Simulating the Visual Experience of Very Bright and Very Dark Scenes

DAVID E. JACOBS
Stanford University
ORAZIO GALLO
NVIDIA Research
EMILY A. COOPER
Stanford University
KARI PULLI
NVIDIA Research
and
MARC LEVOY
Stanford University

The human visual system can operate in a wide range of illumination levels, due to several adaptation processes working in concert. For the most part, these adaptation mechanisms are transparent, leaving the observer unaware of his or her absolute adaptation state. At extreme illumination levels, however, some of these mechanisms produce perceivable secondary effects, or epiphenomena. In bright light, these include bleaching afterimages and adaptation afterimages, while in dark conditions these include desaturation, loss of acuity, mesopic hue shift, and the Purkinje effect. In this work we examine whether displaying these effects explicitly can be used to extend the apparent dynamic range of a conventional computer display. We present phenomenological models for each effect, we describe efficient computer graphics methods for rendering our models, and we propose a gaze-adaptive display that injects the effects into imagery on a standard computer monitor. Finally, we report the results of psychophysical experiments, which reveal that while mesopic epiphenomena are a strong cue that a stimulus is very dark, afterimages have little impact on perception that a stimulus is very bright.

Categories and Subject Descriptors: I.3.3 [**Computer Graphics**]: Picture/Image Generation—*Viewing Algorithms*; I.4.0 [**Image Processing and Computer Vision**]: General—*Image Displays*

Additional Key Words and Phrases: Adaptive displays, afterimages, mesopic vision, gaze-aware displays.

ACM Reference Format:

Jacobs, D. E., Gallo, O., Cooper, E. A., Pulli, K. and Levoy, M. YYYY. Simulating the visual experience of very bright and very dark scenes. ACM Trans. Graph. VV, N, Article X (M), P pages.
DOI = 10.1145/XXXXXXX.YYYYYY
http://doi.acm.org/10.1145/XXXXXXX.YYYYYY

David E. Jacobs acknowledges support from Google Inc.

Authors' addresses: dejacobs@cs.stanford.edu, ogallo@nvidia.com, eacooper@stanford.edu, karip@nvidia.com, levoy@cs.stanford.edu.

DOI 10.1145/XXXXXXX.YYYYYY
http://doi.acm.org/10.1145/XXXXXXX.YYYYYY

1. INTRODUCTION

The human visual system is optimized to perceive the world dynamically. For instance, a combination of saccadic eye movements and integration at higher levels of the visual system gives us the perception of a uniform sharpness across our field of view, despite the fact that for any given gaze position the visual acuity drops significantly just a few degrees away from the line of sight. A similar consideration applies to luminance: complex adaptation mechanisms allow us to perceive high-dynamic-range (HDR) environments, even though we only perceive a couple of orders of magnitude of dynamic range at any given moment [McCann and Rizzi 2011].

Standard displays, such as computer monitors, can reproduce only a fraction of the dynamic range typically encountered in natural environments. Images and videos viewed on these displays, therefore, generally do not engage the luminance adaptation mechanisms that the visual system regularly undergoes in real environments. For a display to realistically reproduce the visual experience of HDR content, it would be desirable to also reproduce the visual experience associated with adapting to a wide dynamic range.

One way to tackle this problem is to design monitors that support larger bit-depths and maximum brightnesses [Seetzen et al. 2004], thus placing the burden of performing adaptation on the observer's visual system. However, aside from the complications of turning such prototypes into commercial products, some limitations are inherently unsolvable: if very bright objects are displayed with proportionally strong radiances by the screen, the viewer may experience discomfort. Moreover, such strategies only address the case of very bright scenes: a display accurately reproducing an extremely

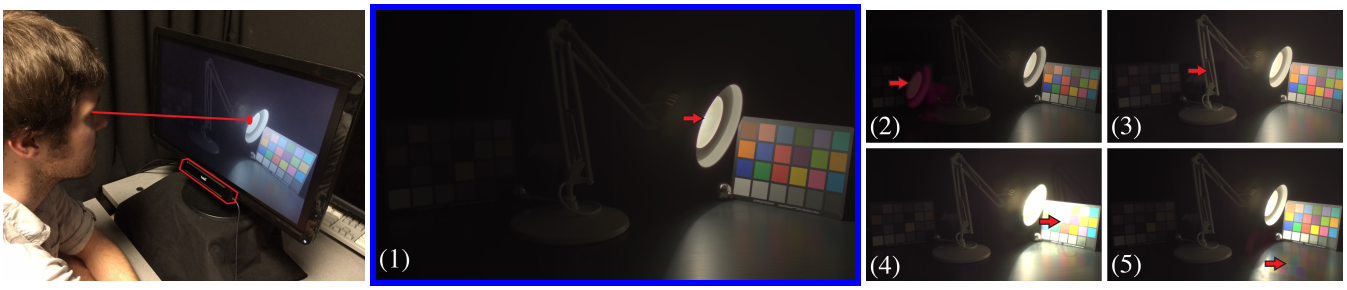

(a) Our system with gaze tracker (b) Animated interaction session (c) Additional keyframes

Fig. 1. **Animated** demonstration of our system. We propose to attach a gaze tracker to a standard display (a). With the knowledge of the user's gaze fixations, our system can synthesize several cues we experience when looking at very bright or dark objects. A demonstration of our system can be seen in (b), an animation that can be viewed in a media-enabled PDF viewer, such as Adobe Reader. (All the animated figures in this paper are marked with a blue border.) In the animation, the red dot indicates a sequence of user gaze fixations; please zoom in on the picture, click on it to start the animation, and follow the red dot. A representative subset of frames is shown in (c), where arrows indicate the fixations. Frame (2) shows the appearance of a simulated bleaching afterimage generated by the fixation in frame (1); we propose explicitly presenting this afterimage to the viewer as shown in the figure. Frame (3) shows simulated adaptation to the dark area of the image; we propose displaying the image synthetically adapted as shown in the figure. Frame (4), the very next frame in the animation, shows that local adaptation to a dark region can briefly boost the perceived brightness of other regions. Finally frame (5) shows adaptation afterimages for colorful stimuli. Lamp image courtesy of Fairchild [2008].

low-light scene could only be viewed in total darkness, as any ambient light would prevent the viewer from completely adapting to the display.

Instead of asking the viewer to adapt to the display, one can alternatively ask the display to adapt to the viewer. Such a display can be driven by the viewer's point of interest, as reported by a gaze-tracking module. Gaze-adaptive displays have been well-studied in the context of rendering [Guenter et al. 2012; Luebke and Hallen 2001]. In the context of HDR imaging, however, gaze-adaptive displays are relatively new. A recent technology demonstration simulated adaptation to global lighting changes, but did not attempt to reproduce the visual experience brought on by viewing very bright or dark objects [Rahardja et al. 2009]. If we are to improve the quality of such approaches, we must first consider three questions:

What phenomena are associated with the perception of very bright or dark objects?

At the extreme ends of the illumination spectrum, the adaptation processes of human vision manifest perceivable secondary effects, or epiphenomena. The nature and appearance of these epiphenomena is tied to the absolute radiance of the object inducing them. Therefore, it is theoretically possible that one could determine absolute brightness based on these features.

For high-photopic (extremely bright) situations, these effects include physical discomfort, glare, and retinal afterimages. In this work, we only explicitly model retinal afterimages because some glare is inherently included by the camera capturing the scene—synthetic scenes can add glare using a technique such as [Ritschel et al. 2009]—and physical discomfort would be undesirable in a practical system, even if it were a strong cue for brightness. For mesopic and scotopic (low-light) situations, these effects include a gradual change in the appearance of colors and a loss of spatial acuity.

What is the correct way to simulate these phenomena?

Recently, Ritschel and Eisemann [2012] proposed a computational model for afterimages and Mikamo et al. [2013] proposed an alternative, data-driven, model based on observations of the temporal appearance of afterimages reported by Padgham [1968]. We have developed a novel computational model that improves on both the flexibility of Ritschel's model and the colorful appearance of Mikamo's model. Our model also explicitly handles negative afterimages (see Section 2.2.3), which are not included in either of these prior works. We describe our afterimage model in Section 2.2. Although not a part of the afterimage model per se, one important property of natural afterimages is that they follow the observer's gaze. Accordingly, any implementation of an afterimage model must make an assumption about the viewer's gaze position or incorporate a gaze-tracking module. While the display system we describe in this paper does the latter, our investigation into the role of afterimages in brightness perception does the former (via an onscreen fixation target provided for experiment participants). We discuss this choice further in Section 4.1.1.

Mesopic rendering has been well-studied. Works by Kirk and O'Brien [2011] and Shin et al. [2004] have examined the effects of low-light on color perception. Low-light acuity loss has also been modeled in the graphics community [Ferwerda et al. 1996; Jensen et al. 2000; Thompson et al. 2002]. In our work, we elaborate on Shin et al.'s color model and present our own phenomenological loss-of-acuity model, which captures the time-varying appearance of objects in low-light situations—an effect that persists even once fully adapted. We describe our mesopic rendering models in Section 2.3.

Does simulation of these phenomena provide a perception of a correspondingly bright or dark stimulus?

The answer, as we will discuss further when we describe our psychophysical experiments in Section 4, varies by phenomenon. Although results presented by Ritschel and Eisemann suggest that observers can use synthetic afterimages as a cue for brightness, our experiments were unable to confirm their conclusion for naïve observers. Accurately simulating afterimages is still a worthwhile endeavor, however. With training, an observer may learn to utilize the presence and specific appearance of afterimages to determine absolute brightness even if it is not an instinctual cue. Mesopic effects, on the other hand, are strong cues for brightness and can easily "trick" a viewer into believing a scene is darker than the pixel luminances should suggest. Cinematographers have relied on this effect for decades, but until now it was unclear whether this technique takes advantage of a learned cultural norm or an inherent property

of our visual system. Our study suggests that it is an inherent property of the human vision.

In this work, our objective is not to accurately model the physiological mechanisms behind these epiphenomena—a task made even more difficult by the many questions that are yet to be answered by the vision science community. Rather, we take inspiration from what is known about the visual system and develop phenomenological models for these cues that match our experience to a level of approximation at which naïve observers may confuse them with naturally occurring epiphenomena.

To test our models, we propose a display system that tracks the user's gaze and adapts the image rendition accordingly. We artificially inject several physiologically-motivated artifacts, including adaptation to global light levels, retinal afterimages (both those due to adaptation and to the bleaching of the photopigments), visual acuity loss in low-light conditions, the Purkinje shift, and the mesopic hue shift. Some of these effects are based on previously known models, while the others we offer as novel contributions. In particular, we introduce a new model for afterimages, and show that it is more realistic than prior work. We also present a novel model for loss of visual acuity in mesopic conditions, which again compares favorably with the state-of-the-art. Finally, we describe psychophysical experiments that measure the significance of these phenomena to the perception of brightness and report our findings. An animated demonstration of our system in action is presented in Figure 1—some of the figures in this paper include animations. These figures are denoted by a blue border and feature captions beginning with "**Animated**". Please view these figures electronically using a media-enabled PDF viewer such as Adobe Reader. (Most web browser PDF plug-ins are not media-enabled.) With the exception of the looping animation in Figure 12, all the animated figures require the reader to click on the image.

2. GAZE-AWARE DYNAMIC RENDERING

In this section we enumerate the visual effects our system implements, and describe our models for them.

2.1 Global adaptation

Although the focus of this work is on the extreme ends of the illumination spectrum, a thorough analysis cannot leave out the mechanisms that govern adaptation and dynamic range compression throughout the full range of luminances. These mechanisms include the pupillary reflex, neural adaptation at various levels of the brain, center-surround photoreceptor signal encoding, and chemical adaptation of the photoreceptors themselves. The net effect of all these mechanisms is to parameterize a function which maps a large dynamic range in the world to a smaller range of electrical signals sent elsewhere in the brain, that is, a tonemapping function. We will use this tonemapping function as the base on which we will apply the rest of our effects.

Perceptually-motivated tonemapping functions have been extensively studied. Many of the proposed methods share some variation on the Naka-Rushton equation [Reinhard et al. 2010], which predicts the instantaneous response of a photoreceptor to a new stimulus after being adapted to a particular background illumination. The Naka-Rushton response $R(I)$ is given by

$$R(I) = \frac{I^n}{I^n + \sigma^n}, \tag{1}$$

where I is the luminance of the stimulus, n is a contrast control, and σ is the intensity at which the response of the photoreceptor is equal to one half. In the log-domain, this equation represents an S-curve, or a soft step function. The adaptation state is encoded in σ. If a steady stimulus $\bar{I}$ is applied, the response will weaken and eventually rest at a value given by a plateau function $M(\bar{I})$. We can solve for the appropriate σ for a background illumination $\bar{I}$ by setting $R(\bar{I}) = M(\bar{I})$ to get

$$\sigma(\bar{I}) = \left(\frac{\bar{I}^n}{M(\bar{I})} - \bar{I}^n \right)^{1/n}. \tag{2}$$

The plateau function $M(\bar{I})$ can be measured and has a form similar to Equation 1 [Dowling 1987].

We treat Equations 1 and 2 together as a tonemapping curve, which we apply globally to the photopic luminance channel of the image. The final displayed pixel value is $J = R(I)^\gamma$, where γ is chosen to compensate for non-linearity in the display's response. Color images are handled by enforcing a consistent ratio between color channels before and after the tonemapping curve is applied.

In our system, we represent the user's current adaptation state as the illumination level A (in log-units) to which he or she is adapted. We determine A from the history of the user's gaze positions and the content of the scene. Let the log of the average scene luminance in a small patch around the user's gaze position be called the target A_T. For a timestep of duration Δt, we find the updated adaptation state A' by moving A towards A_T as follows

$$A' \leftarrow \begin{cases} A + a_1 \Delta t & A < A_T \\ A - a_2 \Delta t & A > A_T \end{cases}, \tag{3}$$

where a_1 and a_2 are the rates of adaptation to brighter and darker stimuli, respectively. Our model includes separate adaptation rates for the two directions because the visual system takes longer to adapt to dark environments than to bright ones [Kalloniatis and Luu 2007]. If the update step would cause A' to overshoot its target A_T in either direction, we set it equal to A_T. A' can then be converted to linear luminance $\bar{I}$ and used in Equation 2 to appropriately parameterize the tonemapping curve. An example of our system's global adaptation is given in Figure 2. The effect of the simulated global adaptation is to darken the scene overall when a relatively bright point is fixated and to brighten the scene when a relatively dark point is fixated.

Fig. 2. Global adaptation. The user's fixations, indicated by the red arrows, determine the adaptation level used to render the scene. Image courtesy of Fairchild [2008].

The method described above computes only the expected *response* of the early visual system when exposed to a particular scene. Ultimately, our goal is to reproduce the *appearance* of such a scene in the visual system of the viewer, which would require compensating for the properties of both the display and the viewer's real-world visual system. Our global adaptation model makes the simplifying assumption that photoreceptor response and visual appearance are equivalent, but is still able to produce results of good quality for our purposes (i.e., believable images into which our

Fig. 3. Inducing afterimages. Bleaching afterimages are caused by extremely bright light sources, such as the camera flash shown in (a). To observe a bleaching afterimage most clearly, watch a camera flash firing from a safe distance (1-2 meters), then quickly close your eyes and cover them with your hand. Depending on the brightness of the flash, you should see a bright green or cyan afterimage which slowly fades to a pink or deep purple color. Adaptation afterimages are produced by long fixations on steady stimuli. To induce an adaptation afterimage, fixate on the cross at the center of (b) for several seconds, then fixate on the cross at the center of (c). You should see a faint checkerboard pattern with the brightness relationships of the quadrants reversed. Note that the adaptation afterimage is not strictly brighter or darker than the background, but rather brighter for some parts and darker for others. The same process can be performed with the colored stimulus in (d) to induce a color-opponent afterimage.

epiphenomena models can be injected). A more involved appearance model, like those described by Pattanaik et al. [2000] or Irawan et al. [2005], might further improve the quality of the results, but would not impact the main contribution of this work.

All of the parameter settings used in this paper are reported in Appendix A. These suggested parameters were selected to strike a balance between the realism of the simulated phenomena and usability of the display. Where possible, we started with values known to the vision science community and adjusted them until a reasonable result was produced. For example, requiring a viewer to dark adapt for several minutes before viewing a scene, as would be necessary in the real world, would make the system impractical; therefore we tune its parameters to speed up this process while maintaining a plausible feeling of adaptation. We do not claim these parameter choices to be optimal, but rather offer them as a reasonable starting point.

2.2 Afterimages

In an effort to provide the perception of looking at radiances brighter than the physical capabilities of the display, we create and display synthetic afterimages.

"Afterimage" is a general term used to describe any latent image of a real-world object that persists after the original stimulus ceases. After being induced, they are "locked" to a region of the retina and will therefore move with the viewer's gaze, until they completely fade. We observe that afterimages vary in strength and appearance depending upon the nature of the stimulus that induced the afterimage. For example, a short-duration, very bright, color-neutral source like a camera flash will produce a strong cyan afterimage, fading to green, and finally to magenta [Weve 1925]. Alternatively, fixating on a dark, color-neutral region enhances the receptors' sensitivity, while fixating on a solid color will induce a cast of the opponent color (e.g., a red stimulus produces a cyan cast). Figure 3 gives additional instructions for inducing these commonly-seen types of afterimages. We can help explain the distinct appearance of the two kinds of afterimages by taking a moment to discuss the details of a physiological process known as phototransduction.

2.2.1 *Phototransduction.* Phototransduction is the process by which photoreceptors convert light into electrical signals. An understanding of this process can help explain the appearance and behavior of the retinal afterimages that we model in this work. An abstracted view of the process is sufficient for our purposes; for more details, see the survey by Fain et al. [2001].

Photoreceptors contain hundreds of thousands of light-sensitive proteins called photopigments. The aggregate behavior of these photopigments determines the response of the photoreceptor. The photopigment life-cycle can explain a particular kind of afterimage that we call a *bleaching afterimage*. Shortly after being struck by photons, photopigments enter a bleached state, where they are no longer sensitive to light, but continue to contribute to the photoreceptor's response [Hodgkin and Nunn 1988]. This means that the photoreceptor continues producing a response even after the end of the stimulus, i.e., it produces an afterimage. Eventually, bleached photopigments are restored to a receptive idle state, which ceases their contribution to the transduction cascade and therefore to the afterimage. Baylor et al. [1974] showed that the time dynamics of this process in the cones of turtles can be modeled as decaying exponentials. Our parameterization of their model produced the curves shown in Figure 4.

The other kind of afterimage, which we refer to as a *local adaptation afterimage*, can be partially attributed to the role of calcium ions as a regulating chemical in transduction. The role of calcium in phototransduction is complex [Matthews 1996] and still an active area of research; but for our purposes we assume a net effect equivalent to a simple gain, because of the strong correlation between calcium concentration and the overall sensitivity of the cell [Torre et al. 1986]. The calcium influx and efflux rates are finite, so changes in the cell's calcium concentration will lag behind a changing stimulus, temporarily boosting or weakening photoreceptor's response until it can reach a new equilibrium, i.e., producing an afterimage.

Recently, Ritschel et al. [2012] presented a computational model for simulating afterimages. Their model does not explicitly distinguish between the two kinds of afterimages we describe above and therefore fails to reproduce important behaviors. Similarly, the model presented by Mikamo et al. [2013] does not include adaptation afterimages. In the following subsections, we present a unified model for bleaching and local adaptation afterimages. While the appearance of these epiphenomena is certainly influenced by higher-level visual processing, we choose to develop our model purely at the photoreceptor level. We make this choice for two reasons: (1) the physiology governing retinal contributions to afterimages is better understood, and (2) a photoreceptor-centric model is sufficient for producing believable afterimages.

2.2.2 *Bleaching afterimages.* Bleaching afterimages are caused by sources of light that are bright enough to bleach large portions of photopigments in a photoreceptor. These afterimages are positive, meaning they are always brighter than the background on which they are superimposed. A bright background, however, can easily mask a bleaching afterimage induced by a stimulus of similar brightness. The appearance of a bleaching afterimage varies with time and with the radiance of the light source. The time-dependent nature of afterimage appearances has been well-studied and is commonly referred to as the "flight of colors" [Weve 1925]. We demonstrate the relationship between source brightness and afterimage appearance in an experiment we describe in Section 4.3. Figure 5 shows our model's color appearance prediction as a function of time and stimulus intensity.

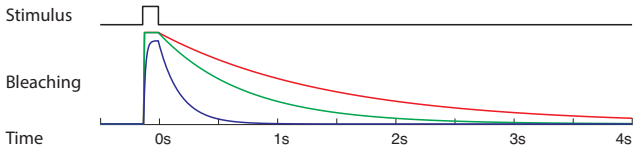

Fig. 4. Photoreceptor bleaching time course. A bright stimulus "bleaches" a portion of a photoreceptor's photopigments; these molecules remain bleached even after the stimulus is removed, see Section 2.2.1. The differential change in bleaching level induced by a stimulus is predicted by Equation 4 [Baylor et al. 1974]. Above we show the predicted bleaching levels over time of an L-, M-, and S-cone (shown in red, green, and blue, respectively) exposed to a 1/8 second (s), 10^7 candela/meter^2 (cd/m^2), white (stimulating all cone types equally) flash, ending at time $t = 0$s. We selected the parameters governing the three cone types' bleaching behaviors to produce phenomenologically plausible afterimages. The corresponding afterimage color appearance for these bleaching levels is shown below in the 10^7 cd/m^2 line of Figure 5.

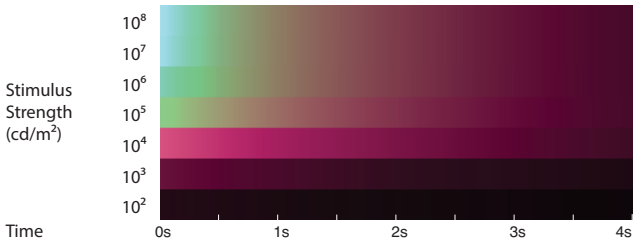

Fig. 5. Bleaching afterimages and stimulus intensity. The bleaching levels of the three cone types determines the appearance of a bleaching afterimage. Here we show our predicted bleaching afterimage appearance over time in response to brief flashes ending at time $t = 0s$. The corresponding flash intensities are shown on the left side of each bar. As the intensity of the flash grows, the afterimages increasingly start from cyan, always fading to magenta. The sudden transition between 10^4 and 10^5 cd/m^2 is a result of the relatively strong impact M-cones have on the final RGB afterimage appearance. The first hint of green begins appearing around $10^{4.5}$ cd/m^2, but is not shown here. The corresponding cone bleaching levels for the 10^7 cd/m^2 flash are shown above in Figure 4.

We define the bleaching level B of a photoreceptor as the fraction of photopigments in the bleached state. The fraction of bleached pigments increases as incident light strikes idle pigments and decreases as bleached pigments are restored to their idle state. According to Baylor et al. [1974], the change in bleaching level B for a photoreceptor is given by the differential equation

$$\frac{dB}{dt} = b_1(1 - B)I - b_2B, \tag{4}$$

where b_1 is a bleaching sensitivity parameter, I is the incident luminance, and b_2 is the recovery rate of the photoreceptor's pigments. We solve it analytically to find the updated bleaching level B' after a simulation timestep of size Δt as

$$B' \leftarrow (B - B_\infty)e^{-(b_1 I + b_2)\Delta t} + B_\infty, \tag{5}$$

where B_∞ is the equilibrium bleaching level assuming a constant stimulus, given by

$$B_\infty = \frac{b_1 I}{b_1 I + b_2}. \tag{6}$$

At the equilibrium state for any given light level a photoreceptor will have some portion (B_∞) of its photopigment bleached. However, we believe that the contribution of these bleached pigments to the overall response of the photoreceptor is normally imperceptible and thus we consider it implicitly included as part of Equation 1. Therefore, to compute how visible a bleaching afterimage should be, we consider the *deviation* from equilibrium bleaching rather than the absolute bleaching level. Specifically, to compute the linear brightness J for an output pixel, we add bleaching to the default photoreceptor output

$$J = R(I) + j_1 \cdot \max(B - B_\infty, 0), \tag{7}$$

where j_1 is a parameter controlling the strength of the bleaching afterimages. This linear brightness will be later modified by a gamma transform for display.

In practice, we apply these equations for each cone type separately after converting the scene to the LMS colorspace. Each type of cone has different bleaching and recovery parameters [Bedggood and Metha 2012], which after converting back to RGB (see color matrix H in Appendix A) results in colorful, temporally-changing afterimages ("flight of colors"). As a result, we only need to modify a few parameters per photoreceptor type to tune the appearance of the afterimages; by comparison, Mikamo et al.'s data-driven afterimage model requires hundreds of datapoints to specify the time-varying appearance of afterimages. Additionally, our model handles novel situations like compounding afterimages on top of afterimages, which are not supported by Mikamo et al.'s model.

The method proposed by Ritschel and Eisemann [2012] is similar to ours in principle, as it uses an update rule akin to Equation 4. However, they do not account for the different time dynamics of the L, M, and S cones; moreover, they compute the afterimages after performing a color opponency transform. As a result, the afterimages of a color-neutral stimulus synthesized by their model are neutral as well, and fail to capture the colorful nature of the bleaching afterimages. See Figures 6 and 7 for visual comparisons.

Fig. 6. **Animated** bleaching afterimages. When viewed in a media-enabled PDF reader, such as Adobe Reader, this figure shows a comparison of the time dynamics of the afterimages generated by a flash, as predicted by our model (left), Ritschel and Eisemann's model [2012] (middle), and Mikamo et al.'s [2013] model (right). Mikamo et al.'s model predicts a longer-lasting afterimage, so we show it at 2× playback speed. Please zoom in and click on one image at a time; before starting each animation, wait for the other animations to finish. Notice that the model by Ritschel and Eisemann fails to reproduce the color casts that are commonly observed when exposed to a very bright stimulus. Readers of the printed edition of this paper should refer to Figure 7 for key frames from the above animations.

2.2.3 *Local Adaptation Afterimages.* Local adaptation afterimages are caused by exposures to stimuli lasting long enough for individual photoreceptors to adapt locally. Unlike bleaching afterimages, this second type of afterimage can also be negative, in the sense that adaptation can reduce the response of the photoreceptors; this is apparent when looking at the crosshair in Figure 3(c) after fixating the crosshair in Figure 3(b). The physiological source of local adaptation is a combination of sensitivity loss due to bleaching (bleached pigment does not react to light) and the varying cellular

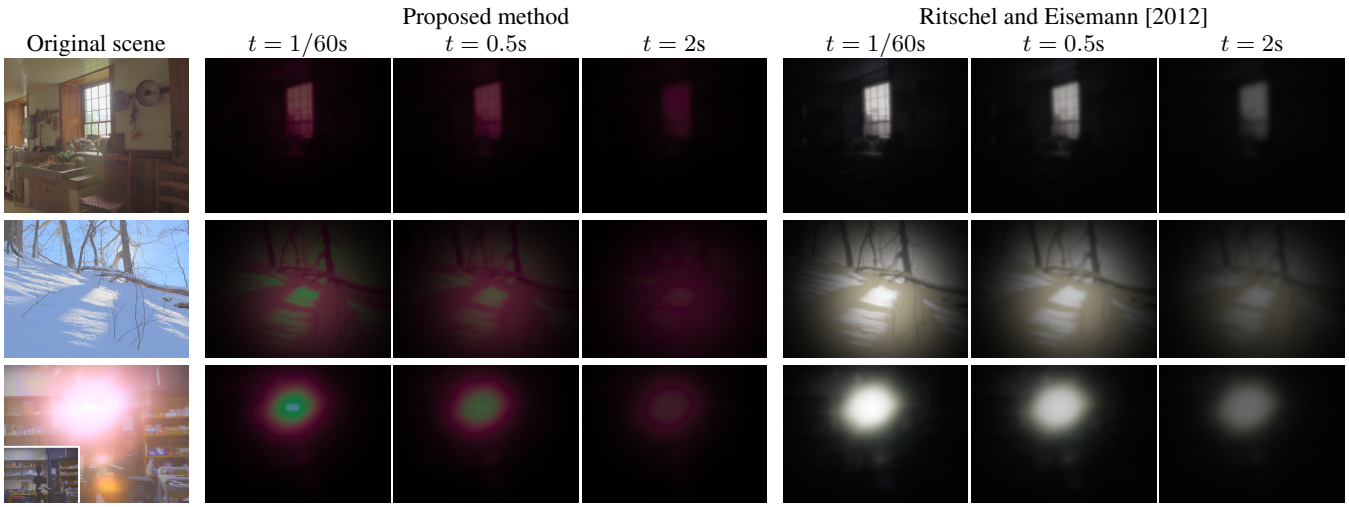

Fig. 7. Bleaching afterimages. To illustrate the appearance of bleaching afterimages as predicted by our model, we simulate a viewer shutting his or her eyes after a brief exposure to a scene (i.e., we suddenly drop the scene radiance to zero). The appearance of the afterimage t seconds after the "blink" is shown in each column as marked. The gaze position is also held constant in this figure—recall that afterimages will move with the viewer's gaze position. From top to bottom, the maximum radiance of the scenes are approximately 3.5, 4.5, and 8.5 log cd/m^2, respectively. These scenes are representative of the range of bleaching afterimage appearances our model can produce. In particular, note the presence of cyan, green, and pink elements in each afterimage. We also provide Ritschel and Eisemann's result for comparison; note that their model cannot capture the colorful nature of afterimages produced by a color-neutral stimulus. An animated version of the flash afterimage is shown in Figure 6. The images in the top two rows are courtesy of Fairchild [2008].

concentration of calcium ions. (Neural adaptation plays a role as well.) Our model ignores the effect of bleaching on photoreceptor sensitivity. This lets us decouple the parameters controlling bleaching and local adaptation afterimages.

Calcium enters the photoreceptor through transduction-controlled ion channels, and is constantly moved out by active pumps. According to Fain et al. [2001], the differential change of intracellular calcium concentration C over time is given by

$$\frac{dC}{dt} = c_1 S - c_2 C, \tag{8}$$

where c_1 and c_2 control the influx and efflux of calcium, respectively, and S is the fraction of open ion channels in the photoreceptor's membrane (which is inversely correlated to the magnitude of the photoreceptor's response). Using this equation directly can lead to numerical instabilities, so instead we approximate S by $1 - R$, where R is from Equation 1. The main difference here is that $1 - R$ does not include afterimages, while S does. In our model, we apply the analytical solution to this differential equation to find the updated calcium concentration C' after a simulation timestep of size Δt as follows:

$$C' \leftarrow (C - C_\infty) e^{-c_2 \Delta t} + C_\infty, \tag{9}$$

where C_∞ is the equilibrium calcium concentration assuming a constant R, i.e., constant stimulus and no global adaptation. C_∞ is given by

$$C_\infty = \frac{c_1(1 - R)}{c_2}. \tag{10}$$

Similar to the case of bleaching afterimages, we primarily concern ourselves with deviations from the equilibrium calcium concentration C_∞ when determining the appearance of a local adaptation afterimage. We apply a local adaptation gain to the photoreceptor response as computed by:

$$\alpha = \frac{C_{\max} - C_\infty}{C_{\max} - C}, \tag{11}$$

where $C_{\max} = c_1/c_2$, the maximum possible equilibrium calcium concentration (which occurs in total darkness). Note that when $C < C_\infty$, $\alpha < 1$, and vice versa. We choose this particular definition of α because we find it gives plausible results; other functions that ensure a similar relationship between C, C_∞, and α would be equally reasonable. We now modify Equation 7 to include the effects of local adaptation:

$$J = \alpha^{j_2} \cdot R(I) + j_1 \frac{C}{C_{\max}} \max(B - B_\infty, 0), \tag{12}$$

where j_2 controls the strength of the adaptation afterimages. Note also that the bleaching afterimage term is multiplied by the normalized calcium concentration. This modification weakens the appearance of bleaching afterimages in bright situations and strengthens them in low-light situations. Figure 8 shows the contribution of this term visually.

Just as with bleaching states, we track the calcium concentration for each of the cone types at each pixel separately. Accordingly, fixating on a strong blue stimulus will lower the sensitivity of the S-cones, while the L- and M-cones adapt to become more sensitive. If later a neutral gray stimulus is presented, the S-cone response will be dampened while the L- and M-cones' response will be boosted, resulting in a yellow percept. This is the source of the color-opponent nature of adaptation afterimages predicted by our model. The same principle applies for neutral-color stimuli, which can reduce or increase the sensitivities of small regions of the retina. Figure 9 shows our model simulating these kinds of afterimages. We use the same parameters c_1 and c_2 for each cone type. This ensures that adaptation afterimages do not produce unnatural color casts.

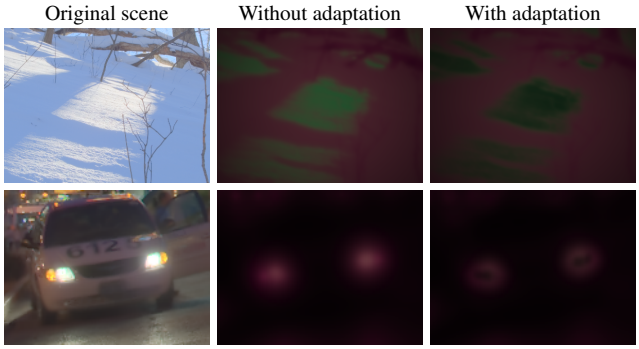

Fig. 8. Combining afterimage types. The gain induced by calcium affects the entire transduction process, including cascades initiated by bleached photopigments. We can combine adaptation afterimages with bleaching afterimages to reflect this. From left to right: the original scene, bleaching afterimages without local adaptation applied, and bleaching afterimages with local adaptation applied (the C/C_{max} term in Equation 12). The images are courtesy of Fairchild [2008].

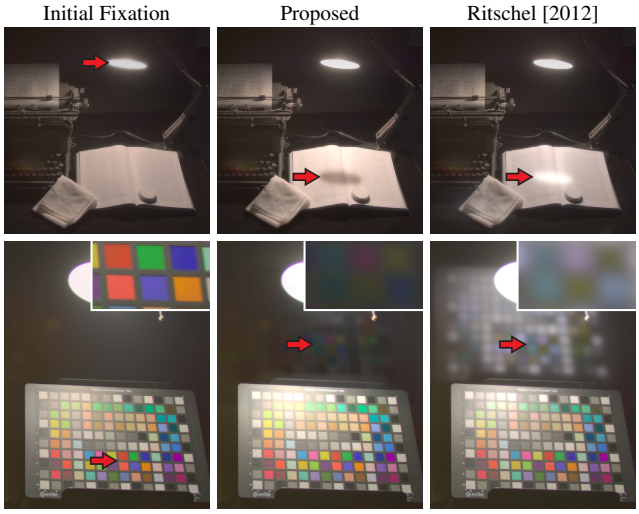

Fig. 9. Local adaptation afterimages. Long fixations give individual photoreceptors time to adapt. When the gaze moves to a new position, the spatially-varying photoreceptor sensitivities manifest themselves as an afterimage of the first fixation. The top row shows the adaptation afterimage of a bright white lamp, which appears as a dark spot on the book. The model by Ritschel and Eisemann predicts a less plausible, bright white spot. The bottom row shows the adaptation afterimage of a strongly colored stimulus. Both models produce color-opponent afterimages (e.g., the blue squares become yellow), but our model does so by modulating the background appearance, while Ritschel and Eisemann overlay an afterimage which is strictly brighter than the background. In other words, their model only predicts positive afterimages. The image in the top row is courtesy of Fairchild [2008].

In contrast to our model, the approach proposed by Ritschel and Eisemann [2012] can only generate positive afterimages. As shown in Figure 9, this causes bright objects to generate afterimages that are always brighter than the background, an unnatural result. Overall, our model is the first to combine simulations of bleaching and local adaptation afterimages and to capture the variations in color and time course.

2.3 Mesopic rendering

The human visual system operates in three different regimes depending on the light level of the environment. In brightly lit situations (>10 cd/m^2), the photopic regime is dominant: cones work well, colors are easily distinguished, and spatial visual acuity is maximal. In very dark situations ($<10^{-3}$cd/m^2), the scotopic regime is dominant: rods work best, color sensitivity is lost, and spatial acuity is low. Any illumination level between photopic and scotopic is referred to as mesopic, and is characterized by a gradual loss of color sensitivity and acuity.

In our system, we insert three cues intimating the darkness level of the scene: the Purkinje effect, the mesopic hue shift, and the low-light loss of spatial acuity. Some of these cues have been studied before in the context of traditional tonemapping [Kirk and O'Brien 2011]. Recall, however, that the goal of our system is not to produce a single tonemapped image, but rather to provide a dynamic, adaptive display. Wanat and Mantiuk's recent work [2014] studies these effects as well, but their goal is distinct from ours. They aim to make a dark display appear bright, while we aim to make a bright display appear dark.

2.3.1 *Purkinje effect.* The Purkinje effect shifts the luminous efficiency function from a peak around 555nm in photopic conditions to a peak near 507nm in scotopic conditions, due to the transition from the operating domain of cones to that of the rods [Wade and Swanston 2013]. The perceptual effect is to alter the brightness relationships of blue and red stimuli depending on the overall illumination: a red stimulus judged "brighter" than a blue stimulus in photopic conditions may appear "darker" in scotopic conditions. Strictly speaking, the Purkinje effect does not refer to any perceived change in color appearance; it only refers to the change in brightness perception—we will discuss color effects in Section 2.3.2.

In our model, we incorporate the Purkinje effect by modifying the RGB→luminance transform used before applying Equation 1. Photopic luminance, for example, is computed as a weighted sum of RGB values with most weight on the R and G channels. Scotopic luminance, on the other hand, puts little to no weight on R and favors G and B instead. (Refer to Appendix A for the specific weights.) We define the mesopic luminance to be a blend of the photopic and scotopic luminances. Although the interaction between rods and cones in the mesopic regime is an active area of study [Pokorny and Cao 2010], we are not aware of any work that fully measures the luminous efficiency curve as a function of both wavelength and adaptation level. For simplicity, we assume this function to be a linear interpolation weighted by the normalized log-illumination level ρ, which is given by

$$\rho = \frac{L - m_2}{m_1 - m_2}, \tag{13}$$

where L is the log-photopic luminance; m_1 and m_2 denote the upper and lower bounds of the mesopic luminance range in log-units, respectively. Note that when $L \geq m_1$, $\rho \geq 1$, and when $L \leq m_2$, $\rho \leq 0$. For use in our Purkinje effect model, we clamp ρ to lie within the range [0,1]. Figure 10 shows the appearance of a Purkinje-shifted image. The Purkinje effect is difficult to notice

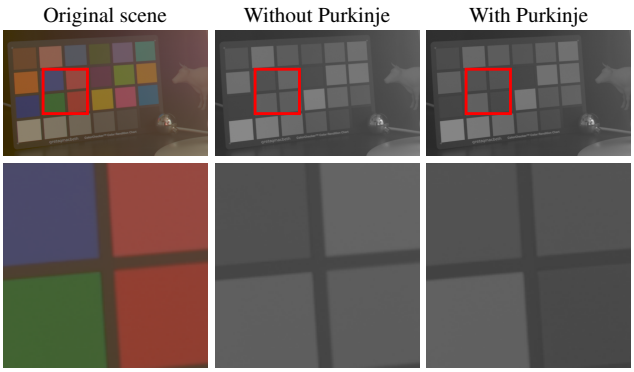

Fig. 10. Purkinje effect. As the illumination decreases, the perception of the relative intensities of different colors is altered. The leftmost column shows the original scene without any mesopic cues. The two other columns show the luminance channel of the scene (middle column) and the luminance channel of the scene after applying the Purkinje effect (rightmost column). Notice that the blue square becomes slightly brighter and the red square becomes darker when the Purkinje shift cue is active. Image courtesy of Fairchild [2008].

and probably has little impact on the perceived realism or brightness of a scene—we do not include it in the perceptual experiment described in Section 4.2. However, it is a well-documented phenomenon, so we include it in our full display system.

2.3.2 *Hue shift and desaturation.* In addition to the Purkinje effect, a loss of color discrimination and a shift in perceived hue accompany mesopic vision. Shin et al. [2004] performed a color-matching experiment at various illumination levels. They found that, as illumination decreased, the reported color of swatches, including neutral gray ones, would drift towards a dull purple color, with some small variation depending on the photopic hue. The model recently proposed by Kirk and O'Brien [2011], while better motivated physiologically, predicts a saturated cyan color for mesopic scenes, which does not match well with Shin et al.'s results or our own personal observations. Alternative models have also been proposed in the vision science community [Cao et al. 2005; Buck et al. 1998; Nerger et al. 2003; Kim et al. 2009]. Because our goal is only to produce believable results, we choose to use a simplified version of Shin et al.'s model. In our system, this means replacing the ratio between color channels used to reconstruct a color image after global adaptation with a weighted average of the scene's actual color ratio and the scotopic color ratio of a neutral gray color chip (i.e., a dull purple color), weighted by ρ clamped to [0,1]. Figure 11 shows the strong desaturation and tendency towards purple predicted by our system.

2.3.3 *Spatial acuity loss.* The final cue we inject into our system is the loss of spatial acuity in low light. Throughout the mesopic range, spatial acuity drops linearly with log-luminance [Riggs 1965]. This can be seen in the inset shown in Figure 11. The loss in acuity can be explained by the reduced response of the cones and the fact that signals of many rods are aggregated as part of the input to any given ganglion cell on its way to visual centers of the brain.

The effect is often modeled with a spatially-invariant, Gaussian blur [Ferwerda et al. 1996] or a bilateral blur [Jensen et al. 2000]. Thompson et al. [2002] suggest that temporal noise is an important component of a low-light rendering model. In studying this phenomenon, we noticed that the acuity loss itself is temporally-varying. More specifically, we did not observe a *uniform* and *constant* blur; rather, the percept was a severely blurred representation of the scene with small irregularly shaped patches of locally higher resolution appearing and disappearing at random. This effect is best witnessed in a environment just dark enough to make reading a book difficult, but not quite impossible. In our experience, we found an office with no windows or interior lights (lit only by the small amount of light entering from underneath the closed door) to be near the correct brightness. We also recommend using a resolution chart held at arm's length as a stimulus after ten or more minutes of adaptation time. Our observation on the time-varying nature of the low-light loss of acuity is not mentioned in the literature as far as we know. It is possible that this phenomenon may be caused by the eyes' difficulty of maintaining a consistent focus in the dark, or an interaction between fast microsaccades and the relatively slow light response of rod photoreceptors. Without knowing the underlying cause, in an informal study, we confirmed that this effect was clearly visible to multiple individuals. We therefore describe here how we incorporate it into the mesopic model.

We have found we can replicate the time-varying appearance of this phenomenon as a stochastic, progressive loss of high-frequency detail, where the probability of including a particular frequency component is a function of the local luminance level. This acuity-loss transform is performed as a preprocess to the scene on each frame. Taking inspiration from Laplacian pyramids [Burt and Adelson 1983], we first create a bank of Gaussian kernels with exponentially increasing standard deviations σ_i (see Appendix A). We then take their differences and convolve the scene to create bandpass images G_i. To reconstruct the acuity-corrected scene G, we compute the sum of bandpass images

$$G = \sum w_i(\rho, X) \cdot G_i, \tag{14}$$

where ρ is given by Equation 13, X is a random sample drawn from the normal distribution $\mathcal{N}(0, 1)$, and w_i is a binary weight function. We choose

$$w_i(\rho, X) = (\rho + g \cdot X > \rho_i), \tag{15}$$

where ρ_i is the expected normalized log-luminance cut-off level for the i^{th} frequency band and g is a parameter controlling the variance of the stochastic acuity loss. Because acuity drops linearly with log-luminance, ρ_i are distributed uniformly (in the log-domain) across the mesopic luminances. Drawing independent samples at each pixel and each frame gives an unwanted appearance of typical image noise; instead, we draw samples X at a lower spatial and temporal resolution than the scene and interpolate between the samples. The sampling rate of X should be a function of the display resolution and distance to the viewer; in this work we assume a typical office viewing environment (a 0.6m-diagonal 1920×1080 monitor, 1m away from the viewer). Figure 12 shows an animated demonstration of our stochastic acuity model and compares it to prior approaches.

2.4 Adapting our model to available displays

In creating our implementation, we discovered a number of practical considerations that necessitated adjustments to our models.

As we mentioned in Section 2.3, the viewing environment of our system may affect the quality of the experience. Our global adaptation cue is particularly sensitive to ambient illumination. Specifically, the illusion of darkness breaks down whenever the global adaptation level A falls below that of the ambient light, as doing so renders dark objects with implausibly bright pixel values. To prevent this from happening, we artificially restrict the viewer's ability

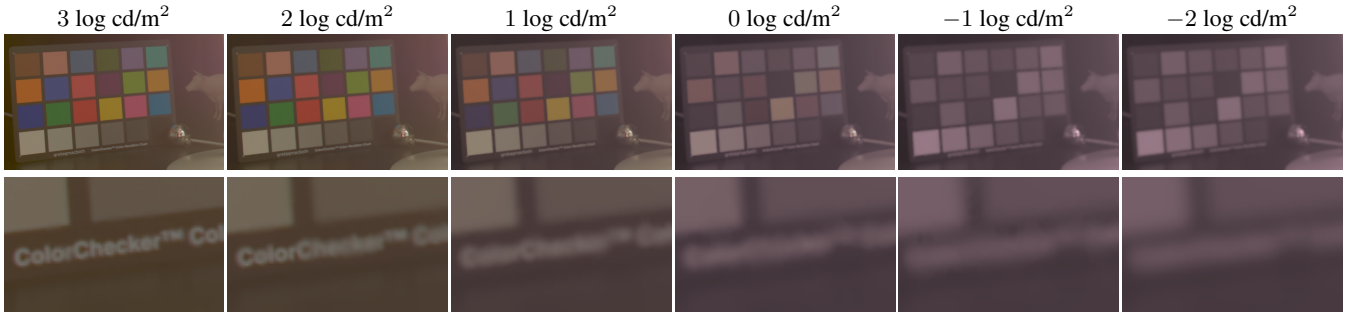

Fig. 11. Mesopic effects. Above we show a scene rendered with varying levels of mesopic cues injected, while keeping the mean pixel brightnesses constant. The simulated luminance of the white square at the bottom left corner of the color checker is shown above each column. As the luminance of a scene decreases, the scene loses color saturation and tends towards a dull purple color. This transition is best appreciated on a color-calibrated monitor when viewing the PDF file. The bottom row of inset images shows the gradual reduction in visual acuity trending with luminance as well. Image courtesy of Fairchild [2008].

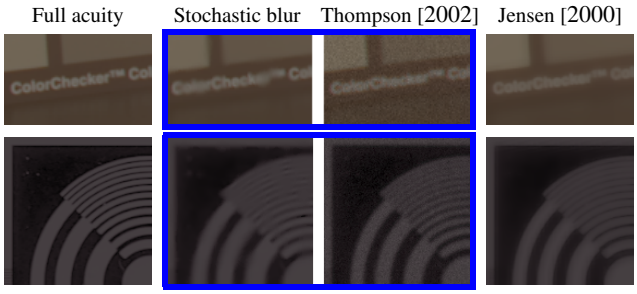

Fig. 12. **Animated** stochastic loss of acuity. Our model produces a spatially- and temporally-varying loss of visual acuity. We compare our model with Thompson et al.'s [2002] and bilateral filtering [Jensen et al. 2000]. Please view the blue-outlined part of this figure in a media-enabled PDF reader such as Adobe Reader to see the temporal dimension of the acuity model. Note that the variance of the acuity loss is exaggerated here for illustrative purposes.

to globally adapt to any luminance lower than 100 cd/m^2, which roughly corresponds to a dimly-lit office.

One consequence of this choice is that for darker scenes large portions of the screen map to near-zero pixel values. This frequently results in visible quantization bands. In order to hide banding, we add a temporal dithering stage before final display. For an 8-bit display, a pixel with unquantized value $J = J_I + J_F$ where J_I is an integer in $[0, 255]$ and J_F is a real number in $[0, 1)$ will be rounded up to $J_I + 1$ with probability $p = J_F$ (otherwise rounded down to J_I). This probabilistic rounding is applied independently to each color channel of each pixel for each frame. Note that any banding in the figures of this paper or supplemental material are compression artifacts not present in the live system.

The rest of our considerations revolve around afterimages.

Ritschel and Eisemann point out that natural afterimages often appear blurry (this is a separate observation from the mesopic loss of acuity). They suggest this could be due to chemical diffusion processes in the retina. We believe another reasonable explanation could be retinal "motion blur," brought on by microsaccades moving the eye during the integration time of a photoreceptor. In either case, we incorporate this effect by applying a Gaussian blur to the bleaching and calcium maps between each displayed frame. If our motion blur theory is right, a gaze tracker with a higher sampling rate may obviate the need for this step.

Another issue we discovered while prototyping our system is the importance of the periphery of the visual field. The periphery is especially sensitive to motion, so any afterimage motion far from the fovea can become distracting, especially when exaggerated by gaze tracker latency. To reduce this effect, we introduce a Gaussian fall-off on the visibility of afterimages, centered on the fovea.

One last concern is the bleaching and adaptation properties of photoreceptors (which in turn determine the strength and duration of afterimages). Although the vision science community has performed explicit measurements of these characteristics [Rushton and Henry 1968], using their experimental values directly in our phenomenological models does not result in a plausible appearance. We simply do not model all of the components of the visual system requisite for producing good results with these parameters. We have shown in this paper, however, that one can find parameters for our model that do produce good phenomenological results, even if they do not directly correspond to physical measurements.

3. IMPLEMENTATION

We implemented our gaze-aware rendering model on a desktop computer equipped with a 3.5GHz quad-core processor, NVIDIA GTX670 GPU, and a Tobii Rex Developer Edition gaze tracker [Tobii 2014]. This compact gaze tracker (visible in Figure 1(a)) affixes to the bottom-edge of the computer monitor and reports gaze positions at 30Hz to an accuracy of about 1°, or ∼0.5cm on our screen at normal viewing distances. Although 30Hz is too slow for foveated rendering applications, it works well for our application due to the relatively slow speed of adaptation processes. The accuracy is also not a major concern, as most of our large, diffuse cues work well as long as the offset between the actual gaze position and reported gaze position remains consistent.

Our software is a collection of OpenGL shader programs with OpenCV calls for preprocessing. The image files are calibrated radiance maps stored in OpenEXR format. Without heavy optimization, this code runs in excess of 60FPS at 1920x1080 resolution on our system. All of the images in this paper (excluding Figures 1(a) and 3) were collected by dumping raw frames from our software.

4. PERCEPTUAL ASSESSMENTS

We conducted three perceptual experiments to assess the realism and effectiveness of our phenomenological display manipulations. The goal of these experiments was twofold: (1) to determine if observers would be tricked into believing that the afterimage and mesopic manipulations were being created by their own visual sys-

tem, rather than synthesized on the display, and (2) to measure the effects of these manipulations on the perception of brightness.

4.1 Afterimages and perceived brightness

We wanted to determine if applying an afterimage would make a flashed image appear realistically brighter. We conducted a perceptual experiment in which participants judged the brightness relationship between two versions of a flashed image—one with a strong synthetic afterimage and one without. We then analyzed the results to determine the relative intensity at which these two images were subjectively equivalent to each other—that is, the point at which a participant could not reliably detect which image was brighter. If afterimages are a brightness cue, we predict that a flashed image followed by an afterimage should be subjectively equivalent to a flashed image that is physically brighter but is not accompanied by a correspondingly powerful afterimage. The procedure and results are described below.

4.1.1 *Methods. Participants:* Thirteen participants were recruited from a university student population (age range 18-22, eight females). All participants reported having normal or corrected-to-normal vision, and were naïve to the experimental manipulation. Two of these participants were excluded from the experiment based on performance during a training session, as described below. The research protocol was approved by the Stanford University Institutional Review Board and all participants gave informed consent.

Stimuli: We used the method of constant stimuli experimental design. For each trial, participants were sequentially presented with two images (in randomized, counter-balanced order): a "standard" image, which was held constant throughout the experiment, and a "test" image, which varied with each trial. Each image was presented on-screen for 50ms and followed by a black screen (upon which afterimages would be displayed) shown for 2.5s. The standard image was the camera flash scene (Figure 6) scaled to an intermediate intensity level (27 cd/m^2, as measured by a Photo Research PR-715 spectroradiometer [Photo Research Inc. 2014]). Test images were drawn from a set of images identical to the standard, but scaled to different intensity levels distributed evenly in log space around the standard: two were brighter, two were darker, and one was identical (14, 19, 27, 35, and 49 cd/m^2).

The control portion of the experiment (one half of the trials) consisted of simple relative brightness comparisons which we used to establish a baseline for the precision of perceptual judgments. Participants simply judged which of the two images shown on a given trial was brighter. In these control trials, both the standard and the test images were followed by a very weak synthetic afterimage (2.2×10^{-2} cd/m^2, $j_1 = 0.0003$). This was included to prevent participants from catching on to the experimental manipulation included in the manipulated trials. For the manipulated trials, the test images were identical, but were given an afterimage that was strong (6.5×10^{-2} cd/m^2, $j_1 = 0.001$). See our supplemental material for a video example of our stimuli. All stimuli were presented in a completely darkened room (1.8×10^{-4} cd/m^2 with black screen displayed) on a Sony PVM-2541 OLED monitor. We chose this monitor for its excellent performance near black [Ito et al. 2013; Cooper et al. 2013], which improves the appearance of subtle afterimage effects (and mesopic scenes, which is critical for the experiment we describe in Section 4.2).

Procedure: Participants were told they would be viewing a series of videos of camera flashes firing on a special monitor that can get brighter than a normal display. We claimed to have adjusted the power of the flash unit in each video, and asked participants to indicate by keyboard whether the first or second flash was set to be brighter in each trial. We did not use the gaze-tracker for this experiment, instead, we provided a gaze-fixation target on the screen just below the position of the flash unit. We made this choice to minimize the risk of breaking the illusion with a calibration procedure and potential tracking jitter. Fixation targets are a commonly accepted practice in psychophysical experiments where gaze position is critical, so while in principle this experiment does not explicitly test gaze-enabled afterimages, in practice we believe our results should apply to gaze-aware implementations as well. The experiment consisted of three blocks of trials, with the trial presentation order randomized within each block. The initial set of 20 trials was used for training: participants compared pairs of unmanipulated images with large brightness differences. These trials were used for assuring that the participants understood the task, and their results were not retained for analysis. Two participants were excluded because they responded correctly less than 70% of the time during training—indicating that their discrimination threshold was too high or they were simply not paying close enough attention to provide meaningful data. The 70% cutoff was determined prior to beginning the experiment. The second set of trials contained the control measurements. Each test image was paired with the standard 20 times. Before the final set of trials, the experimenter claimed to adjust the monitor's brightness settings. This last set of trials contained 30 repetitions of each manipulated test condition and 10 additional repetitions of each control test condition.

4.1.2 *Results. Quantitative analysis:* We estimated the point of subjective equality (PSE) separately for the control and main experiments. Conceptually, the PSE is the relative intensity of the test for which a participant cannot tell the difference from the standard. This is reflected in the responses as the point at which they select the test to be brighter for 50% of the trials, i.e., they are performing at chance.

Figure 13 illustrates how this point is estimated. For each participant (three representative examples for this experiment are shown in the top row), we plotted the percentage of trials in which the test image was selected as brighter than the standard image for each of the five test relative intensities. The results for the control experiment are shown in blue, while green shows the results for the trials with strong synthetic afterimages. The data for each experiment were fit with a cumulative Gaussian using MATLAB's `glmfit` function to estimate a continuous psychometric function. The inverse of each psychometric function evaluated at 50% is our estimate of the test intensity that is subjectively equivalent to the standard, shown as a dashed vertical line in Figure 13.

As expected, the control results (blue) show that participants tended to select the darker test images (relative intensity less than zero) at rates below-chance, the brighter test images at rates above-chance, and the equiluminant test images at rates near-chance. The steeper the S-curve, the better the subject could discriminate brightness differences between the control and test images; a step function would represent perfect discrimination, and a horizontal line at 50% (or scattering of points with no obvious fit) would represent random answers. Test image luminance values were selected with the goal of making the task relatively difficult, which is reflected by the fact that performance rates are mostly between 10-90%. The data thus only spanned the middle portion of the fitted cumulative Gaussian (shown as solid), but typically included the key range of values around chance. The results from the manipulated trials (green) could be shifted either to the left or to the right of the control data.

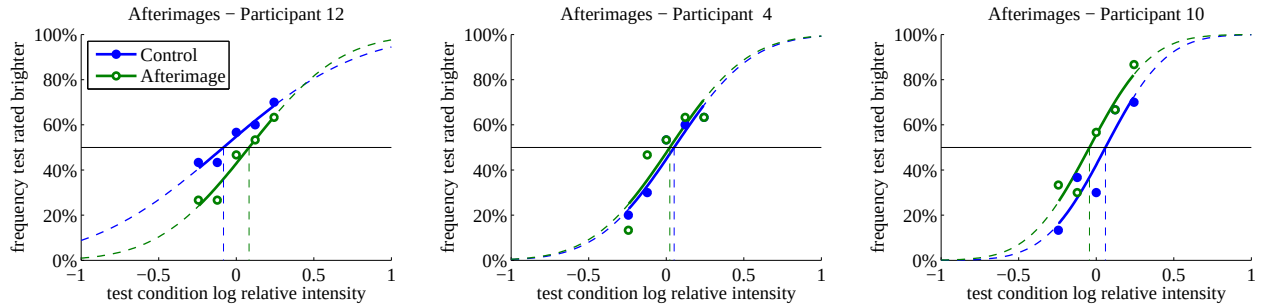

(a) Afterimage experiment representative individual results

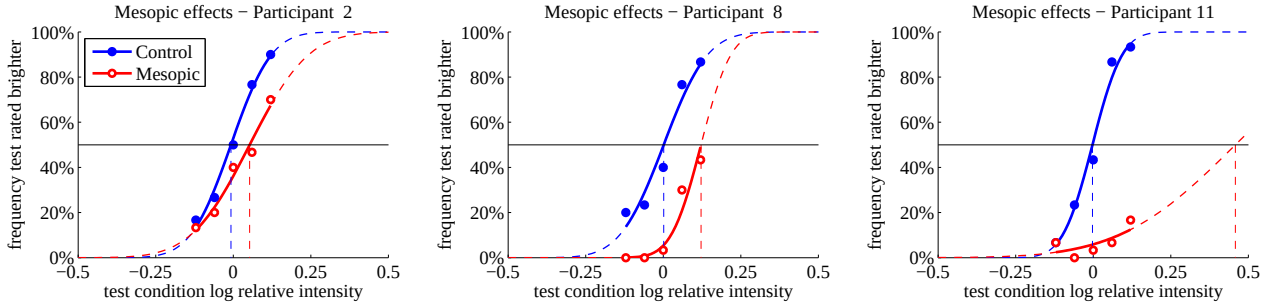

(b) Mesopic effects experiment representative individual results

Fig. 13. Representative individual results for the afterimage and mesopic effects experiments. Each plot shows raw datapoints fit with a cumulative Gaussian curve for the control and manipulated condition trials. The dashed vertical lines in each plot show the point of subjective equality (PSE) for each condition, i.e., the relative intensity at which an observer cannot reliably distinguish between the standard and the test image. Within each row, from left to right, we show results for observers near the weakest, median, and strongest responses to the manipulated condition as a brightness cue. If afterimages were a cue for brightness, the afterimage condition curves (top row, green) would be consistently shifted to the left of the control condition curves (blue). If mesopic effects were similarly effective as a cue for dark stimuli, the mesopic condition curves (bottom row, red) would consistently be shifted right of the control condition curves (blue). Our results (summarized in Figure 14) suggest that the second is true, but the first is not.

Figure 14 (top) shows the PSEs for all participants for the control and main afterimage experiments. The median of each sample set is shown as a box. As expected, the control PSEs are very near zero—recall that the test image at zero is identical to the standard for control trials. For the main experiment, the median is slightly negative (i.e., in the direction indicating afterimages are a positive cue for brightness). A ranksum test on the PSE samples from both experiments determined that the difference between the medians of the two distributions was not significant ($p = 0.449$), so our data do not support the conclusion that afterimages are a cue for brightness.

Qualitative analysis: Following the completion of all trials, we debriefed participants to discover whether they became aware of the manipulation in the course of the experiment. If so, they would have been able to choose the test or standard by way of a cognitive strategy and not because they actually perceived a stimulus as brighter or darker. Our debrief included questions such as "What aspects of the images did you use to judge their brightness?" and "Did you notice any strange things in any of the images?" Most of the eleven participants were unable to describe a concrete strategy beyond choosing whichever stimulus "felt" brighter. One participant did explicitly mention using afterimage duration as a cue for brightness, and two others reported being suspicious that some afterimages were synthetic (the other nine participants assumed the afterimages were natural). Despite the reports from these three participants, their results were similar to the rest of the participants'. The results from these interviews suggest that our model of afterimage appearance is realistic enough to deceive the majority of observers into believing that the visual effect is created by their visual system and not synthesized on-screen. The reported realism of our synthetic afterimages also gives us confidence that our quantitative results would be similar had we used a gaze-enabled implementation for this experiment. We will explore the realism of different types of afterimages further in Section 4.3.

4.2 Mesopic effects and perceived brightness

Just as in the afterimages experiment, our goal here is to determine whether applying our mesopic rendering techniques affects the apparent brightness of a displayed image. We used a similar procedure to the afterimage experiment, substituting mesopic rendering effects in place of strong afterimages. We describe the differences between the experiments and our results below.

4.2.1 *Methods. Participants:* The same set of participants described in Section 4.1.1 took part in this experiment. We excluded one additional subject from analysis because upon debriefing it was clear he had seen through the experimental manipulation.

Stimuli: In this experiment, the standard image was the color checker (Figure 11) shown for 1.25s and scaled to an intermediate intensity level (5.5×10^{-2} cd/m^2 at the white tile), followed by a black screen shown for 0.75s. For control trials, the test images were again scaled to different intensity levels around the standard (4.1×10^{-2}, 4.8×10^{-2}, 5.5×10^{-2}, 6.8×10^{-2}, and 7.5×10^{-2} cd/m^2). For the manipulated trials, the test images were rendered with the color and acuity mesopic effects described in Sections 2.3.2 and 2.3.3 before intensity scaling.

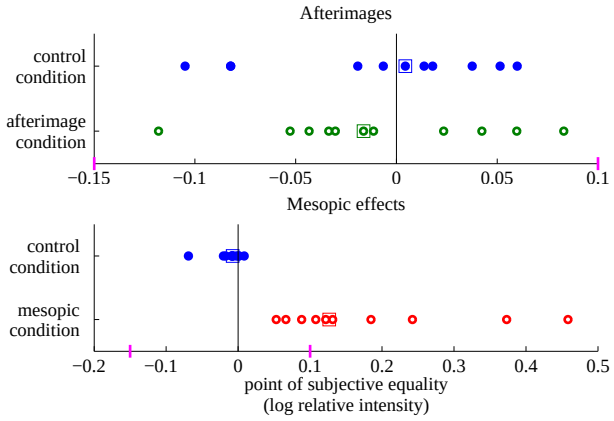

Fig. 14. Perceived brightness: Top, the distributions of points of subjective equality (PSE) for the control and main afterimage experiments. The median of the main afterimage experiment PSE samples was -0.02 log units, but a ranksum test of the two distributions revealed no statistically significant difference. We do not have sufficient evidence to claim that afterimages have an effect on brightness perception. Bottom, the distributions of PSEs for the control and main mesopic experiments. The median of the main mesopic experiment samples was 0.13 log units, which was significantly different from the control experiment. These results provide strong evidence that simulating mesopic visual effects can convincingly make a displayed image appear darker than it actually is. Note that the horizontal scales of the two plots are not the same—the overlapping region is indicated by the magenta tick marks on each plot's horizontal axis.

We did not include the Purkinje effect (Section 2.3.1) in the manipulated stimuli. Unlike the other mesopic epiphenomena, the simulated Purkinje effect modifies the physical, objective brightness of a stimulus. Accordingly, the impact of the Purkinje effect on perceived brightness may vary with the part of the scene each participant used to make their brightness judgments—e.g., a participant looking at the red tile of the color checker would likely see the manipulation as a negative cue for brightness (because the tile becomes physically darker), while a participant looking at a blue tile may see it as a positive cue (because the tile becomes physically brighter). These cues would impose a constant bias on the results, thus masking the effect of other epiphenomena. For these reasons we do not incorporate the Purkinje effect with the other epiphenomena in our experiment. A separate experiment isolating the Purkinje effect could confirm the relationship between Purkinje-based objective brightness changes and subjective perception, but it would not alter the importance of the results we present here. We leave performing such an experiment for future work.

Procedure: Participants were told they would be viewing a series of images of a very dark scene on a special monitor. We claimed to have adjusted the brightness of an out-of-frame lamp in the scene, and asked participants to indicate by keyboard whether the first or second image had the lamp set to be brighter in each trial. Before beginning the experiment, participants spent five minutes in complete darkness. The participants were led to believe this was long enough for full dark adaptation (full adaptation takes much longer). This element was important for maintaining the illusion that the display was dark enough for the mesopic effects to be natural. The rest of the experimental procedure was identical to that of the afterimages experiment.

4.2.2 *Results. Quantitative Analysis:* We performed a similar analysis on the results of this experiment. Representative individual results are shown in Figure 13(b). As in the previous experiment, the control results (blue) appear to be near-chance at the zero luminance difference test condition. These results tended to have steeper slopes than the control results from the previous experiment, which simply indicates that the task was easier for participants to perform. The manipulated condition data (red) showed a much larger bias for this experiment, and were consistently shifted rightward, indicating that the mesopic images had a strong tendency to appear dark. The distribution of PSEs for this experiment is shown in Figure 14 (bottom). Obviously, it is difficult to assess the exact PSEs for datasets such as the one shown for Participant 11, where the 50% point must be extrapolated from data far from the PSE. However, it is clear that this participant perceived the mesopic images as much darker than the standard, so we use the extrapolated PSE as a reasonable metric of this bias. Unlike the afterimages experiment, the ranksum analysis of these results shows strong evidence that these mesopic effects are a cue for brightness (significant at $p = 2.84 \times 10^{-6}$).

Qualitative Analysis: Our debriefing procedure for this experiment was identical to the procedure for the afterimage. Ten of the eleven participants said that darker images were less colorful. Only five of the participants mentioned a similar relationship between dark images and blurriness. Many participants expressed suspicion that something in the images was being manipulated (which is natural, since they are coming in for an experiment), however, when asked to elaborate, they were either unable to describe it, or were clearly confusing true brightness changes with mesopic cues (i.e., they reported seeing manipulations during the control experiment, in which none were actually present). One participant was able to correctly identify the manipulated stimuli nearly every time (and act on it), so his results were excluded from the quantitative analysis above. The interview results suggest that the mesopic effects not only make images appear to be darker, but do so in a way that the majority of participants cannot deliberately distinguish them from true darkening.

4.3 Natural afterimage appearance

Having established that our model of afterimages created realistic visual effects, we were interested to know what properties of a model's predicted afterimage appearance makes it look more or less realistic. Specifically, we wanted to compare the results of our model to the model of Ritschel and Eisemann [2012]. In one regard, the temporal properties of the appearance of bleaching afterimages have been well studied in the psychology literature [Padgham 1968; Weve 1925]. To our knowledge, however, no one has reported on the relationship between afterimage appearance and the brightness of the light source that induced the afterimage. In order to investigate the apparent colors and time course of real afterimages, we performed an additional study which we describe below. Rather than viewing on-screen images of a camera flash, subjects were exposed to real camera flashes of varying brightness, and asked to describe the real afterimage appearance in relation to options displayed on a computer screen. If there is a relationship between source brightness and afterimage appearance, we would expect participants to describe a different appearance for each flash intensity.

4.3.1 *Methods. Participants:* We recruited ten participants from a university student and faculty population (age 23-59, two females). All participants reported having normal color vision and normal or corrected-to-normal visual acuity.

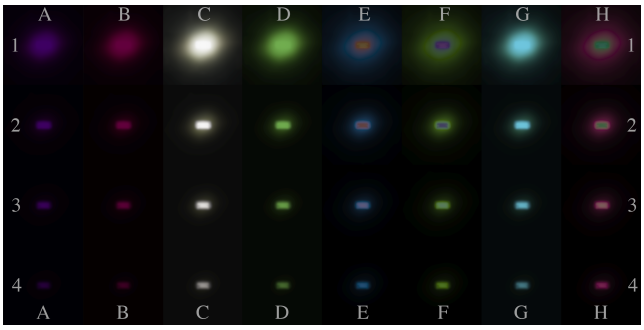

Fig. 15. Afterimage appearance candidates. Participants used this image to help describe the appearance of natural afterimages induced by a real camera flash. The images in column H were produced by our model with simulated flash brightnesses of 10^7, 10^5, 10^4, and 10^3 cd/m^2 (from top to bottom). Column C was produced similarly by Ritschel and Eisemann's model. The remaining images were produced by applying color transforms to either our or Ritschel's model. The image reproduced above has been brightened (relative to what participants saw) to improve visibility under office lighting conditions.

Stimuli: The stimulus for each trial was a single camera flash produced by a Canon Speedlite 430EX II flash unit (maximum guide number 43m) set to 1/16th power located 2m from the observer. Three conditions were tested: an unobstructed flash, a flash filtered through a 4x neutral density (ND) filter, and a flash filtered through a cumulative 32x ND filter. The afterimage options displayed on screen were still images of the initial afterimage appearance following a simulated camera flash. Four options were produced by our model, using different simulated flash brightnesses. Four options were similarly produced by Ritschel and Eisemann's model. The remaining options were produced by applying color transformations to the images produced by our and Ritschel's models. Figure 15 shows the onscreen options provided to subjects. All afterimage options were displayed on a Sony PVM-2541 OLED monitor. The experiment was performed in a completely dark room (excepting the monitor and flashes).

Procedure: Prior to starting the experiment, participants were shown the afterimage options on the computer screen and given a short explanation of what (natural) afterimages are and how they are produced. Each flash condition was presented as many times as the participant requested, starting with the brightest, unobstructed flash. Participants were told to pay attention to the initial appearance of the afterimage after the flash fired. They then verbally described the afterimage appearance in reference to the options shown on the computer screen. Between each flash condition, participants were given one minute to recover while the experimenter "adjusted" the flash.

4.3.2 *Results.* All ten participants agreed that afterimage appearance changed with flash brightness. All participants also reported afterimages with a strong color contrast between the very bright bulb of the flash and the dimmer glare surrounding it—for the dimmest flash, the glare rarely produced an afterimage, so no color contrast could be reported. There was less consensus about the afterimage colors themselves. Most participants liked to mix and match options from the palette in Figure 15, e.g., "a center like G3 with a fuzzy outside like B1", but each would use different components. Eight of the participants described colors identical or similar to those predicted by our model, with only slight variations, e.g., a more "purple" periphery than our predicted magenta, a more "light blue" center than our predicted cyan, etc.

Across flash conditions, seven of the participants reported a difference in hue. In a follow up session with one of the three remaining participants, a stronger flash (1/4th power) produced an afterimage of a different hue, so it is possible that the thresholds for afterimage hue transitions vary from person to person and were not necessarily captured by our three original test conditions. A wider range of flash brightnesses would likely result in an even larger proportion of people reporting afterimage hue transitions. As is, both the within and across condition results of this experiment are sufficient to support the conclusion that afterimage appearance, specifically color, is a function of the absolute luminance of the light source.

Following the experiment, we asked the participants to compare the realism of our afterimage model to that of Ritschel and Eisemann. All participants agreed that the lack of color contrast in Ritschel's afterimages made them less believable than those produced by our model. Interestingly, this includes five participants who, prior to the experiment, selected Ritschel's as the afterimage they expected to see. This reversal may provide some insight regarding the poor performance of afterimages as a brightness cue. A person without an accurate mental model of how an afterimage ought to appear probably does not pay close attention to them in the course of daily life and thus would be less likely to use them as a brightness cue.

5. CONCLUSIONS

In this paper we presented computational models for simulating the appearance of extreme illumination epiphenomena, including a new model for afterimages that distinguishes between those resulting from photoreceptor bleaching and from local adaptation. We also introduced a novel stochastic model for mesopic visual acuity, which demonstrates the potential of dynamic alternatives to effects formerly studied primarily in the context of static images. In order to determine the relationship between these epiphenomena and the perception of brightness, we built a display system that implements our models and then used it to perform psychophysical evaluations of the epiphenomena's effects. Ultimately, we discovered afterimages have little impact on the brightness percept. Mesopic rendering effects, on the other hand, seem to be a strong cue for brightness in naïve observers.

One issue that merits further discussion is the seemingly conflicting results of our study and the one presented by Ritschel and Eisemann [2012]. One possible explanation is that afterimages are indeed a cue, but perhaps are too weak a cue to be detected with the datapoints that we collect. Ritschel and Eisemann's main experiment collected data only for pairs of stimuli of identical intensity (and one augmented with an afterimage). This is more or less equivalent to a subset of trials in our experiment (the main experiment trials where the test intensity is zero). In our experiment, we collected 30 responses per participant for that particular test condition, while Ritschel and Eisemann collected 300 responses per participant. Another potential explanation is that the experimental design used by Ritschel and Eisemann left room for bias in the subject responses. Because the brightness of the stimuli was always identical, the subject was never actually making any real brightness judgments. Accordingly, it is possible that a naïve subject could have become aware of this and then misinterpreted the goal as a "find the afterimage" task rather than a true brightness comparison (which was impossible). Whether afterimages are a natural cue for brightness or not, viewers may be able to be trained to recognize

and understand them to still gain some benefit from their presence. It is also possible that the presence of synthetic afterimages adds to the general sense of immersion or realism in a display system, and thus should be considered for inclusion on this basis alone.

One limitation is that our system assumes that only one user at a time is looking at the display. Even if a multi-person eye tracker were available, the proposed method could not be extended as is to multiple viewers, as the rendition of the scene is tailored to the fixations of a specific user. This can be a limitation in some contexts, but it would not affect applications such as games, where the images displayed are typically consumed by a single player.

We end with a few suggestions based on our personal experience observing these effects. While repetitively flashing lights at ourselves to study the appearance of real afterimages, we informally observed that after blinking, or after a saccade, the perceived strength of the afterimage is temporarily boosted, perhaps because we are consciously attending to the afterimage. In future work we would like to further study this phenomenon and possibly integrate it in our system.

One cue that our system does not implement relates to the behavior of the fovea in scotopic conditions. The fovea contains no rods, so, in scotopic conditions, the fovea becomes a blind spot. This means that objects that are visible in the periphery of the field-of-view can disappear completely once they are foveated. We tried simulating this effect as a modification to our stochastic loss-of-acuity model, specifically by reducing the probability of high acuity in the fovea under mesopic conditions. However, contrary to all the other cues we implement, this phenomenon does require a higher spatial accuracy and faster framerate than our gaze tracker can offer. Therefore, we decided to remove the cue. As higher-quality gaze trackers become more commonly available, incorporating foveal loss of acuity could improve the experience of using our system to view low-light scenes.

Finally, in the general framework of a display that reacts and adjusts to the user's gaze, the perceived brightness of the monitor is only one of the axes that can be explored. Jacobs et al.'s [2014] recent work explored mobile touch interfaces for navigating additional dimensions, such as focus and white balance. The latter is a particularly interesting dimension as it would support the display of scenes lit by different illuminants, where a single white balance may be impossible to define. Shifting control of these dimensions to a natural gaze interface as well may help create an even more immersive and realistic viewing experience.

REFERENCES

BAYLOR, D. A., HODGKIN, A. L., AND LAMB, T. D. 1974. Reconstruction of the electrical responses of turtle cones to flashes and steps of light. *Journal of Physiology 242,* 3 (9).

BEDGGOOD, P. AND METHA, A. 2012. Variability in bleach kinetics and amount of photopigment between individual foveal cones. *Multidisciplinary Opthalmic Imaging 53,* 7 (6).

BUCK, S. L., KNIGHT, R., FOWLER, G., AND HUNT, B. 1998. Rod influence on hue-scaling functions. *Vision Research 38,* 21 (11).

BURT, P. AND ADELSON, E. 1983. The Laplacian pyramid as a compact image code. *IEEE Transactions on Communications 31,* 4 (4).

CAO, D., POKORNY, J., AND SMITH, V. C. 2005. Matching rod percepts with cone stimuli. *Vision Research 45,* 16 (7).

COOPER, E. A., JIANG, H., VILDAVSKI, V., FARRELL, J. E., AND NORCIA, A. M. 2013. Assessment of OLED displays for vision research. *Journal of Vision 13,* 12.

DOWLING, J. E. 1987. *The Retina: An Approachable Part of the Brain*. Belknap Press.

FAIN, G. L., MATTHEWS, H. R., CORNWALL, M. C., AND KOUTALOS, Y. 2001. Adaptation in vertebrate photoreceptors. *Physiological reviews 81,* 1.

FAIRCHILD, M. D. 2008. *The HDR Photographic Survey*. MDF Publications.

FERWERDA, J. A., PATTANAIK, S. N., SHIRLEY, P., AND GREENBERG, D. P. 1996. A model of visual adaptation for realistic image synthesis. In *Proceedings of ACM SIGGRAPH 1996*.

GUENTER, B., FINCH, M., DRUCKER, S., TAN, D., AND SNYDER, J. 2012. Foveated 3D graphics. *ACM Transactions on Graphics (Proc. ACM SIGGRAPH Asia 2012) 31,* 6 (11).

HODGKIN, A. L. AND NUNN, B. J. 1988. Control of light-sensitive current in salamander rods. *Journal of Physiology 403,* 1.

IRAWAN, P., FERWERDA, J. A., AND MARSCHNER, S. R. 2005. Perceptually based tone mapping of high dynamic range image streams. In *Proceedings of Eurographics 2005*.

ITO, H., OGAWA, M., AND SUNAGA, S. 2013. Evaluation of an organic light-emitting diode display for precise visual stimulation. *Journal of Vision 13,* 7.

JACOBS, D., GALLO, O., AND PULLI, K. 2014. Dynamic image stacks. *Computer Vision and Pattern Recognition Workshops (CVPRW), 2014 IEEE Conference on*.

JENSEN, H. W., PREMOZE, S., FERWERDA, J. A., STARK, M. M., AND THOMPSON, W. B. 2000. Night rendering. Tech. Rep. UUCS-00-016, University of Utah.

JUDD, D. B. 1933. The 1931 I. C. I. standard observer and coordinate system for colorimetry. *Journal of the Optical Society of America 23,* 10 (10).

KALLONIATIS, M. AND LUU, C. 2007. Light and dark adaptation. In *Webvision*. University of Utah. `http://webvision.med.utah.edu/`, accessed on January 1, 2015.

KIM, M. H., WEYRICH, T., AND KAUTZ, J. 2009. Modeling human color perception under extended luminance levels. *ACM Transactions on Graphics (Proc. SIGGRAPH 2009) 28,* 3 (7).

KIRK, A. G. AND O'BRIEN, J. F. 2011. Perceptually based tone mapping for low-light conditions. *ACM Transactions on Graphics (Proc. ACM SIGGRAPH 2011) 30,* 4 (7).

LUEBKE, D. AND HALLEN, B. 2001. Perceptually driven simplification for interactive rendering. In *Rendering Techniques 2001*. Springer.

MATTHEWS, H. R. 1996. Static and dynamic actions of cytoplasmic Ca^{2+} in the adaptation of responses to saturating flashes in salamander rods. *Journal of Physiology 490,* 1.

MCCANN, J. AND RIZZI, A. 2011. *The Art and Science of HDR Imaging*. Wiley.

MIKAMO, M., SLOMP, M., RAYTCHEV, B., TAMAKI, T., AND KANEDA, K. 2013. Perceptually inspired afterimage synthesis. *Computers & Graphics 37,* 4.

NERGER, J. L., VOLBRECHT, V. J., AND HAASE, K. A. 2003. The influence of rods on colour naming during dark adaptation. In *Normal and defective colour vision*, J. D. Mollon, J. Pokorny, and K. Knoblauch, Eds.

PADGHAM, C. 1968. Measurements of the colour sequences in positive visual after-images. *Vision Research 8,* 7.

PATTANAIK, S. N., TUMBLIN, J., YEE, H., AND GREENBERG, D. P. 2000. Time-dependent visual adaptation for fast realistic image display. In *Proceedings of ACM SIGGRAPH 2000*.

PHOTO RESEARCH INC. 2014. PR-715 SpectraScan spectroradiometer. `http://www.photoresearch.com/`, accessed on December 10, 2014.

POKORNY, J. AND CAO, D. 2010. Rod and cone contributions to mesopic vision. In *Proceedings of CIE 2010 "Lighting Quality and Energy Efficiency"*.

RAHARDJA, S., FARBIZ, F., MANDERS, C., ZHIYONG, H., LING, J. N. S., KHAN, I. R., PING, O. E., AND PENG, S. 2009. Eye HDR: Gaze-adaptive system for displaying high-dynamic-range images. In *ACM SIGGRAPH ASIA 2009 Art Gallery & Emerging Technologies*.

REINHARD, E., HEIDRICH, W., DEBEVEC, P., PATTANAIK, S., WARD, G., AND MYSZKOWSKI, K. 2010. *High Dynamic Range Imaging: Acquisition, Display, and Image-Based Lighting*. Morgan Kaufmann.

RIGGS, L. 1965. *Vision and Visual Perception*. Wiley, Chapter Visual acuity.

RITSCHEL, T. AND EISEMANN, E. 2012. A computational model of afterimages. *Computer Graphics Forum (Proc. Eurographics 2012) 31,* 2.

RITSCHEL, T., IHRKE, M., FRISVAD, J. R., COPPENS, J., MYSZKOWSKI, K., AND SEIDEL, H.-P. 2009. Temporal glare: Real-time dynamic simulation of the scattering in the human eye. In *Proceedings of Eurographics 2009*.

RUSHTON, W. AND HENRY, G. 1968. Bleaching and regeneration of cone pigments in man. *Vision Research 8,* 6.

SEETZEN, H., HEIDRICH, W., STUERZLINGER, W., WARD, G., WHITEHEAD, L., TRENTACOSTE, M., GHOSH, A., AND VOROZCOVS, A. 2004. High dynamic range display systems. *ACM Transactions on Graphics (Proc. ACM SIGGRAPH 2004) 23,* 3 (8).

SHIN, J. C., YAGUCHI, H., AND SHIOIRI, S. 2004. Change of color appearance in photopic, mesopic and scotopic vision. *Optical review 11,* 4.

THOMPSON, W. B., SHIRLEY, P., AND FERWERDA, J. A. 2002. A spatial post-processing algorithm for images of night scenes. *Journal of Graphics Tools 7,* 1 (11).

TOBII. 2014. Tobii Rex developer edition. `http://www.tobii.com/`, accessed on December 10, 2014.

TORRE, V., MATTHEWS, H. R., AND LAMB, T. D. 1986. Role of calcium in regulating the cyclic GMP cascade of phototransduction in retinal rods. *Proceedings of the National Academy of Sciences 83,* 18.

WADE, N. AND SWANSTON, M. 2013. *Visual Perception: An Introduction, 3rd Edition*. Psychology Press.

WANAT, R. AND MANTIUK, R. K. 2014. Simulating and compensating changes in appearance between day and night vision. *ACM Transactions on Graphics (Proc. ACM SIGGRAPH 2014) 33,* 4 (08).

WEVE, H. 1925. The colours of after-images, following strong light-stimuli. *British Journal of Opthalmology 9,* 12.

ACKNOWLEDGMENTS

We would like to thank Anthony Norcia for his assistance and support in designing our perceptual assessments as well as Martin Banks and Brian Wandell for their helpful suggestions and feedback. We would also like to acknowledge Tobias Ritschel and Michihiro Mikamo for providing source code for their afterimage models. Finally, we would like to thank our experiment participants and the anonymous reviewers for their invaluable contributions to this work.

APPENDIX

A. SUGGESTED PARAMETERS

We found the following parameters to give good results under typical office viewing conditions. Different choices of parameters can also exaggerate the appearance of our cues to create alternate percepts, e.g., darker or brighter than the scene's actual radiance suggests.

Symbol	Value(s)	Description
γ	0.455	Display gamma compensation
n	0.7	Naka-Rushton exponent
a_1	0.75	Adaptation rate to brighter stimuli
a_2	0.2	Adaptation rate to darker stimuli
b_1	1.2×10^{-3} 1.4×10^{-4} 5.5×10^{-7}	LMS photopigment bleaching rate
b_2	6.9×10^{-1} 1.4×10^{0} 5.5×10^{0}	LMS photopigment regeneration rate
c_1	1	Calcium influx rate
c_2	1	Calcium efflux rate
j_1	0.03	Bleaching afterimage strength
j_2	0.15	Adaptation afterimage strength
m_1	0	Maximum mesopic log-luminance
m_2	-2	Minimum mesopic log-luminance
w_P	0.2126 0.7152 0.0722	RGB photopic luminance weights [Judd 1933]
w_S	0.000 0.8451 0.1459	RGB scotopic luminance weights (Adapted from [Kirk and O'Brien 2011])
σ_i	$0.75\times2^{i/4}$	Gaussian kernel standard deviations ($i = 0..7$)
ρ_i	$0.286 \times i$	Acuity loss thresholds in log-luminance ($i = 0..7$)
g	0.25	Stochastic acuity variance
H	$\begin{bmatrix} 1.14 & -1.18 & 0.36 \\ -0.08 & 1.02 & 0.18 \\ 0.30 & -0.11 & 1.25 \end{bmatrix}$ LMS bleaching levels $\rightarrow$ RGB afterimage color matrix	