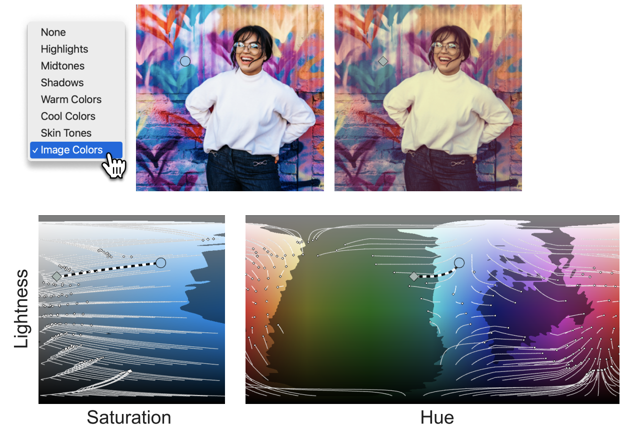

Figure 7: User selects the Image Colors. The colors used in the original image are highlighted in the Color Field.

These questions evidently overlap because *understanding* is typically a pre-requisite to *applying*. To help separate the goals of these questions by analogy, RQ1 is similar to asking how the system fulfills the role of a textbook, by teaching domain knowledge. Although textbooks can teach how to apply a coding scheme through curated examples, they cannot directly help readers apply the coding scheme to their current task. RQ2 asks how the system helps its users apply the coding scheme to in new contexts.

5.1 Participants

We recruited 4 expert participants (E1–E4; 1 female, 3 male; ages 26–33) and 5 motivated novices (N1–N5; 3 female, 1 male, 1 preferred not to say; ages 19–24) through Reddit, Slack and email. We also piloted the study with one pilot participant to iterate on study design. All experts had experience working on multiple professional projects that involved color grading. Since Color Field is designed for users who want to develop expertise, we sought to recruit "motivated" novices. We selected people with a hobbyist level of experience in visual media, minimal experience with color grading, and a self-identified interest in learning about the subject.

5.2 Procedure

The study is an informal observation of participants using Color Field to make sense of color filters.

5.2.1 Study Task. Based on the findings from the preliminary interview study (Section 3.3), this study focuses on using Color Field to analyze color filters. Although we hoped to informally observe participant's use of the system in a natural context, we decided that a structured task would ground them with a consistent objective and make the experiences of participants more comparable.

At each stage, participants were asked to describe two color filters (purchased from Rocketstock [61] and Eldamar Studio [63]). We manually selected three pairs of filters, where each pair consisted of a warm-toned and a cool-toned filter (see Figure 8). The order of filter pairs was counterbalanced to reduce learning effect biases. Each filter was named after an arbitrary musical instrument to avoid potential biases due to color associations with the names (eg. "desert" with a warm-toned filter). Participants could examine the effect of filters on two images, a portrait and a natural landscape.

5.2.2 Study Outline. The study was conducted remotely over remote video conferencing and lasted around an hour. Participants were compensated $30 for their time. Study materials (protocol, survey, etc.) are provided in the supplemental materials.

Introduction. Participants completed a consent form, a demographics form, and were asked for permission to record the video call and their screen. They were then given a study overview.

Introducing Vocabulary. Participants were first given a list of terms related to color filters with definitions synthesized from a number of color resources [7, 12, 20, 58, 72]: hue, saturation, lightness, highlights, midtones, shadows, skin tones. After reading the vocabulary sheet, participants were asked to rate their confidence in their understanding of the terms and their ability to apply them on a 5-point Likert scale.

Before Color Field. Participants launched the tool with a minimal UI showing only a before and after image. In this version of the tool, participants could only select images and filters, and adjust the intensity of the color filter. Participants were tasked with thinking aloud while analyzing two filters, one after the other. For each analysis, participants were asked to rate their confidence in the accuracy of their description on a 5-point Likert scale.

Tutorial. Participants were then walked through a tutorial that enables the rest of the features in the Color Field step-by-step: the field visualizations, the sample color before and after, the natural language description, and the selection drop-down (Figure 1). To help participants understand how to interpret Color Field, the experimenter walked through a simple filter that uniformly increases lightness. Participants were then asked to analyze a simple desaturation filter using the system. They then examined a real color filter,

Figure 8: Images provided to the participants with each pair of the warm and cool filters applied.
Images from Unsplash by Kalen Emsley (top) and Tyler Nix (bottom).

distinct from the filters used for evaluation. During the tutorial, participants were allowed to ask clarifying questions about the system and how to interpret the visualization.

With Color Field. Using the full Color Field UI, participants were asked to think aloud when analyzing two more filters.

After Color Field. The experimenter then asked the participant to disable Color Field features to describe two more filters with the simplified UI, like the *Before Color Field* condition.

Vocabulary Check. After completing the description tasks, participants were once again asked to rate their confidence in the provided vocabulary (as in *Introducing Vocabulary*).

Interview. The study concluded with an open-ended interview in which participants reflected on their use of Color Field, what they liked and disliked about the tool, difficulties they encountered, if they missed the tool in the "After Color Field" condition, and the tools overall impact on their color filter knowledge (eg. understanding of terminology, confidence in ability to describe filters).

5.3 Analysis Methods

We collected color filter descriptions, confidence measures for the descriptions and vocabulary, interview notes, as well as screen-recordings of the video conference. Here we briefly described how the qualitative data is analyzed.

Interviews. Two authors collaboratively synthesized participant quotes exhaustively into themes using affinity mapping. These themes were generated separately for novice and expert participants. When coding interviews, we focused on the impact of the embedded terminology on *understanding* the coding scheme (**RQ1**), and the spatial nature of the visualization for assisting in *applying* the coding scheme (**RQ2**). Since the primary goal of this study is to qualitatively understand how Color Field affected participants, this analysis forms the core of our results.

In our pilot study, we found that the length and content of participant's descriptions of filters varied substantially between filters for a variety of factors, such as the complexity of the filter or how long the analysis was taking. As a result, and due to the small number of participants, we refrain from thorough quantitative of filter descriptions.

6 RESULTS

First, we report evidence for each research question. Then we report other findings that characterize the participant's overall experience with Color Field.

6.1 How is Color Field used to understand the coding scheme?

In this section, we focus on ways in which Color Field helped users make sense of color grading terminology. Unlike the other research question, this question could be applied to a textbook because it focuses on communicating the coding scheme.

6.1.1 Understanding HSL. Most novices (4 of 5) appreciated the ability to visualize the relationship between hue, saturation, lightness, and the various selections. For example, N4 mentioned that *Color Field "helped me realize that I hadn't fully worked out the differences in relationship between hue and saturation"*. E4 also said that the visual representation of the vocabulary could help them communicate more clearly with clients, and make the process *"much more procedural and actionable to get me the [desired] output"*.

All novices reported that the user study left them with a deeper appreciation of Color Filters. N3 realized that *"there are more degrees of color than just red, green and blue"*, while N2 found that Color Field made them *"think more technically"*. N4 said *"I'll be reflecting on this, so thank you I learned something new today"*.

6.1.2 Validating Color Field's Coding Scheme. Most experts (3 of 4) validated that Color Field successfully represented color filters as a function. E2 found that the system was a *"very intuitive way to understand what you're doing to colors, not just what [values] you're increasing but also where those [colors] are going"*. Comparing to current professional tools, E1 points out that *"change is not something that's illustrated in scopes correctly"*. As a result, they argue that traditional software makes it too easy to focus on the color palette of the final image, forgetting the original image. They recalled a project in which a video's performer was unhappy that their dyed hair was misrepresented by the color edit.

6.1.3 Understanding Color Regions. Color region selections (eg. Midtones, Cool Colors or Skin Tones) were used as scaffolding to learn the associated terms. All novices used the selection feature and primarily used the Highlights, Midtones, Shadows, Cool Colors and Warm Colors selections because they identified those regions as important. Experts used selections less often, but when they did they almost exclusively used the Skin Tones and Image selection. E2 explains *"the location of highlights, midtones, shadows, were kind of obvious, so I didn't feel like I needed to turn those on"*, which also validates that the coding scheme is accurately represented. As a novice, N2 noted that they were developing the same perspective: *"I really only used them for the first filter [description], [...] for the second, I was kind of just like oh the highlights are up here"*. This suggests that the participants picked color regions based on their prior knowledge, which could change as they learned from the tool.

6.2 How is Color Field used to apply the coding scheme?

In this section, we focus on ways in which Color Field helped users explicitly apply the coding scheme to color filters.

6.2.1 Analyzing the Shape of the Field. By manipulating the sample color with the color picker, participants could see how the field responded to change in dimensions of HSL. Many participants (6 of 9) could extract the overall effect of the filter by looking at the *"general movement of colors"* (E2). Participants described these motions using spatial language such as *"point to"*, *"gravitate towards"*, *"bending and leaning"* (E1, E2, N3). E4 uses such language to motivate Color Field: *"What you lose with levers is how they interconnect; that's the benefit, you can literally see how the greens are going to move"*. Traditionally, spatial language is sometimes used in reference to controls in color grading software, such as *"hue shift"* in color wheels or *"lifting black levels"* with tone curves. The new spatial language found during the study suggests that Color Field can help users articulate what a filter is doing in a unified representation. In Goodwin's terms, Color Field can serve as a material representation that helps its users highlight and communicate.

6.2.2 Confirming with Text Description. Experts in the interview study (Section 3.3) predicted that only novices would benefit from the text description of the sample color. Indeed, many novices (3 of 5) used the description to confirm their interpretation of the field. However, most experts (3 of 4) also use the descriptions to *"double check that [their] eyes are saying the same thing that it says"* (E3). E1 especially benefited from this strategy when the summary was unexpected. They explain: *"I expected everything was desaturated and I looked at colors in the jeans, and it said there was a slight increase in saturation. Little things like that I generally didn't see it and it was nice to have it pointed out"*. This shows that the text description could serve as an alternate representation of the vector field both to help put a pattern in the vector field into words and to convert a verbalized hypothesis into a shape on the vector field.

6.2.3 Referring to Color Field after Withholding. Even when Color Field was no longer available, some (3 of 4) novices and one expert were still framing filters in terms of the field visualization. For example, while working through the third round of filters, N2 used their mouse cursor to draw imaginary lines where the Color Field was in the previous condition: *"You can tell I'm pretending I'm going through the color field here, like I'm imagining it in my head"*. In subsequent review of the recording, we found that their simulation was indeed accurate. N1 also tried to mentally simulate Color Field, although they struggled to remember the names of the axes, which hindered their ability to apply the proper vocabulary in context. Nonetheless, it suggests that Color Field was sufficiently intuitive to lead some participants to use it to externalize cognitive effort.

Color Field's utility was also implied by the participant's reactions when it was withheld for the *After Color Field* condition. Most participants (7 of 9) immediately expressed missing Color Field. When asked if they missed the prototype, E3 recalled thinking *"Oh no I no longer have my buddy, my assistant here"*.

6.3 Other Participant Behaviors

In this section, we report how participants learned to use the system, which features they used, and the strategies they employed to explore Color Field.

6.3.1 Onboarding. Experts in the interview study were concerned that Color Field is too dense with information for novices to understand (Section 3.3). Indeed, most (4 of 5) novices reported that the system was overwhelming at first. Despite the learning curve, all novices were eventually able to interpret Color Field, and one even found it immediately intuitive. Others required significantly more practice during the tutorial and sometimes during the first filter analysis. Experts learned how to use Color Field more easily. At first, many (3 of 4) expected that moving the sample color would change the filter, and were confused that they saw no change in the image. This is likely because color picker interfaces are typically used for editing rather than visualizing.

6.3.2 Features Used. While all participants had the opportunity to use and practice using all of Color Field's functionality during the tutorial, many (6 of 9) reported forgetting about certain features. Experts (3 of 4) mainly forgot the selection and text description features. Notably, these were the features deemed less useful to experts in our preliminary interview study. Meanwhile, novices forgot that they could pick colors from the image and about the intensity slider.

6.3.3 Exploration Strategies. Participants used multiple strategies to explore and interpret Color Field. Participants sometimes used the sample color to investigate how the filter affected the image or objects in the image, such as skin tones. Since experts were used to understanding filters from example images, this feature seemed to complement their abilities. Alternatively, participants often swept the vector field to analyze color regions. For instance, N4 preferred to select colors from the field because *"using the image made it harder to be systematic in a way that my brain was able to process"*.

Some novices (2 of 5), but none of the experts, used the text descriptions to guide their exploration. By moving the sample color around the field, they looked at the descriptions for patterns, then investigated the field lines. This behavior demonstrates how simultaneous representations can lead to strategies that would be impossible with either representation on its own.

6.3.4 Participant Confidence. For each filter, participants reported their confidence in the accuracy of their description. The results are shown in Figure 10 in the appendix. Due to the small sample and effect size, these quantitative results are on their own inconclusive and require further inquiry. However, they led to qualitatively interesting reflections during the interview.

Overall, the participant's confidence in their filter descriptions changed from filter to filter, but did not consistently increase or decrease for each participant. Similarly, participants' confidence in their knowledge of color grading vocabulary started rather high and did not significantly change. However, most experts (3 of 4) slightly lowered their rating in their confidence. When asked about the change, they explained that Color Field revealed gaps in their explicit knowledge of color grading. E4 elaborates: *"When using Lightroom and Photoshop, it's just kind of levers that you pull. You can kind of see whats going on in terms of the output on the screen, but you don't get exposed to the internal mechanism dissected like it is in the color field. [...] In the process of explaining it, I realize that I don't actually know what I'm talking about. I know the terms, but having to verbalize it to someone else is quite difficult"*. This suggests that Color Field can help make tacit knowledge about color grading more explicit by directly representing the effect of a color filter.

7 DISCUSSION

Color Field was designed to visualize color filters based on the Professional Vision of color grading experts. We next discuss the limitations and future work for this design. We then show that Color Field can also be used with other types of color filters. We conclude by arguing that the theory of Professional Vision can be applied to many other domains and serve as a powerful tool in HCI.

7.1 Limitations & Future Work

7.1.1 Representing Color Filters. Color Field's visualization conveys how colors are affected by filters at a glance. Although the curved vectors help guide the eye, they can sometimes overlap and result in visual clutter. To mitigate this problem, the source of the vectors could be optimally sampled from the color space to reduce vector overlap. Alternatively, other vector field visualizations, such as flow visualizations [44], may be appropriate. However, it is important to show the start and end colors of the filter, which rules out many flow visualization techniques.

Color Field also employs a very simple language representation of the filter's effect on a sample color. A natural next step would be to describe regions of the color field by clustering colors that behave similarly. However, language descriptions can be substantially more expressive. For instance, it could describe color filters at a higher level of abstraction like the interview study participants, with words like *"dreamy"* and *"polaroid"*. However, these descriptions are highly subjective and may be distracting if the user disagrees. Furthermore, participants in the user study primarily used the text descriptions to cross reference with the vector field.

Additionally, although color grading is often a final pass over the entire image, experts sometimes make adjustments to sections of the image using masks. By representing images in the HSL colorspace, Color Field cannot describe filters that affect spatial regions of the image. Since each mask has its own color grading adjustment, they would require independent Color Fields.

7.1.2 Editing in Color Field. Although color grading is often associated with video post-production, our prototype is limited to viewing images. In practice however, color grading is usually done using static frames from the video. As a result, Color Field's design could easily be adopted within video editors. Furthermore, in the same way that the image can be used to highlight its colors in the color field, users could import the video into Color Field to highlight all pixels in the video on Color Field and therefore visualize its color palette. This would help users provide a holistic view of the content of the video in terms of its color.

A natural extension of Color Field is to directly manipulate the vector field to create new filters. The user could select sections of the field to define the source of vectors in the filter, and move that selection to define the target of those vectors. For example, users could select the dark blue region and move it towards purple, leaving other colors unaffected. This mechanism would cover a lot of functionality in professional color grading software. In current tools, these controls are split across multiple panes, tabs and representations. Color Field could serve as a single tool to carry out all of these operations. However, such fine-grained control could easily overwhelm novices. For this paper, we decided not to pursue color filter editing because resulting design challenges exceed the purview of Professional Vision.

7.2 Other Applications of Color Field

Because we focused on essential domain knowledge for generic color editing, Color Field can be applicable to other types of color filters. We present two other applications of Color Field outside of traditional color grading.

7.2.1 Visualizing Blending Modes. Blending modes are used for a variety of tasks across digital arts. For example, digital painters sometimes overlay a translucent solid color with a specific blending mode over their work to make the image's color palette more cohesive. There are many algorithms to blend one image on top of another: some brighten the image, some darken it, others invert the colors. Because they operate in RGB space, the results are not always predictable. Explanations of blending modes, for example in Adobe Photoshop's documentation, often cryptically describe their mathematical functions with examples [21]. As a result, artists commonly resort to tweaking the top color with every mode until they are satisfied with the result. Color Field can visually and interactively represent the difference between these blending modes in terms of hue, saturation and lightness. For example, it becomes clear that *screen* primarily lightens and desaturates the image while *color dodge* lightens and significantly saturates the image.

7.2.2 Visualizing Color Vision Deficiency. Color Vision Deficiency (CVD) is a condition that affects the ability to perceive color. In most cases, one or more retinal cones are less sensitive than in individuals with common trichromatic vision. To illustrate what the world looks like to someone with CVD to people with trichromatic vision, researchers have developed simulations of various

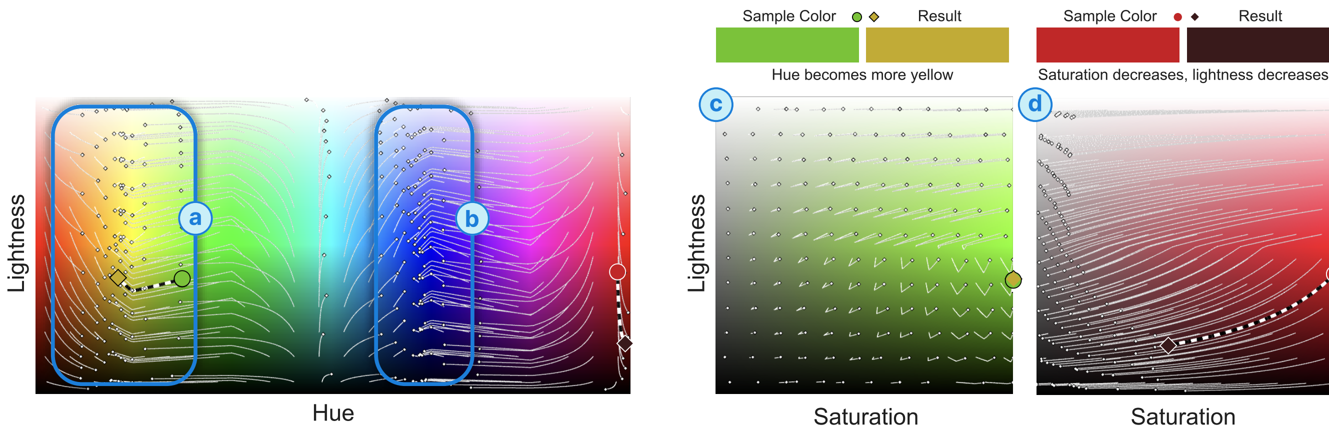

Figure 9: Color Field can represent simulations of Color Vision Deficiency (CVD). In this case, the vector field shows that protanomaly, also known as red-green color blindness, concentrates hues towards orange (a) and blue (b). The Saturation-Lightness graphs on the right show that for a similar lightness of green (c) and red (d), the change in both saturation and lightness is drastic—greens barely move, while reds desaturate and darken significantly. This is because protanomaly reduces sensitivity to red colors [27].

CVDs [27, 55]. Since they are color filters, they can be analyzed with Color Field. Figure 9 shows a simulation of protanomaly, a common form of red-green CVD, using the results of Fernandes et al. [27]. Colors clearly gravitate from red and green towards orange, implying that red and green may be harder to distinguish, as the term "red-green colorblind" suggests. Additionally, color maps designed to be accessible often use yellow and blue at the extremes to maintain high contrast for most people [10]. Color Field explains why those hues are common in this context: they are largely unaffected by protanomaly. Furthermore, one might intuitively come to the conclusion that red-green colorblindness would simply mean that red and green hues cannot be distinguished. This is true, but doesn't capture the fact that green hues maintain their lightness while red hues become much darker. This is consistent with the physiological explanation for protanomaly. Since red retinal cones are less sensitive, red light is perceived as darker [27]. This has implications for graphic design: a green and a red hue with the same lightness for common trichromatic vision will have high contrast for individuals with protanomaly because the red color will be darker, even if hues converge. However, a mid-toned red and a dark green will have high contrast for people with common trichromatic vision, but low contrast for people with protanomaly.

7.3 Other Applications of Professional Vision

Although we only applied Professional Vision to a single domain, here we argue that it can be applied more broadly.

7.3.1 Applying Professional Vision to Other Domains. Understanding the medium is central to creative endeavors. In this paper, we focused on color spaces for color grading, but other domains exhibit similar practices, such as composition in photography [23], data science [53], or even visual chunking of code blocks when programming [62].

Although Professional Vision implies a visual capacity, its central ideas are independent from modality. For example, audio engineers in music production train their ears to identify which audio frequencies instruments produce, how they overlap, and how to make them more audible. This is process involves the same highlighting principle as Professional Vision, even if it is purely auditory. In fact, audio engineers often caution beginners against relying on spectral visualizations as it can only serve as a sometimes misleading proxy for audio perception. In other words, relying on your eyes can become a crutch for your ears. This tension suggests that this area is ripe for further investigation.

7.3.2 Applying Professional Vision in HCI. We developed two design objectives based on Goodwin's theory of Professional Vision. Are they effective design tools? On the one hand, the theoretical grounding in Professional Vision helped us guide Color Field's design. On the other hand, it is possible that color grading was uniquely positioned to produce a singular representation that embeds its coding scheme. Additionally, they provide goals to reach, but little guidance on how to achieve them. To validate the general effectiveness of these design objectives, they must be applied to other domains. In doing so, we may find recurring strategies that designers can employ to fulfill the objectives and develop actionable principles for systems designed to support Professional Vision [6].

More broadly, how can Professional Vision serve as a framework in HCI? Goodwin focused on the discursive practices of Professional Vision in social contexts. As a result, his work highlights communicable explicit knowledge rather than invisible tacit knowledge. Because tacit knowledge is a hallmark of expertise, this may sound like a substantial limitation [54]. However, experts in our user study reported that Color Field helped them make their tacit knowledge more explicit, which they posited could help them communicate with others. Therefore, these design objectives might have the secondary benefit of making tacit knowledge more explicit.

As generative AIs produce better and better content, it is tempting to develop higher and higher level controls for them, such as natural language input or latent space representations. However, by eliminating granular controls we rob novices of the opportunity to learn to make low-level decisions. This is important not only to dial in details, but also to teach novices how to "see" like experts—and by extension how to appreciate their work. Goodwin's work suggests that the path to expertise involves much more than overcoming workflow inefficiencies. Therefore HCI can contribute more than creativity support tools that helps novices achieve agreeable results with minimal effort. Developing Professional Vision also increases one's awareness for the details which, pragmatically, can be useful to achieve better results, and holistically improves appreciation of crafts. If generative AIs are to eventually replace executing creative tasks, may we at least maintain our ability to appreciate the results.

8 CONCLUSION

When confronted with a creative task, some novices want the best results with minimal effort. Instead, this paper focuses on motivated novices who want to become experts in manipulating color filters. To support their journey, we turned to Goodwin's work on Professional Vision and studied how experts communicate about color filters to help novices "see" like experts. This motivated two design objectives for systems that strive to help novices develop Professional Vision: 1) *Help users understand the coding scheme* and 2) *Help users apply the coding schemes*. Using these objectives, we designed Color Field, a system to help novices develop Professional Vision for color grading. Color Field fulfills them by using a unified representation of color filters that interactively illustrates expert concepts and vocabulary for specific filters and images. We demonstrated how it can help color grading novices analyze filters and that it can lead to useful insights about blending modes and color vision deficiencies.

This paper explores how we can use Professional Vision in the context of a single domain. We envision that by applying the presented design objectives to other domains, we can better understand how to use Professional Vision as a generative framework [6].

ACKNOWLEDGMENTS

We would like to thank the reviewers for their input and helping us improve the paper. We also thank Devamardeep Hayatpur for guidance on developing the system, Jack Terwilliger for helpful discussions about Professional Vision, and members of the Creativity Lab for their feedback and generous assistance.

REFERENCES

[1] Maneesh Agrawala and Chris Stolte. 2001. Rendering effective route maps: Improving usability through generalization. In *Proceedings of the 28th annual conference on Computer graphics and interactive techniques*. Association for Computing Machinery, New York, NY, USA, 241–249. https://doi.org/10.1145/383259.383286
[2] Soonmin Bae, Aseem Agarwala, and Frédo Durand. 2010. Computational Rephotography. 29, 3 (2010). https://doi.org/10.1145/1805964.1805968
[3] Jongmin Baek, Dawid Pająk, Kihwan Kim, Kari Pulli, and Marc Levoy. 2013. WYSIWYG Computational Photography via Viewfinder Editing. *ACM Trans. Graph.* 32, 6, Article 198 (nov 2013), 10 pages. https://doi.org/10.1145/2508363.2508421
[4] Reynold Bailey, Ann McNamara, Nisha Sudarsanam, and Cindy Grimm. 2009. Subtle Gaze Direction. *ACM Transactions on Graphics* 28, 4, Article 100 (2009), 14 pages.
[5] Olivier Bau and Wendy E. Mackay. 2008. OctoPocus: A Dynamic Guide for Learning Gesture-Based Command Sets. Association for Computing Machinery, New York, NY, USA. https://doi.org/10.1145/1449715.1449724
[6] Michel Beaudouin-Lafon, Susanne Bødker, and Wendy E. Mackay. 2021. Generative Theories of Interaction. *ACM Trans. Comput.-Hum. Interact.* 28, 6, Article 45 (nov 2021), 54 pages. https://doi.org/10.1145/3468505
[7] B. Block. 2020. *The Visual Story: Creating the Visual Structure of Film, TV, and Digital Media.* Taylor & Francis. https://books.google.com/books?id=RC7yDwAAQBAJ
[8] James H Block and Robert B Burns. 1976. Mastery learning. *Review of research in education* 4 (1976), 3–49. https://doi.org/10.2307/1167112
[9] Benjamin S Bloom. 1968. Learning for Mastery. Instruction and Curriculum. Regional Education Laboratory for the Carolinas and Virginia, Topical Papers and Reprints, Number 1. *Evaluation comment* 1, 2 (1968), n2.
[10] Noam Ross Bob Rudis and Simon Garnier. 2023. Introduction to the viridis color maps. https://cran.r-project.org/web/packages/viridis/vignettes/intro-to-viridis.html
[11] Nicolas Bonneel, Kalyan Sunkavalli, Sylvain Paris, and Hanspeter Pfister. 2013. Example-Based Video Color Grading. *ACM Trans. Graph.* 32, 4, Article 39 (jul 2013), 12 pages. https://doi.org/10.1145/2461912.2461939
[12] Steven Bradley. 2013. The fundamentals of color: Hue, saturation, and lightness. https://vanseodesign.com/web-design/hue-saturation-and-lightness/
[13] Julia Cambre, Scott Klemmer, and Chinmay Kulkarni. 2018. Juxtapeer: Comparative peer review yields higher quality feedback and promotes deeper reflection. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. Association for Computing Machinery, New York, NY, USA, 1–13. https://doi.org/10.1145/3173574.3173868
[14] Huiwen Chang, Ohad Fried, Yiming Liu, Stephen DiVerdi, and Adam Finkelstein. 2015. Palette-Based Photo Recoloring. *ACM Trans. Graph.* 34, 4, Article 139 (jul 2015), 11 pages. https://doi.org/10.1145/2766978
[15] Youngha Chang, Suguru Saito, Keiji Uchikawa, and Masayuki Nakajima. 2005. Example-Based Color Stylization of Images. *ACM Trans. Appl. Percept.* 2, 3 (jul 2005), 322–345. https://doi.org/10.1145/1077399.1077408
[16] William G Chase and Herbert A Simon. 1973. The mind's eye in chess. In *Visual information processing*. Elsevier, 215–281.
[17] Fanny Chevalier, Pierre Dragicevic, and Christophe Hurter. 2012. Histomages: Fully Synchronized Views for Image Editing. In *Proceedings of the 25th Annual ACM Symposium on User Interface Software and Technology* (Cambridge, Massachusetts, USA) *(UIST '12)*. Association for Computing Machinery, New York, NY, USA, 281–286. https://doi.org/10.1145/2380116.2380152
[18] Pei-Yu Chi, Sally Ahn, Amanda Ren, Mira Dontcheva, Wilmot Li, and Björn Hartmann. 2012. MixT: Automatic Generation of Step-by-Step Mixed Media Tutorials. Association for Computing Machinery, New York, NY, USA. https://doi.org/10.1145/2380116.2380130
[19] Blackmagic Design. 2004-2023. DaVinci Resolve. https://www.blackmagicdesign.com/products/davinciresolve
[20] Adobe Discover. 2023. Learn photography: Tips, tricks & techniques. https://www.adobe.com/creativecloud/photography/discover.html
[21] Adobe Photoshop Documentation. 2023. Blending modes. https://helpx.adobe.com/photoshop/using/blending-modes.html
[22] Jane L. E, Ohad Fried, and Maneesh Agrawala. 2019. Optimizing Portrait Lighting at Capture-Time Using a 360 Camera as a Light Probe. In *Proceedings of the 32nd Annual ACM Symposium on User Interface Software and Technology* (New Orleans, LA, USA) *(UIST '19)*. Association for Computing Machinery, New York, NY, USA, 221–232. https://doi.org/10.1145/3332165.3347893
[23] Jane L. E, Ohad Fried, Jingwan Lu, Jianming Zhang, Radomír Mech, Jose Echevarria, Pat Hanrahan, and James A. Landay. 2020. Adaptive Photographic Composition Guidance. Association for Computing Machinery, New York, NY, USA. https://doi.org/10.1145/3313831.3376635
[24] Jane L. E, Kevin Y. Zhai, Jose Echevarria, Ohad Fried, Pat Hanrahan, and James A. Landay. 2021. Dynamic Guidance for Decluttering Photographic Compositions. In *The 34th Annual ACM Symposium on User Interface Software and Technology* (Virtual Event, USA) *(UIST '21)*. Association for Computing Machinery, New York, NY, USA, 359–371. https://doi.org/10.1145/3472749.3474755
[25] Elliot W Eisner. 1972. Educating artistic vision. (1972).
[26] K Anders Ericsson, Ralf T Krampe, and Clemens Tesch-Römer. 1993. The role of deliberate practice in the acquisition of expert performance. *Psychological review* 100, 3 (1993), 363.
[27] L. F. Fernandes, M. M. Oliveira, and G. M. Machado. 2009. A Physiologically-based Model for Simulation of Color Vision Deficiency. *IEEE Transactions on Visualization & Computer Graphics* 15, 06 (nov 2009), 1291–1298. https://doi.org/10.1109/TVCG.2009.113
[28] C. Ailie Fraser, Tricia J. Ngoon, Mira Dontcheva, and Scott Klemmer. 2019. RePlay: Contextually Presenting Learning Videos Across Software Applications. Association for Computing Machinery, New York, NY, USA. https://doi.org/10.1145/3290605.3300527
[29] Charles Goodwin. 1994. Professional Vision. *American Anthropologist* (1994), 606–633.
[30] Arthur C Graesser, Mark W Conley, and Andrew Olney. 2012. Intelligent tutoring systems. *APA educational psychology handbook, Vol 3: Application to learning and teaching.* (2012), 451–473.

[31] Tovi Grossman and George Fitzmaurice. 2010. ToolClips: An Investigation of Contextual Video Assistance for Functionality Understanding. Association for Computing Machinery, New York, NY, USA. https://doi.org/10.1145/1753326.1753552
[32] Andrew Head, Amber Xie, and Marti A Hearst. 2022. Math Augmentation: How Authors Enhance the Readability of Formulas using Novel Visual Design Practices. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems*. Association for Computing Machinery, New York, NY, USA, 1–18. https://doi.org/10.1145/3491102.3501932
[33] Adobe Inc. 1990-2023. Adobe Photoshop. https://www.adobe.com/products/photoshop.html
[34] Adobe Inc. 2017-2023. Adobe Lightroom. https://www.adobe.com/products/photoshop-lightroom.html
[35] Meta Inc. 2010-2023. Instagram. https://www.instagram.com
[36] Nikhita Joshi, Justin Matejka, Fraser Anderson, Tovi Grossman, and George Fitzmaurice. 2020. Micromentor: Peer-to-peer software help sessions in three minutes or less. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. Association for Computing Machinery, New York, NY, USA, 1–13. https://doi.org/10.1145/3313831.3376230
[37] Joy Kim, Maneesh Agrawala, and Michael S Bernstein. 2017. Mosaic: designing online creative communities for sharing works-in-progress. In *Proceedings of the 2017 ACM Conference on Computer Supported Cooperative Work and Social Computing*. 246–258. https://doi.org/10.1145/2998181.2998195
[38] Juho Kim, Philip J. Guo, Carrie J. Cai, Shang-Wen (Daniel) Li, Krzysztof Z. Gajos, and Robert C. Miller. 2014. Data-Driven Interaction Techniques for Improving Navigation of Educational Videos. Association for Computing Machinery, New York, NY, USA. https://doi.org/10.1145/2642918.2647389
[39] David Kirsh. 2010. Thinking with external representations. *AI & society* 25 (2010), 441–454.
[40] Yuki Koyama, Daisuke Sakamoto, and Takeo Igarashi. 2016. Selph: Progressive learning and support of manual photo color enhancement. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*. Association for Computing Machinery, New York, NY, USA, 2520–2532. https://doi.org/10.1145/2858036.2858111
[41] James A Kulik and JD Fletcher. 2016. Effectiveness of intelligent tutoring systems: a meta-analytic review. *Review of educational research* 86, 1 (2016), 42–78.
[42] Gordon Kurtenbach and William Buxton. 1994. User Learning and Performance with Marking Menus. In *Conference Companion on Human Factors in Computing Systems* (Boston, Massachusetts, USA) *(CHI '94)*. Association for Computing Machinery, New York, NY, USA, 218. https://doi.org/10.1145/259963.260376
[43] Gierad P Laput, Mira Dontcheva, Gregg Wilensky, Walter Chang, Aseem Agarwala, Jason Linder, and Eytan Adar. 2013. Pixeltone: A multimodal interface for image editing. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. Association for Computing Machinery, New York, NY, USA, 2185–2194. https://doi.org/10.1145/2470654.2481301
[44] Robert S Laramee, Helwig Hauser, Helmut Doleisch, Benjamin Vrolijk, Frits H Post, and Daniel Weiskopf. 2004. The state of the art in flow visualization: Dense and texture-based techniques. In *Computer Graphics Forum*, Vol. 23. Wiley Online Library, 203–221.
[45] Chunggi Lee, Sanghoon Kim, Dongyun Han, Hongjun Yang, Young-Woo Park, Bum Chul Kwon, and Sungahn Ko. 2020. GUIComp: A GUI design assistant with real-time, multi-faceted feedback. In *Proceedings of the 2020 CHI conference on human factors in computing systems*. 1–13. https://doi.org/10.1145/3313831.3376327
[46] Jun-Tae Lee, Han-Ul Kim, Chul Lee, and Chang-Su Kim. 2018. Photographic composition classification and dominant geometric element detection for outdoor scenes. *Journal of Visual Communication and Image Representation* 55 (2018), 91–105.
[47] JOHN H. LUNN. 1948. Chick Sexing. *American Scientist* 36, 2 (1948), 280–287. http://www.jstor.org/stable/27826226
[48] Kurt Luther, Jari-Lee Tolentino, Wei Wu, Amy Pavel, Brian P Bailey, Maneesh Agrawala, Björn Hartmann, and Steven P Dow. 2015. Structuring, aggregating, and evaluating crowdsourced design critique. In *Proceedings of the 18th ACM conference on computer supported cooperative work & social computing*. 473–485. https://doi.org/10.1145/2675133.2675283
[49] Janet Maybin, Neil Mercer, and Barry Stierer. 1992. Scaffolding learning in the classroom. *Thinking voices: The work of the national oracy project* 186 (1992), 195.
[50] Hiroko Mitarai, Yoshihiro Itamiya, and Atsuo Yoshitaka. 2013. Interactive photographic shooting assistance based on composition and saliency. In *International Conference on Computational Science and Its Applications*. Springer, 348–363.
[51] Tricia J Ngoon, C Ailie Fraser, Ariel S Weingarten, Mira Dontcheva, and Scott Klemmer. 2018. Interactive guidance techniques for improving creative feedback. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. 1–11. https://doi.org/10.1145/3173574.3173629
[52] Hyacinth S Nwana. 1990. Intelligent tutoring systems: an overview. *Artificial Intelligence Review* 4, 4 (1990), 251–277.
[53] Samir Passi and Steven Jackson. 2017. Data Vision: Learning to See Through Algorithmic Abstraction. In *Proceedings of the 2017 ACM Conference on Computer Supported Cooperative Work and Social Computing* (Portland, Oregon, USA) *(CSCW '17)*. Association for Computing Machinery, New York, NY, USA, 2436–2447. https://doi.org/10.1145/2998181.2998331
[54] Michael Polanyi and Amartya Sen. 1966. The Tacit Dimension. (1966).
[55] Katharina Reinecke, David R. Flatla, and Christopher Brooks. 2016. Enabling Designers to Foresee Which Colors Users Cannot See. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems* (San Jose, California, USA) *(CHI '16)*. Association for Computing Machinery, New York, NY, USA, 2693–2704. https://doi.org/10.1145/2858036.2858077
[56] E. Reinhard, M. Adhikhmin, B. Gooch, and P. Shirley. 2001. Color transfer between images. *IEEE Computer Graphics and Applications* 21, 5 (2001), 34–41. https://doi.org/10.1109/38.946629
[57] Roger C Schank, Tamara R Berman, and Kimberli A Macpherson. 1999. Learning by doing. *Instructional-design theories and models: A new paradigm of instructional theory* 2, 2 (1999), 161–181.
[58] Anne Morgan Spalter. 1998. *The computer in the visual arts*. Addison-Wesley Longman Publishing Co., Inc.
[59] Aaron Springer and Steve Whittaker. 2019. Progressive Disclosure: Empirically Motivated Approaches to Designing Effective Transparency. Association for Computing Machinery, New York, NY, USA. https://doi.org/10.1145/3301275.3302322
[60] Srinivas Sridharan, Reynold Bailey, Ann McNamara, and Cindy Grimm. 2012. Subtle Gaze Manipulation for Improved Mammography Training. In *Proceedings of the Symposium on Eye Tracking Research and Applications* (Santa Barbara, California) *(ETRA '12)*. ACM, New York, NY, USA, 75–82. https://doi.org/10.1145/2168556.2168568
[61] Rocket Stock. 2023. Rocket Stock LUTs. https://www.rocketstock.com/free-after-effects-templates/35-free-luts-for-color-grading-videos/
[62] M-AD Storey, F David Fracchia, and Hausi A Müller. 1999. Cognitive design elements to support the construction of a mental model during software exploration. *Journal of Systems and Software* 44, 3 (1999), 171–185.
[63] Eldamar Studio. 2023. Eldamar Studio LUTs. https://eldamar-studio.com/product-category/luts/
[64] Masaki Suwa and Barbara Tversky. 2002. External Representations Contribute to the Dynamic Construction of Ideas. In *Diagrammatic Representation and Inference*, Mary Hegarty, Bernd Meyer, and N. Hari Narayanan (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 341–343.
[65] Masaki Suwa and Barbara Tversky. 2002. How do designers shift their focus of attention in their own sketches? *Diagrammatic representation and reasoning* (2002), 241–254.
[66] Edward R Tufte. 1983. The visual display of quantitative information. (1983).
[67] Barbara Tversky. 2019. *Mind in motion: How action shapes thought*. Hachette UK.
[68] Barbara Tversky, Maneesh Agrawala, Julie Heiser, Paul Lee, Pat Hanrahan, Phan Doantam, Chris Stolte, and Marie-Paule Daniel. 2020. Cognitive design principles for automated generation of visualizations. In *Applied Spatial Cognition*. Psychology Press, 53–74.
[69] Barbara Tversky, Julie Bauer Morrison, and Mireille Betrancourt. 2002. Animation: can it facilitate? *International journal of human-computer studies* 57, 4 (2002), 247–262.
[70] Scott Valentine. 2012. *The hidden power of blend modes in Adobe Photoshop*. Adobe Press.
[71] Baoyuan Wang, Yizhou Yu, Tien-Tsin Wong, Chun Chen, and Ying-Qing Xu. 2010. Data-Driven Image Color Theme Enhancement. *ACM Trans. Graph.* 29, 6, Article 146 (dec 2010), 10 pages. https://doi.org/10.1145/1882261.1866172
[72] Wikipedia. 2023. HSL and HSV. https://en.wikipedia.org/wiki/HSL_and_HSV
[73] Ruyu Yan, Jiatian Sun, Longxiulin Deng, and Abe Davis. 2022. ReCapture: AR-Guided Time-Lapse Photography. Association for Computing Machinery, New York, NY, USA. https://doi.org/10.1145/3526113.3545641
[74] Alvin Yuan, Kurt Luther, Markus Krause, Sophie Isabel Vennix, Steven P Dow, and Bjorn Hartmann. 2016. Almost an expert: The effects of rubrics and expertise on perceived value of crowdsourced design critiques. In *Proceedings of the 19th ACM Conference on Computer-Supported Cooperative Work & Social Computing*. 1005–1017. https://doi.org/10.1145/2818048.2819953

A SELF-REPORTED CONFIDENCE FROM THE USER STUDY

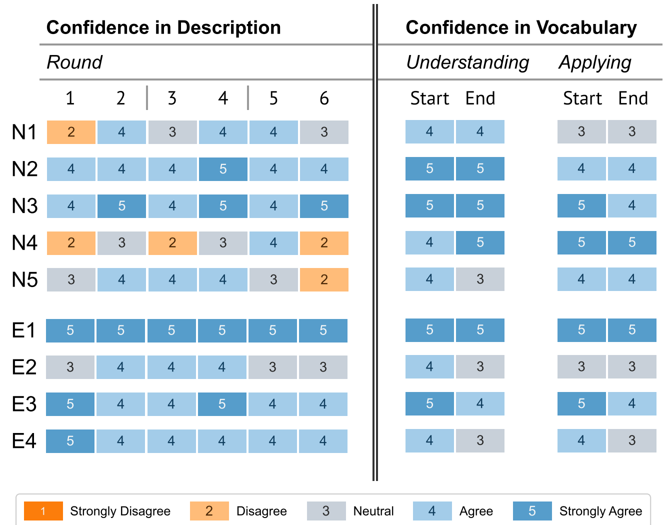

Figure 10: User study participants rated their confidence in the accuracy of their descriptions for each color filter they analyzed (rounds 1 and 2 for the *Before Color Field* condition, 3 and 4 for the *Color Field* condition, and 5 and 6 for the *After Color Field* condition). Participants also rated their understanding and ability to apply color grading vocabulary at the start and end of the study.