

Color-Perception-Guided Display Power Reduction for Virtual Reality

BUDMONDE DUINKHARJAV* and KENNETH CHEN*, New York University, USA

ABHISHEK TYAGI, University of Rochester, USA

JIAYI HE, University of Rochester, USA

YUHAO ZHU†, University of Rochester, USA

QI SUN†, New York University, USA

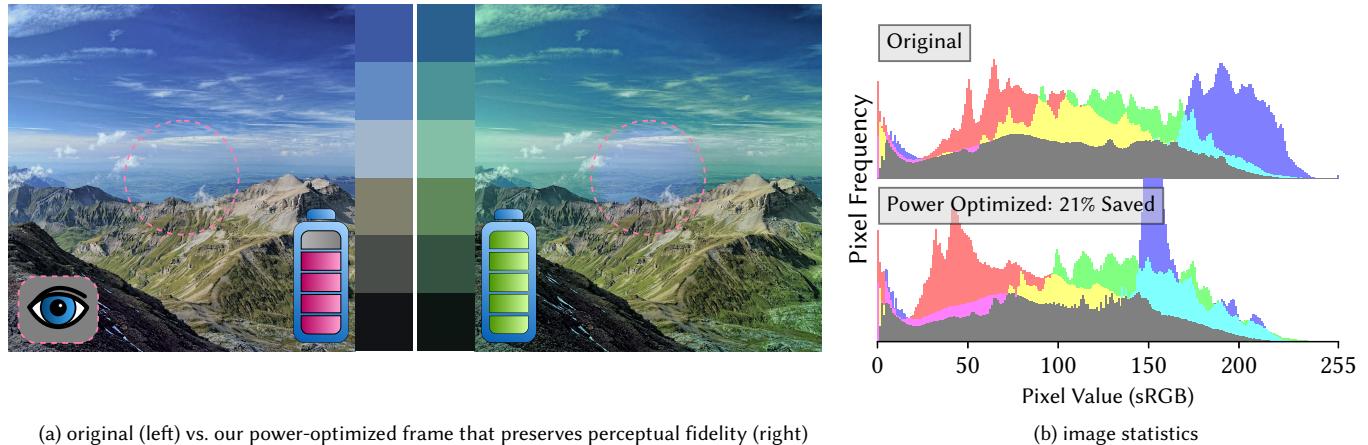


Fig. 1. Illustration of our gaze-contingent and color-perception-aware display power saving model. We present a perceptually-guided, real-time, and closed-form model for minimizing the power consumption of untethered VR displays while preserving visual fidelity. We apply a gaze-contingent shader onto the original frame ((a) left) to produce a more power efficient frame ((a) right) while preserving the luminance level and perceptual fidelity during active viewing. The dashed circles indicate the user's gaze. Our method is jointly motivated by prior literature revealing that: i) the power cost of displaying different colors on LEDs may vary significantly, even if the luminance levels remain unchanged [Dong and Zhong 2011]; ii) human color sensitivity decreases in peripheral [Hansen et al. 2009] and active vision [Cohen et al. 2020]. The color palette of the original frame is modulated using our peripheral filter and is shown in (a) for a visual comparison. While the color palettes appear different when gazed upon, an observer cannot discriminate between them when shown in their periphery. (b) visualizes how our model shifts the image's chromatic histograms to minimize the physically-measured power consumption. The blue LEDs consume more energy than the red/green in our experiment display panel. Image credits to Tim Caynes © 2012.

Battery life is an increasingly urgent challenge for today's untethered VR and AR devices. However, the power efficiency of head-mounted displays is naturally at odds with growing computational requirements driven by better resolution, refresh rate, and dynamic ranges, all of which reduce the sustained usage time of untethered AR/VR devices. For instance, the Oculus Quest 2, under a fully-charged battery, can sustain only 2 to 3 hours

*equal contribution.

†corresponding authors.

Authors' addresses: Budmonde Duinkharjav, budmonde@gmail.com; Kenneth Chen, kc4906@nyu.edu, New York University, USA; Abhishek Tyagi, atyagi2@ur.rochester.edu, University of Rochester, USA; Jiayi He, jhe36@u.rochester.edu, University of Rochester, USA; Yuhao Zhu, yzhu@rochester.edu, University of Rochester, USA; Qi Sun, qisun@nyu.edu, New York University, USA.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2022 Association for Computing Machinery.

0730-0301/2022/12-ART210 \$15.00

<https://doi.org/10.1145/3550454.3555473>

of operation time. Prior display power reduction techniques mostly target smartphone displays. Directly applying smartphone display power reduction techniques, however, degrades the visual perception in AR/VR with noticeable artifacts. For instance, the “power-saving mode” on smartphones uniformly lowers the pixel luminance across the display and, as a result, presents an overall darkened visual perception to users if directly applied to VR content.

Our key insight is that VR display power reduction must be cognizant of the gaze-contingent nature of high field-of-view VR displays. To that end, we present a gaze-contingent system that, without degrading luminance, minimizes the display power consumption while preserving high visual fidelity when users actively view immersive video sequences. This is enabled by constructing 1) a gaze-contingent color discrimination model through psychophysical studies, and 2) a display power model (with respect to pixel color) through real-device measurements. Critically, due to the careful design decisions made in constructing the two models, our algorithm is cast as a constrained optimization problem with a *closed-form* solution, which can be implemented as a real-time, image-space shader. We evaluate our system using a series of psychophysical studies and large-scale analyses on natural images. Experiment results show that our system reduces the display power by as much as 24% (14% on average) with little to no perceptual fidelity degradation.

CCS Concepts: • Computing methodologies → Perception; Virtual reality; Mixed / augmented reality.

Additional Key Words and Phrases: Visual Perception, VR/AR, Color Perception, Power Consumption, Gaze-Contingent Rendering

ACM Reference Format:

Budmonde Duinkharjav, Kenneth Chen, Abhishek Tyagi, Jiayi He, Yuhao Zhu, and Qi Sun. 2022. Color-Perception-Guided Display Power Reduction for Virtual Reality. *ACM Trans. Graph.* 41, 6, Article 210 (December 2022), 16 pages. <https://doi.org/10.1145/3550454.3555473>

1 INTRODUCTION

Virtual and Augmented Reality (VR and AR) devices are increasingly becoming untethered for portability, outdoor usage, and unrestricted locomotion to enable ultimate immersion. At the same time, the demands for higher resolution, framerate, and dynamic range are steadily increasing, which is directly at odds with the limited energy capacity of untethered AR/VR devices. For example, when fully charged, both the Oculus Quest 2 and Hololens 2 can actively run only for 2-3 hours [Microsoft 2021]. Since the total energy capacity increases only marginally because “there is no Moore’s law for batteries” [Schlachter 2013], power consumption has become a primary concern in the design process of AR/VR devices [Debattista et al. 2018; Wang et al. 2016; Zhang et al. 2021b].

This paper specifically focuses on reducing the display power. In our measurement of HTC Vive Pro Eye and Oculus Quest 2, the display consumes as much as half of the total power consumption by comparing the power when the display is on vs. off. The results are consistent with data reported in other measurement studies [Leng et al. 2019; Yan et al. 2018]. Display power will only become more important in the cloud rendering paradigm, where the computation is offloaded to the cloud, heightening the contribution of display to the total device power.

Conventional display power optimizations are geared toward smartphones, which, when directly applied to VR devices, lead to significant visual quality degradation. This is because smartphone display optimizations are fundamentally gaze-agnostic, rightly so because smartphone displays have very narrow field-of-view. These optimizations either modulate pixels *uniformly* across the display [Shye et al. 2009; Yan et al. 2018] or are purely based on the content (e.g., UI elements) [Dong et al. 2009; Dong and Zhong 2011; Ranganathan et al. 2006]. Classic gaze-contingent optimizations in AR/VR such as foveated rendering, while reducing the rendering load [Krajancich et al. 2021; Patney et al. 2016], do not (directly) reduce the display power.

We present a gaze-contingent rendering approach that reduces the power consumption of untethered VR displays by as much as 24% while preserving visual quality during active viewing. We achieve this by only modulating the chromaticity of the display output without changing luminance.

Our method is jointly motivated by hardware research that revealed the variation of power consumption of displaying different colors on LEDs [Dong and Zhong 2011], as well as the recently discovered limitations of human peripheral color perception during active vision [Cohen et al. 2020]. That is, given an original frame such as in a 360 video, we seek a computational model that guides a

gaze-contingent color “shift” that (1) requires the minimal power cost, and (2) preserves the perceived fidelity.

To accomplish this, we conducted two pilot studies. First, we quantitatively model how our color sensitivity degrades with higher retinal eccentricities. Second, we physically measure the LED display power consumption as a function of the displayed color. Given the perceptual and the power model, our system performs a constrained optimization that identifies, for each pixel, an alternative color that minimizes the power consumption while maintaining the same perceptual quality. Critically, the optimization problem has a *closed-form* solution because of the judicious design decisions we made in constructing the perceptual and power models. As a result, our perception-preserving color modulation can be implemented as a real-time shader.

We validate our method with both subjective studies on panoramic videos, as well as an objective analysis on large-scale natural image data. We demonstrate the model’s effectiveness in display power reduction and perceptual fidelity preservation, relative to an alternative luminance-based “power saver”. Our objective analysis concludes that this model shows generalizability to a large variety of natural scenes and save, on average, 14% power. In summary, our main contributions include:

- a psychophysical study and data that measure human color discrimination sensitivity at various retinal eccentricities and reference colors in the Derrington-Krauskopf-Lennie (DKL) color-space;
- a physical system and data that measure the power consumption of VR-alike stereoscopic displays as a function of displayed color;
- a closed-form formulation that suggests the optimal (in terms of lowest display power cost without compromised visual fidelity) per-pixel chromaticity modulation by leveraging two learned models from our two aforementioned datasets;
- a real-time shader for gaze-tracked VR headsets and natural content viewing applications, as well as a demonstration of its general benefits with a large-scale analysis.

We provide the source code for our model regression and shader implementations at www.github.com/NYU-ICL/vr-power-saver.

2 RELATED WORK

2.1 Energy-Aware Graphics and Display

The graphics rendering pipeline requires heavy computation to execute in real-time. Reliably maintaining the performance requirements of these applications consumes considerable power. As such, energy-aware methods have been developed to minimize power while maintaining rendering quality. Most prior work in the graphics literature, however, focuses on reducing the rendering power [Debattista et al. 2018; Wang et al. 2016; Zhang et al. 2018, 2021b].

Complementary to prior work on reducing the rendering power, our work reduces the *display* power – by modulating the display color while preserving perceptual fidelity. The relationship between display power and color is studied in the mobile computing community, mostly in the context of smartphones. Dash and Hu [2021], Dong et al. [2009], and Dong and Zhong [2011] model the power consumption of smartphone displays with respect to color. Miller

et al. [2006] discusses how various hardware design optimizations can affect the power modeling accuracy.

Other work focuses on reducing the display power by modulating the brightness/luminance of the display rather than color, which our work focuses on. For instance, Shye et al. [2009] gradually dims the smartphone display after a user stares at the display for an extended period of time. Yan et al. [2018] pushes the empirical approach in Shye et al. [2009] a step further, and uses physiological data to derive a quantitative method to adaptively reduce the display luminance. Our work on color modulation is orthogonal and complementary to luminance-modulating techniques. We show that significant power saving is readily obtainable by adjusting only color; combining color and luminance modulation would conceivably lead to higher power savings, which we leave to future work (see Section 8).

Another orthogonal line of work is to reduce OLED power via better hardware design. Shin et al. [2013] propose to dynamically scale the display supply voltage, coupled with image-space color transformation, to minimize the display power while maintaining perceptual color similarity (as quantified by the CIELAB ΔE^* metric). Boroson et al. [2009], Miller et al. [2008] and Miller et al. [2007] propose four-color OLED structures, where a fourth sub-pixel has a higher power efficiency (luminous efficacy) than one of the RGB sub-pixels. The fourth, higher-efficiency sub-pixel allows many colors that usually dominate an image (e.g., neutral or saturated colors) to be produced with less contribution from the blue and red sub-pixels, which have low power efficiency. These prior works, however, do not consider the gaze-contingent color perception in VR.

2.2 Perceptually-Aware Immersive Rendering

Perceptual graphics studies hinge on the idea that the human visual system's ability to receive and process light signals deteriorates for stimuli located at higher retinal eccentricities, due to the significantly denser distribution of retinal cone cells in the fovea [Song et al. 2011]. Studies have taken advantage of this fact to create perceptually-aware models of rendering imagery [Duchowski et al. 2005]. Foveated rendering takes advantage of the drop in visual acuity as retinal eccentricity increases by reducing image quality in the periphery, ultimately improving performance while maintaining perceptual quality [Guenter et al. 2012; Patney et al. 2016; Sun et al. 2017; Walton et al. 2021]. Similar perceptually-aware gaze-contingent algorithms have been proposed to save spatio-temporal data bandwidth [Kaplanyan et al. 2019; Krajancich et al. 2021]. However, other than modulating user-perceived color [Mauderer et al. 2016] or remapping peripheral luminance and color for power saving [Li et al. 2022], there has been little work leveraging gaze-contingent color perception to optimize real-time rendering. That could be mainly because the modulated color does not introduce performance acceleration within the graphics pipeline.

2.3 Visual Perception of Displayed Color

Color and the corresponding human perception is a long-standing and extensive research topic in computer graphics [Rhyne et al. 2018]. Here, we discuss prior research that understands perceptual effects on display color.

First, perception of color is studied from psychophysical and even physiological perspectives. Early color perception research was largely based on psychophysical color matching experiments, which resulted in the *CIE 1931 RGB* [Guild 1931; Wright 1929] and *CIE 1931 XYZ* color spaces [Fairman et al. 1997]. The *CIE XYZ* color space has since become the cornerstone of modern color research, because it presents a hardware/device-independent way of quantifying colors. In fact, the first measurements of human color-discriminative thresholds, widely known as the MacAdam ellipses were presented in the *xy* chromaticity plot [MacAdam 1942].

The *sRGB* color space widely used in graphics, display, and vision fields today is derived from the *XYZ* color space. The *sRGB* color space, however, is rarely directly used in color perception studies, because *sRGB* is a device-dependent color space with non-uniform quantization and has a gamut smaller than the gamut of the human visual system. Color perception studies usually operate on some form of physiologically-based *LMS* color spaces, which are developed from measurements of the surgically removed cone cells [Bowmaker and Dartnall 1980; Dartnall et al. 1983] or from color defective vision [Stockman and Sharpe 2000]. Cone color spaces provide the ability to make “feed-forward” reasoning about our visual perception system.

One such physiologically-based study of color-discriminative thresholds was established in [Krauskopf and Karl 1992]. Similar to the color-opponent theory, upon which CIELAB was based on [Schiller and Logothetis 1990], neural recordings within the visual cortex confirmed the existence of a cone-opponent mechanism where signals from the *L*, *M*, and *S* cones are compared in a feed-forward fashion [De Valois et al. 1966]. The *DKL* color space [Derrington et al. 1984] leverages this cone-opponent mechanism to derive a physiologically-relevant and perceptually uniform color space. Our perceptual study is based on the *DKL* color space.

Second, human color perception exhibits eccentricity effects. Similar to our resolution and depth [Sun et al. 2020] acuity, color sensitivity also decreases as retinal eccentricity increases [Cohen et al. 2020; Hansen et al. 2008, 2009]. For instance, given a reference color, our discrimination thresholds are observed to have the shape of an ellipse in *DKL* space, and this region of sub-threshold colors increases significantly as the retinal eccentricity of stimuli increase ($\sim 4.5\times$ larger ellipse radii at 50° compared to 5°) [Hansen et al. 2009]. In AR, the interference between virtual and physical colors raises new challenges in correctly aligning perceived color [Hassani and Murdoch 2016; Murdoch et al. 2015; Zhang et al. 2021a]. However, to our knowledge, there is no computational model that numerically predicts the sensitivity given an eccentricity and reference color.

Third, chromaticity and luminance have different temporal roles and sensitivities [Hermann et al. 2021]. The eccentricity effects of luminance contrast have been leveraged to perform gaze-contingent rendering [Tursun et al. 2019] and measure spatio-temporal video quality, compared with a reference [Mantiuk et al. 2021]. Chromaticity-based sensitivity has been leveraged to encode high dynamic range displays [Kim et al. 2021]. However, perhaps because modulating colors does not play a role in accelerating performance, the eccentricity effect has not been investigated for advancing VR/AR systems.

Finally, our color sensitivity is also correlated to human status and task nature. For instance, the color sensitivity during fixation

shifting (a.k.a. saccade) uniformly and significantly decreases [Braun et al. 2017]. Prior literature showed notably lower sensitivity in discrimination than detection tasks [Vingrys and Mahon 1998]. More recently, studies by Cohen et al. [2020] revealed the remarkably further decreased color sensitivity during active and natural viewing tasks: in those conditions, even desaturating peripheral images are imperceptible by viewers. Those studies all exhibit the complex nature of color perception, which is still undergoing active scientific discoveries. In this research, we focus on *active viewing* scenarios, which are representative conditions of VR/AR applications such as gaming and video-watching.

3 PILOT STUDY: ECCENTRICITY EFFECTS ON COLOR PERCEPTION

We aim to exploit how human perception of color varies across the visual field, so that we can adjust the appearance of visual stimuli in our peripheral vision in an advantageous way. Hansen et al. [2009] showed that while our ability to discriminate colors significantly deteriorates at high retinal eccentricities, we still maintain some ability to discriminate colors at eccentricities as high as 45° . Drawing inspiration from this work, we designed and performed a psychophysical study on the perceptual *discrimination* thresholds of colors, given various reference colors (5 total) and retinal eccentricities (from 10° to 35°). The experimental data later transforms to a computational model in Section 5.1.

Setup. We perform our study with the HTC Vive Pro Eye head-mounted display as shown in Figure 2a. Participants remained seated during the duration of the study, and interacted with the user study software via the keyboard.

Participants. We recruited 5 participants (ages 20-32, 2 female) for a series of four-alternative forced choice (4AFC) staircase experiments (similar to [Hansen et al. 2009]) to determine the discrimination thresholds. All participants had normal or corrected-to-normal vision and exhibited no color perception deficits as tested by the Ishihara pseudo-isochromatic plates. In this pilot study, we chose 5 participants due to the long duration of our staircase experiment. This is also practiced for similar threshold-determination psychophysical experiments [Krajancich et al. 2021; Sun et al. 2020]. All experiments were approved by an ethics committee and all participants' data was de-identified.

Stimuli. As shown in Figure 2a, the stimuli were four colored disks (with a diameter of 5 degrees). They were rendered simultaneously on top of a neutral gray background (i.e., $[0.5, 0.5, 0.5]$ in the linear sRGB space, or 71.5 cd/m^2). The azimuth position of the disks remained constant throughout the entire study, located at 45° , 135° , 225° , and 315° (i.e. the four diagonals in the participant's visual field), while the radial position (i.e. the retinal eccentricity) varied across sequences to be either 10° , 25° , or 35° . Three of the disks have the same "reference" color, and the fourth has a "calibration" color which changes throughout a sequence of trials. The space of colors that the disks can obtain is visualized as a color-space in Figure 2c. The luminance of all disks is maintained at the same level as the background's luminance.

Tasks. The task was an 1-up-2-down 4AFC staircase procedure. Specifically, the study was conducted in a single session split into 60 sequences ($= 5$ reference colors $\times 3$ eccentricities $\times 4$ color space dimensions, as specified below) of trials. Each sequence was a staircase procedure meant to determine a participant's discriminative threshold of the calibration disk, which was colored with subtly different chromaticities. Each sequence may contain a variable number of trials depending on the staircase convergence.

Each trial was a single 4AFC task where the participants were instructed to identify which one of the 4 color disks appeared different. The participant was instructed to fix their gaze on a white crosshair at the center of the screen for the duration that the stimuli were shown. We used eye tracking to ensure participants maintain their gaze at the central crosshair. We automatically rejected a trial if the user's gaze moves beyond 3° eccentricity, randomized the trial order again, and notified them. At the start of each trial of a sequence, we shuffle the four colored disks, and display them for 500ms (the same stimulus duration used in prior color discrimination literature [Hansen et al. 2009]). Once the stimuli disappear, we prompt the participant to identify and select the disk with the calibration color, using the keyboard. Depending on their answer, the calibration disk's color was made easier or harder to discriminate in the subsequent trial by adjusting the chromaticity of the calibration disk to be closer/farther from the reference color while preserving its luminance. After 6 reversals of this staircase procedure (or a maximum of 50 trials), the sequence terminates, and the next sequence begins. We visualize the progression of an example staircase-procedure in Figure 2b.

Across the sequences, we present 5 different reference colors, as visualized with black crosses in Figure 2c, each presented at 10° , 25° , and 35° retinal eccentricities. For each reference color, we adjust the color from four directions along the two equi-luminant cardinal axes in the *DKL* color-space. Briefly, *DKL* is a perceptually uniform color space that is conducive to color vision research [Hansen et al. 2008, 2009; Krauskopf and Karl 1992]. Please refer to Section 5.1 for a detailed overview of the *DKL* color space and why it is used in our perceptual studies.

The entire study took approximately 1.5 hours for each participant and they were encouraged to take breaks in between sequences. At the beginning of each user study, the participants completed 1 sequence to familiarize with the procedure and equipment.

3.1 Results

In total, 8,123 trials were obtained from our participants (5 participants each with 60 sequences consisting of ≈ 21 trials each on average). We record the color values at each reversal in *DKL* coordinates, and average the last 3 reversals (out of 6 total) to determine the final discrimination threshold for each participant. The average thresholds across all participants are visualized in Figure 2c in red, along with the 75% confidence interval error bars. As we approach the reference color from four directions in *DKL* space, we obtain four different thresholds for each color at each eccentricity. The lines connecting the four thresholds do not represent the shape of the overall threshold, and is only served as a visual guide to group each set of thresholds together. To avoid visual clutter, we

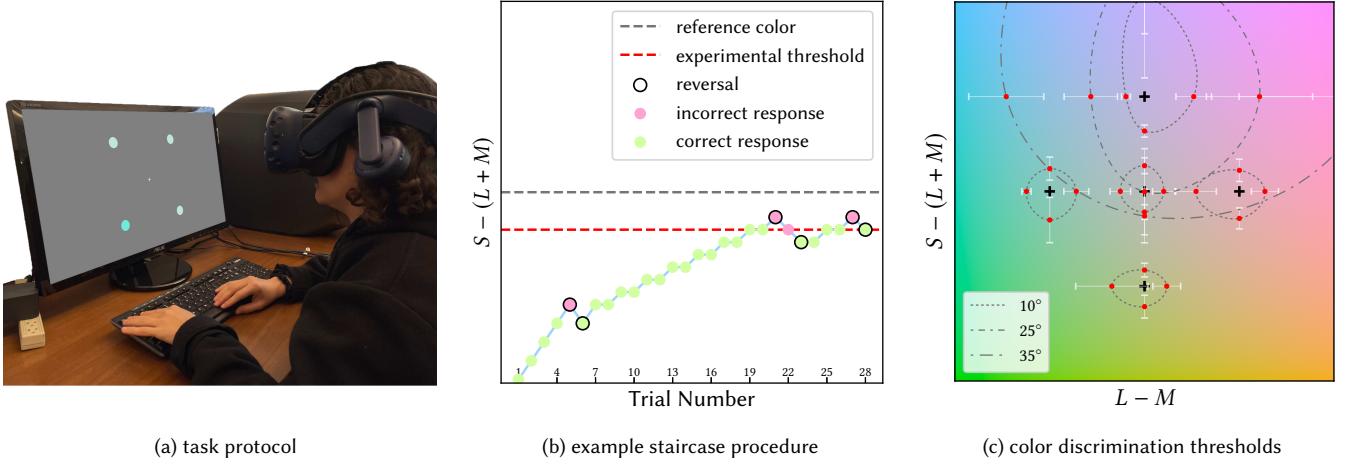


Fig. 2. *Pilot study for color discrimination threshold.* The $S - (L + M)$ and $L - M$ axes are axes in the *DKL* color-space, which compares the difference between S vs $L + M$ and L vs M cone activations [Derrington et al. 1984]. (a) Participants were instructed to fix their gaze at the center of a VR headset display and to distinguish the “calibration” disk out of the four shown in their periphery. The other three “reference” disks were colored with the reference color, while the calibration disk’s color is adjusted via a 1-up-2-down staircase procedure by the participant. (b) The staircase procedure aims to narrow down the peripheral color-discrimination threshold of the participant. As the participant completes a series of 4AFC tasks as depicted in (a), the staircase procedure automatically adjusts the color of the calibration disk to make it harder if the participant identified the calibration disk correctly and vice-versa. We continue each trial until a total of 6 reversals (denoted by outlined points) occur or 50 trials are completed, whichever occurs first. In this example, the experiment continues until the threshold in the $S - (L + M)$ axis is converged. (c) The black crosses indicate our 5 sampled reference/pedestal colors in *DKL* space. The experimental thresholds are displayed with red dots as mean, and white bars as 75% confidence intervals. For the other two sampled eccentricities ($25^\circ/35^\circ$), we only plot the results from one reference color, at 12 o’clock in the figure, to avoid visual clutter. We plot splines (dashed gray lines) for ease of visualization only.

plot discrimination thresholds at 10° eccentricity for each reference color and 10° , 25° , and 35° eccentricity thresholds for one reference color. Refer to Supplement B for all the measured threshold values separated by each participant.

3.2 Discussions

We chose the *DKL* color-space as our color sampling space primarily because prior work suggests that the *DKL* color space yields a perceptually uniform sampling space and, for that reason, is used extensively in perceptual studies [Hansen et al. 2008, 2009; Krauskopf and Karl 1992]. For our work, we only sampled colors on a single equi-luminant plane. In *DKL* space, this corresponds to keeping the third dimension of the color space constant. First, we observed unequal thresholds with different reference colors even if they were displayed at the same eccentricity. However, they all appeared an ellipse shape, as also evidenced by prior literature [Krauskopf and Karl 1992]. That motivates us to develop our computational perceptual model considering the reference color as one of the inputs.

Prior work which utilizes the *DKL* color-space suggests that discriminative thresholds measured with respect to a specific adaption luminance can be extended to arbitrary adaptation luminances due to the linearity of the cone-opponent process [Larimer et al. 1974, 1975]. We use these results in this research and only conducted discriminative threshold measurements at a single adaptation luminance of 71.5 cd/m^2 as mentioned above.

In the scope of our work, we did not study how spatial frequencies of stimuli affect discriminative thresholds. Our experimental data

provides the thresholds for a stimulus with a dominant frequency equal to 0.2 cpd corresponding to the stimulus size used throughout the experiment.

Unsurprisingly, our data shows a decrease of ability to discriminate chromatic discrepancies as the retinal eccentricity increases. The trend agrees with past experiments [Hansen et al. 2009], and is intuitive given the higher density of retinal receptors in the fovea [Song et al. 2011]. Figure 2c shows that the fall-off of discriminative sensitivity is very sharp, and the region of sub-threshold chromaticities at 35° can take up as much as a third of the observable hues. Some participants noted that at high eccentricities, all four disks appeared to be different, even though three of the disks were colored identically. As such, the amount of noisy thresholds at high eccentricities attribute to the larger uncertainty for the overall threshold measurements as shown in Figure 2c. Further investigations into this surprising phenomenon is an interesting future work.

We also observe inter-subject variation in the measured thresholds, as shown in Supplement B. While this could be due to a number of reasons (e.g., observer metamerism [Xie et al. 2020], prereceptoral filtering [Norren and Vos 1974], calibration, experimental setup, etc.), further study is required to understand the reason for these differences. Nevertheless, for developing a computational model, we use the most conservative thresholds across participants, instead of an average fit. This assures generalization to a larger population considering individual variances (see Section 6.1).

Lastly, it is notably critical that those thresholds only hold for discriminative tasks. Using the observed thresholds, we performed

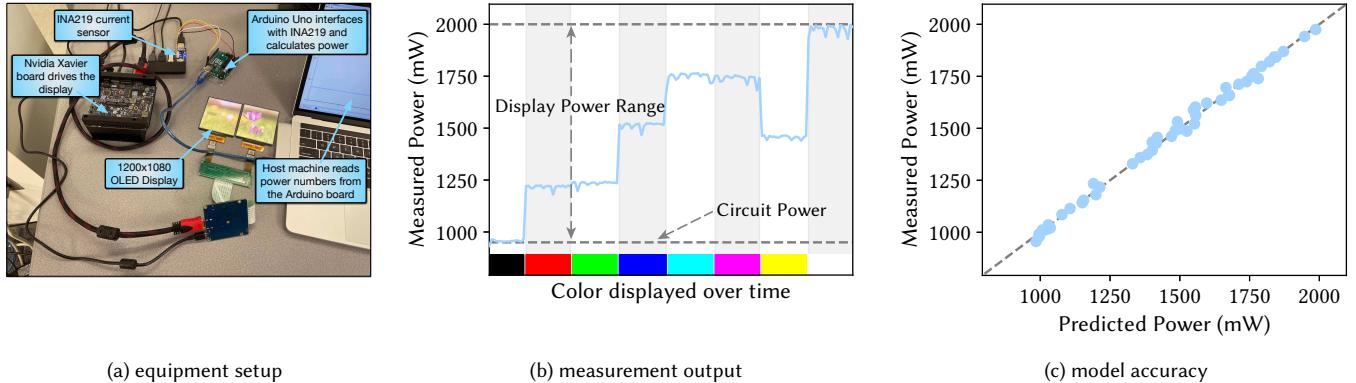


Fig. 3. *Display power measurement and modeling.* (a) The display is connected to an Nvidia Xavier board, which provides power and frame data. The display power is intercepted by a Texas Instruments INA219 current sensor. The current and voltage readings are transmitted to an Arduino Uno board, which calculates the power. A host laptop obtains the power consumption from Arduino Uno through the serial port. (b) The voltage and current readings are multiplied at each timestep to compute the power reading time-series. To measure the power for displaying different colors, we sample the *sRGB* color space; each color is displayed for a period of 5 seconds. Here we show the power trace when cycling through the eight vertices of the *sRGB* color cube. (c) We regress a linear power model by randomly sampling 52 colors in the *sRGB* color space, and plot the measured power (*x*-axis) vs. the estimated power (*y*-axis). The mean relative error of the regression is 0.996%, indicating an accurate power model. The dashed line indicates the line of perfect agreement (i.e. $y = x$).

a preliminary validation with a *sequential detection task* and two-alternative-forced choice (2AFC). In this study, the same group of participants was instructed to observe pairs of stimuli and identify whether they appear identical. Some of the trials consist of one non-altered image, with the other containing peripheral color altering within the identified thresholds. We observed that a majority of users can successfully identify the altered condition, suggesting the distinct perceptual thresholds between discrimination and detection tasks. Nevertheless, during active vision tasks where an observer is instructed to freely observe natural visual content, their sensitivity may significantly reduce [Cohen et al. 2020]. We hypothesize that the color sensitivity during active vision is also lower than during discriminative tasks. We investigate and validate the hypothesis in more detail in Section 7.1.

4 PILOT STUDY: MEASURING DISPLAY POWER WITH VARIED COLORS

To measure the power consumption characteristics of VR displays and how it varies depending on the images displayed on them, we conduct a hardware study, and later use the collected data to derive a model for predicting the power consumption of a display given the image displayed on it.

4.1 Setup

For our power study, we use the Wisecoco H381DLN01.0 OLED [Wisecoco 2022]. The display module has two identical displays, each with a resolution of 1080×1200, matching the aspect ratio of HTC Vive Pro Eyes, which is what we use for perceptual studies.

We do not use the native display modules in Vive Pro Eye HMD and Oculus Quest 2 for power studies, because their displays are physically tightly integrated into the headsets; thus, the display power cannot be easily isolated from the rest of the system. In the

case of the Oculus Quest 2, the headset is powered by a battery that is tightly integrated into the headset, which prevents us from using methods used in studying smartphone display power, where the battery is unplugged and replaced with an external power supply that has internal power sensing capabilities [Dash and Hu 2021; Dong et al. 2009; Halpern et al. 2016].

Figure 3a shows the experimental setup to measure display power. We intercept the display power supply with a SwitchDoc Power-Central board, which has an on-board INA219 module (with a 0.1Ω shunt resistor) to measure the current. The INA219 module is connected to an Arduino board through the I₂C interface. We develop a driver that runs on the Arduino board to get the display current and voltage, from which we can calculate the power.

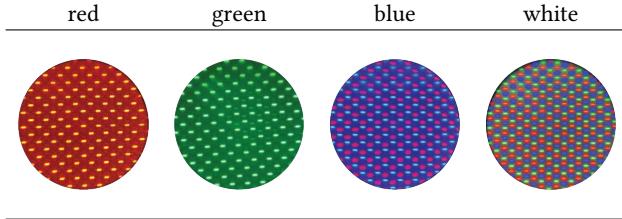
The driver running on the Arduino board configures the INA219 sensor to output a new power measurement every ~ 68 ms; each power reading is internally averaged over 128 samples, resulting in an effective power sampling rate of $\sim 1,882$ Hz.

4.2 Measurement and Discussion

As a preliminary test, we measure the power consumption of the eight vertices of the *sRGB* color cube. For each color, we set all the display pixels to that color, display it for five seconds, and calculate the average power. Figure 3b shows the measured power trace. It is clear that the display power consumption is sensitive to the color.

We make two observations from Figure 3b. First, even when the display is showing black pixels, i.e., when the LEDs are not emitting light, there is a non-trivial amount of *static* power consumption. The power beyond the static portion is consumed by the LEDs, which we dub the *dynamic* display power. This static power is consumed by the peripheral circuitry that drives the LEDs, such as the per-pixel transistors and capacitor as well as the addressing logic [Huang

Table 1. Our microscopic photos of the display under different colors. We image the display under sRGB red, green, blue, and white colors using a Carson MicroFlip microscope with a magnification of 120x. One can observe that the display red and green primaries roughly match the red and green primaries in the sRGB color space, but the sRGB blue requires contributions from both the blue and red sub-pixels from the display.



et al. 2020]. The contribution of the static power is about 50% in display white and is about 80% when displaying red and green.

The trend of semiconductor technology is that the circuit power is decreasing over time with better fabrication technologies [Bohr 2007], but the LED power is much harder to reduce because the display must sustain certain luminance levels to meet brightness requirements, which arguably do not change dramatically over time. Our work aims to reduce the (color-sensitive) dynamic power of the display, which will become more important as the static power reduces in the future.

Second, the dynamic power consumption of red and green colors are roughly half that of blue. This is because displaying the sRGB blue on our display requires contributions from both the blue and red sub-pixels (due to the primaries used by this display) as confirmed by examining the microscopic images of the display (Table 1). As a result, if we expect to see any energy wins, we anticipate that green-, and/or, red-shifting images can decrease the power consumption of the image. We will leverage the measured data to obtain a computational power-vs-color model in Section 5.2.

5 MODEL: PERCEPTUALLY GUIDED POWER OPTIMIZATION

Using the results of our perceptual user study, and hardware power measurements, we develop a display power optimization model under the constraint that the change in the images observed by human subjects is not perceptible. In Section 5.1, we first derive a computational model of human color discrimination (Figure 4) using the data obtained from Section 3. In Section 5.2, we build a linear power consumption model regressed from the physical measurement data in Section 4. Finally, in Section 5.3, we integrate the two models above (as a constrained convex optimization) toward a closed-form display color modulation function. It aims to minimize the display’s power consumption while ensuring the modulation within the human discriminative thresholds.

5.1 Perceptual Model for Color Discrimination

The study of human color vision suggests that the physiological mechanisms which govern our discriminative ability of colors rely on comparing the different cone receptor cell activations in our retinas [De Valois et al. 1966]. Specifically, the two primary comparison

mechanisms are the difference between (1) the L vs M cone and (2) the $L + M$ vs S cone activations [Derrington et al. 1984]. These observations are the foundation of the *cone opponent* theory of color discrimination and are widely accepted for measuring human color discriminative thresholds [Conway et al. 2018].

In this research, we leverage these results from the vision science community and develop a computational framework that quantifies the discriminative threshold of any given color at different retinal eccentricities. That is, given a colored stimulus at some retinal eccentricity, we determine the extent to which we can modulate its color while maintaining its perceptual color appearance to an observer.

The set of colors which are indistinguishable from some *test* color by human observers are modeled as ellipse shaped regions defined over equi-luminant color-spaces [Hansen et al. 2008, 2009; Krauskopf and Karl 1992]. Notably, the MacAdam ellipses [1942] are the first to model discriminative thresholds as such. Additionally, Krauskopf and Karl et al. [1992] show that the sizes of these ellipses are best described in the *DKL* color-space [Derrington et al. 1984]. We first introduce the *DKL* space, followed by how we model the discriminative threshold in the *DKL* space.

DKL color space. *DKL* space defines colors in two steps. First, it quantifies a color using the cone-opponent mechanism. Therefore, a color is first defined over a basis that computes the cone-opponent neural activations. Namely, a color with *LMS* coordinates (t_L, t_M, t_S) is first converted to the basis:

$$\begin{aligned} t_{L-M} &= t_L - t_M \\ t_{S-(L+M)} &= t_S - (t_L + t_M) \\ t_{L+M} &= t_L + t_M, \end{aligned} \quad (1)$$

Second, instead of directly using the cone-opponent basis, the *DKL* space models colors in terms of *color contrast*; that is, colors are defined relative to a reference (a.k.a., the *adaptation*) color. This is different from the more absolute measure of color used in *sRGB*, *XYZ*, and *LMS* color spaces, where each color is defined based on its own characteristics. Specifically, given a *test* color, \mathbf{t} , and an *adaptation* color, \mathbf{b} , we can compute the color-contrast of the test color with respect to the adaptation color as:

$$\kappa(t_i; b_i) = \frac{t_i - b_i}{b_i}, \text{ for } i \in \{1, 2, 3\}. \quad (2)$$

Depending on the specific basis used to represent these colors, the indices $\{1, 2, 3\}$ could be $\{X, Y, Z\}$ for the *XYZ* basis or $\{L, M, S\}$ for the *LMS* basis. In our case, the indices $\{1, 2, 3\}$ are $\{L - M, S - (L + M), L + M\}$, since we have already defined colors in the cone-opponent space, as shown in Equation (1).

It is worth noting Equation (1) can be seen as producing an intermediate color space that is a linear transformation away from the conventional *LMS* space. We will henceforth refer to the intermediate color space given by Equation (1) as *i*(n)termediate-*DKL*.

We use the *LMS* color space as defined by Smith and Pokorny [1975], which is what the original *DKL* space is based on [Derrington et al. 1984]. The particular *LMS* cone fundamentals are so defined that the coordinate t_{L+M} of a color is strictly equal to the luminance of the color, i.e., the Y coordinate in the *XYZ* space.

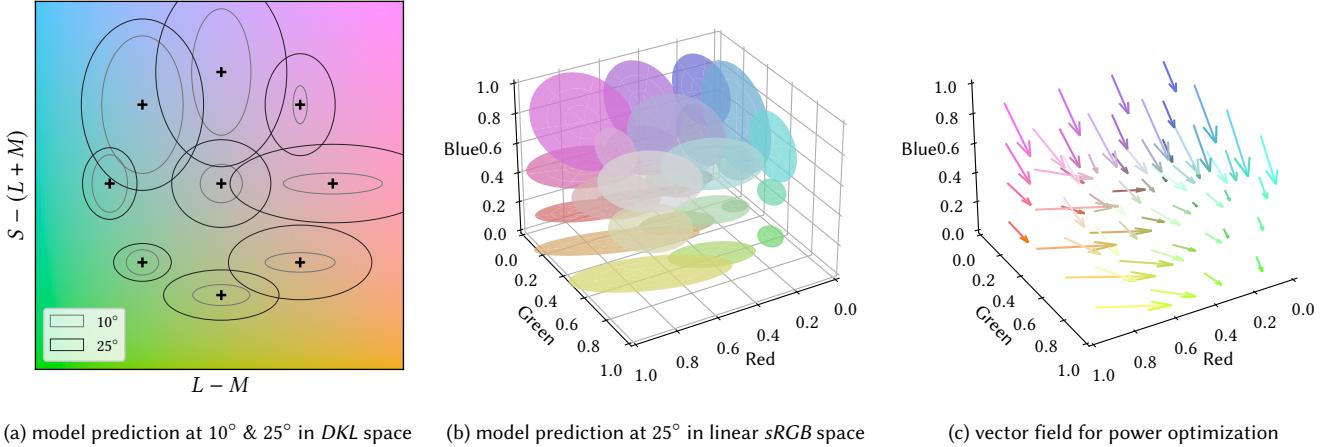


Fig. 4. *Color perception model and power-aware chromaticity optimization.* (a) illustrates our color discrimination model’s ellipse thresholds at nine coordinates. They are evaluated on an equi-luminant plane in the *DKL* color space, with eccentricities 10° (light gray curve) and 25° (dark black curve). All colors within an ellipse are perceptually indistinguishable from the center (black-cross) color. Similarly, these ellipse thresholds can also be sampled in linear *sRGB* space. (b) shows ellipses of samples across a $3 \times 3 \times 3$ grid sampled within the *sRGB* cube. Each ellipse in this illustration is shaded with the color at its center. Since equi-luminant planes are parallel in linear *sRGB*, all the ellipses appear on parallel planes. (c) visualizes the model-guided chromaticity shifting (at 25° eccentricity) to minimize display power consumption. We use our perceptual constraints in combination with our power cost function to shift the chromaticities of various sample colors inside the *sRGB* color cube (illustrated in Figure 5). The original, and power-optimized colors correspond to the tail and head of the vectors, respectively.

Modeling ellipse level sets. In our model, we represent the set of all equi-luminant colors which cannot be discriminated from a test color, $\mathbf{t} \in i\text{-DKL}$, relative an adaptation color, $\mathbf{b} \in i\text{-DKL}$, using an ellipse-shaped region centered around the color contrast of the test color. The boundary of this ellipse region corresponds to the discriminative threshold of $\kappa(t_i, b_i)$, for $i \in \{L-M, S-(L+M)\}$. The set of color coordinates which represent this threshold, $\mathbf{x} \in i\text{-DKL}$, fulfill the system of equations:

$$\begin{cases} x_{L+M} &= b_{L+M} \\ \mathcal{E}(\mathbf{x}; \mathbf{t}, \mathbf{b}, \boldsymbol{\alpha}) &= 0. \end{cases} \quad (3)$$

The first constraint ensures that all the color coordinates on the threshold are equi-luminant to the adaptation color. The second constraint ensures that all \mathbf{x} are on the edge of the ellipse region with major and minor semi-axes equal to $\boldsymbol{\alpha} = (\alpha_{L-M}, \alpha_{S-(L+M)}) \in \mathbb{R}^2$. Formally, the function $\mathcal{E}(\cdot)$ is defined as

$$\mathcal{E}(\mathbf{x}; \mathbf{t}, \mathbf{b}, \boldsymbol{\alpha}) = \sum_{i=\{L-M, S-(L+M)\}} \left(\frac{\kappa(x_i; b_i) - \kappa(t_i; b_i)}{\alpha_i} \right)^2 - 1, \quad (4)$$

Model Regression. Equation (4) requires the knowledge of the ellipse-size parameters, α_i . Prior work has shown that α_i relates to the color-contrasts of various test colors, $\kappa(t_i, b_i)$, as well as the retinal eccentricity, $e \in \mathbb{R}^+$, at which a colored stimulus is displayed [Hansen et al. 2009; Krauskopf and Karl 1992]. We leverage our user study data from Section 3 to learn the relationship

$$\Phi : (\boldsymbol{\kappa}, e) \mapsto \boldsymbol{\alpha} \quad (5)$$

where $\boldsymbol{\kappa} \in \mathbb{R}^2$ are the $L-M$ and $S-(L+M)$ coordinates of the test color in *DKL* space computed using Eqs. (1) and (2). Specifically, we

use our data to optimize a shallow neural network, which estimates the discrimination thresholds, using least-squares regression:

$$\hat{\boldsymbol{\alpha}} = \arg \min_{\Phi} \|\Phi(\boldsymbol{\kappa}, e) - \boldsymbol{\alpha}\|_2^2. \quad (6)$$

The R^2 value of the regression is 0.58 (adjusted R^2 value of 0.51), indicated an acceptable regression accuracy. Note that our raw data from Section 3 is intentionally pre-processed as described in detail in Section 6.1. Briefly, we aim to cover more *conservative* thresholds that are generalizable to broad users instead of an “average fit”.

Neural Network Architecture. We chose the Radial Basis Function Neural Network (RBFNN) with a sigmoid output layer to ensure local smoothness, as well as a positive, localized output range. Mathematically, the network is summarized as

$$\Phi(\boldsymbol{\kappa}, e) = \boldsymbol{\eta} \odot S \left(\sum_{j=1}^N \lambda_j \rho \left(\left\| \begin{bmatrix} \boldsymbol{\kappa} \\ e \end{bmatrix} - \mathbf{c}_j \right\|_2, \sigma_j \right) + \nu_j \right), \quad (7)$$

where \odot is the term-wise multiplication operator. The RBFNN takes the input, and computes the weights of the effect each of the N nodes of the latent representation have on the input. It does so by applying a Gaussian Radial Basis function, ρ , centered at \mathbf{c}_j with std of σ_j , for each node, j . The weights of each node is scaled by a scaling constant λ_j , incremented by the linear bias ν_j , summed up, and passed to the sigmoid function S and multiplied by a scaling factor $\boldsymbol{\eta}$ to produce the final prediction. The trainable parameters of this network are the centres, \mathbf{c}_j , sizes, σ_j of the radial bases, as well as the final scaling factors λ_j , and linear biases ν_j . $\boldsymbol{\eta}$ is a normalization constant and chosen to be the maximum possible value of contrasts within the capability of the display used in our

work, and hence does not change. For our work we keep the number of nodes $N = 5$ low to maintain smoothness of the outputs. Please refer to our source code for more details on the model specifications.

Ellipse re-parameterization. Since the adaptation color, \mathbf{b} , is the same for all variables in Eq. (4), we simplify the function by re-parameterization as $a_i = \alpha_i b_i$ for $i \in \{L - M, S - (L + M)\}$:

$$\mathcal{E}(\mathbf{x}; \mathbf{t}, \mathbf{a}) = \sum_{i=\{L-M, S-(L+M)\}} \left(\frac{x_i - t_i}{a_i} \right)^2 - 1. \quad (8)$$

While the original formulation in Eq. (4) relates the ellipse to variables in DKL space, and are ultimately the variables used to regress the model, as we'll see in Section 5.2, it's helpful to reformulate the model with respect to the differences of color in i - DKL : $\mathbf{x} - \mathbf{t}$, as well as the new parameter, \mathbf{a} , which clearly represents the size of the ellipse in i - DKL space, are both measures defined within i - DKL . In summary, Eq. (8) converts Eq. (4) from a DKL space parameterization to a i - DKL parameterization. The obtained model is visualized in Figures 4a and 4b.

5.2 Power Model for Display Illumination

In this section we derive a computation model that correlates an OLED's power consumption with the pixel color. The display power is modeled as the sum of the LED power, which consists of the powers of its three sub-pixels, and the power of the peripheral circuitry (e.g., the thin-film transistors) [Huang et al. 2020]. It is known that the power of an OLED sub-pixel is roughly proportional to its current, which is proportional to the numerical value of the corresponding channel [Tsujimura 2017]. Thus, given the RGB value of the three sub-pixels, $\mathbf{x}_{disp} \in disp-RGB$ (i.e., the pixel value in the display native color space), its total power consumption is

$$\mathcal{P} = \left(\sum_{i=\{1,2,3\}} p_i x_i \right) + p_{circ} = \mathbf{p}_{disp}^T \mathbf{x}_{disp} + p_{circ}, \quad (9)$$

where $\mathbf{p}_{disp} \in \mathbb{R}^3$ is the vector of unit powers of each sub-pixel, and $p_{circ} \in \mathbb{R}$ is the static power consumption (consumed by the peripheral circuits) when all the pixels are black, i.e., the LEDs do not emit light and, thus, do not consume power.

In most computer graphics applications, it is impractical to use the display's native color space because it varies depending on the manufacturer specifications, and could be unknown. Color-spaces that are commonly used, such as (linear) $sRGB$, can transform to a display's native color-space via some linear transformation, $M \in \mathbb{R}^{3 \times 3}$. Without loss of generality, using this transformation, we can rewrite Eq. (9) in terms of the (linear) $sRGB$ pixels as

$$\begin{aligned} \mathcal{P}(\mathbf{x}_{srgb}) &= \mathbf{p}_{disp}^T M_{srgb2disp} \mathbf{x}_{srgb} + p_{circ} \\ &= \mathbf{p}_{srgb}^T \mathbf{x}_{srgb} + p_{circ}, \end{aligned} \quad (10)$$

where $M_{srgb2disp}$ is the transformation matrix from (linear) $sRGB$'s color-space to the display's, and $\mathbf{x}_{srgb} \in sRGB$ denotes the pixel color in linear $sRGB$ space. For convenience, we define $\mathbf{p}_{srgb}^T = \mathbf{p}_{disp}^T M_{srgb2disp}$, which intuitively denotes the power consumption of the three display sub-pixels under unit $sRGB$ stimuli.

\mathbf{p}_{srgb} depends on the specification of a particular display. In our work, we study an OLED display module from Wisecoco that has two 1080×1200 displays, as described in Section 4. Critically, our methodology is not unique to the specific display at study and, thus, can be extended to build power models for any other three-primary display.

Power model regression. To build an analytical power model, we must find the parameter, \mathbf{p}_{srgb} . We do so by physically measuring the power consumption of 52 randomly sampled colors in the $sRGB$ space, including the eight colors that correspond to the eight vertices of the $sRGB$ color cube, as described in Section 4, and solving an over-determined linear system,

$$\mathcal{P}^{(color)} = \mathbf{p}_{srgb}^T \mathbf{x}_{srgb}^{(color)} + p_{circ}, \quad (11)$$

where $color$ is the 52 sampled colors, via the classic linear least squares method. Figure 3c shows the measured power of these sampled colors (y -axis) and the regressed model outputs (x -axis). The mean relative error of the regression is 0.996%, indicating an accurate model.

5.3 Optimizing Display Energy Consumption under Perceptual Constraints

Finally, using Eq. (3) and Eq. (10) we can minimize the power consumption function of a display, $\mathcal{P}(\mathbf{x})$, while constrained within the perceptual limits set by $\mathcal{E}(\mathbf{x})$. Qualitatively, we notice that the power function is a linear function of the input, \mathbf{x} , so the minimizing power will be on the surface of the discriminative threshold ellipse (as opposed to its interior). Notice that in this optimization problem, it is more convenient to use the definition of the ellipses in i - DKL space, instead of the DKL space definition (cf. Eq. (8) and Eq. (4)) because the i - DKL space is only a linear transformation away from (linear) $sRGB$, therefore making its energy computations a lot simpler.

Formally, we define the optimization process as:

$$\begin{aligned} \mathbf{x}_{idkl}^* &= \arg \min_{\mathbf{x}_{idkl}} \mathcal{P}(M_{idkl2srgb} \mathbf{x}_{idkl}) \\ \text{subject to: } \mathcal{E}(\mathbf{x}_{idkl}; \mathbf{t}_{idkl}, \mathbf{a} = \boldsymbol{\alpha} \odot \mathbf{b}_{idkl}) &= 0, \end{aligned} \quad (12)$$

where the original color of the pixel is \mathbf{t} and, the adaptation color of the display is \mathbf{b} . In our work we choose \mathbf{b} to be equal to a color with a chromaticity equal to the CIE D65 Standard Illuminant (i.e., the reference white in the $sRGB$ color space) and a luminance equal to the luminance of the test pixel \mathbf{t} . While the choice of adaptation color is an interesting question to explore, it is beyond the scope of this work and is left as future work.

Due to the convexity of both the cost and constraint functions, we can apply the method of Lagrange multipliers to find the output color, \mathbf{x}_{srgb}^* , which minimizes the total power consumption in closed form:

$$\mathbf{x}_{srgb}^* = M_{idkl2srgb} \begin{bmatrix} \frac{p_1 a_1^2}{\sqrt{p_1^2 a_1^2 + p_2^2 a_2^2}} \\ \frac{p_2 a_2^2}{\sqrt{p_1^2 a_1^2 + p_2^2 a_2^2}} \\ t_3 \end{bmatrix}, \quad (13)$$

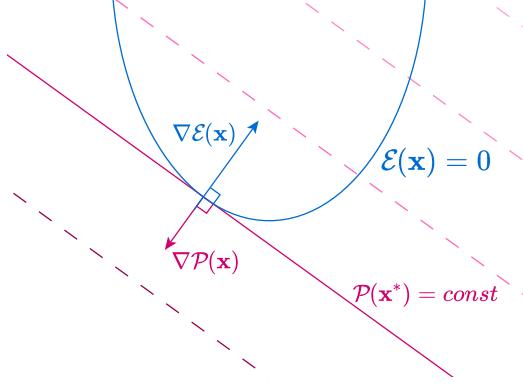


Fig. 5. Illustration of deriving the closed-form and optimal chromaticity. We aim to search the color \mathbf{x}^* which (1) minimizes our first-order power function, $\mathcal{P}(\mathbf{x})$; (2) is under the constraint $\mathcal{E}(\mathbf{x}) = 0$ when the gradients of both functions are co-linear. The straight lines represent the set of all colors consuming the same amount of power. Since the color-spaces *sRGB* and *i-DKL* are related via a linear transformation, this solution works in either color-space.

where $\{1, 2, 3\}$ correspond to $\{L - M, S - (L + M), L + M\}$. Figure 5 visually illustrates how this optimal color is found using the derivatives of \mathcal{E} , and \mathcal{P} . Please refer to Supplement A for the derivation of the above result.

6 IMPLEMENTATION

6.1 Perception Study Data Pre-processing

We take two steps to pre-process the perception study data. Both steps are meant to keep the model’s threshold estimation conservative, which is necessary for two reasons. First, there are natural variances across participants (Section 3.2) and, thus, a conservative estimation allows our model to generalize to large populations. Second, our model is built to modulate the displayed colors to preserve the visual fidelity in active viewing, which we hypothesize to have a lower threshold than that in discriminative tasks.

First, we use the smallest thresholds across participants, instead of an average fit. Second, we observed small asymmetries in the collected thresholds, and we confirmed that this is also the case in Krauskopf and Karl et al. [1992]. We made the engineering decision to keep our model’s thresholds more conservative; thus each threshold is chosen to be the narrower one from the two thresholds approached from opposing sides along a *DKL* axis. That is, given a threshold approached from the positive $L - M$ side, α_{L-M}^+ , and one from the negative $L - M$ side, α_{L-M}^- , the discrimination threshold we pick for model regression is:

$$\alpha_{L-M} = \min(\alpha_{L-M}^+, \alpha_{L-M}^-), \quad (14)$$

and similarly for the $S - (L + M)$ axis.

6.2 Eccentricity Extrapolation

In our perceptual model regression, we restricted the range of valid input eccentricities to be between 10° and 35° because we had only measured discriminative thresholds within this range of eccentricities. We avoided color-shifting content at eccentricities $< 10^\circ$ due

to the low power-saving payoffs for foveal and para-foveal regions. Meanwhile, eccentricities $> 35^\circ$ were clamped down to 35° as a conservative estimate.

6.3 Shader

We implement a post-processing image-space shader in the Unity ShaderLab language to compute per-pixel power-minimizing color. Figure 6 outlines the pseudocode of our shader. We tested our shader on the HTC Vive Pro Eye (relevant specs shown below) powered by an NVIDIA RTX3090 GPU, and observed that processing each frame takes less than 11 ms, which ensures no loss of frames in the displays.

Table 2. Relevant specifications for the HTC Vive Pro Eye

Feature	Specification
Display Resolution	1440 × 1600 pixels per eye
Display Refresh-rate	90 Hz
Peak Luminance	143 cd/m ²
Eye-tracker Accuracy	0.5° – 1.1°
Eye-tracker Frequency	120 Hz.

7 EVALUATION

In this section we evaluate the performance and applicability of our model. In Section 7.1, we first conduct a psychophysical experiment to assess the perceptual fidelity between our method and an alternative luminance-based power reduction approach. The experiment design follows previous literature [Cohen et al. 2020]. Then, in Section 7.2, we measure the model’s generic benefits in broad applications by further analyzing the display power saving ratio with a large scale natural image dataset, ImageNet.

7.1 Psychophysical Study for Perceptual Quality

Motivated by the experiment of Cohen et al. [2020], we conduct a psychophysical user study to measure participant-experienced fidelity deterioration, as well as the corresponding power-saving level during active and real-world viewing. “Active and real-world” is notably a condition where participants may freely rotate their head/eyes and naturally investigate an immersive scene.

Setup and participants. We recruited 13 participants (ages 21-32, 3 female). None of the participants were aware of the research, the hypothesis, or the number of conditions. All participants have normal or corrected-to-normal vision. We used the same hardware setup as our preliminary user study in Section 3. Before each experiment, we ran a five-point eye-tracking calibration for each participant.

Stimuli and conditions. The stimuli were 6 panoramic video sequences as shown in Figure 8. For broader coverage, the tested scenes contain natural/synthetic, static/dynamic, bright/dark, and indoor/outdoor content.

We studied the perceptual quality by applying two gaze-contingent and power-saving shading approaches to the scenes: a baseline luminance-modulated shader, LUM; and the shader with our chromaticity modulation method, OUR (Section 6.3). Specifically, in LUM, we applied a constant scaling factor to all peripheral (eccentricity $> 10^\circ$) pixels’ colors. That is, LUM can be understood as a

Power-Minimizing Color Computation

```

1: function FRAGMENT_SHADER
2:   POS_PIXEL = GETPIXELPOS
3:   POS_GAZE = GETGAZEPOS
4:   t_SRGB = SAMPLETEXTURE(MainTex, POS_PIXEL)
5:   e = GETECCENTRICITY(POS_GAZE, POS_PIXEL)
6:   if e < ECC_MIN then
7:     return t_SRGB
8:   else if e > ECC_MAX then
9:     e = ECC_MAX
10:  end if
11:  t_IDKL = M_SRGB2IDKL t_SRGB
12:  LUM = t_IDKL.Z
13:  ▷ Adaptation color
14:  b_IDKL = M_SRGB2IDKL [LUM, LUM, LUM]T
15:  ▷ ◦: element-wise multiply
16:  κ = (t_IDKL - b_IDKL) ◦  $\frac{1}{b_{IDKL}}$ 
17:  INPUT =  $\begin{bmatrix} \kappa \\ e \end{bmatrix}$  ▷ Model input
18:  INIT RBF[5], LINEAR[2] ▷ Model output
19:  ▷ Lines 14-22: RBFNN from equation (7)
20:  for i in [0, 5) do
21:    RBFi  $\leftarrow \rho(\|\text{INPUT} - \mathbf{c}_i\|_2, \sigma_i)$ 
22:  end for
23:  for i in [0, 2) do ▷ Linear layer
24:    LINEARi  $\leftarrow \lambda_i \odot \text{RBF} + \nu_i$ 
25:  end for
26:   $\alpha \leftarrow \eta \odot \text{SIGMOID}(\text{LINEAR})$ 
27:  a =  $\alpha \odot b_{IDKL}$ 
28:  ▷ Compute power-optimal color
29:  p = GETPOWERMODELCOEFS
30:  xIDKL* = t_IDKL + a2 ◦ p  $\frac{1}{\sqrt{\sum_i a_i^2 p_i^2}}$ 
31:  return M_IDKL2SRGB xIDKL*
32: end function

```

Fig. 6. *Shader implementation pseudocode.* The shader routine optimizes an input color into the optimized color as described by our method from Section 5. Refer to our source code for the ShaderLab implementation.

gaze-contingent version of the “power-saving mode” on mobile devices. The scaling factor was determined in such a way that the power saving (estimated using the power model) of LUM is similar to that of OUR. An example frame of the original stimulus, OUR, and LUM are shown in Figure 7.

Similar to [Cohen et al. 2020], we *temporally* inserted one of the two shaders to the original stimulus during each trial. More formally, let I_o be the original video and I_p be the power optimized version. Then starting at timestamp $\tau_o = 5$ s, we linearly interpolate (i.e. *lerp*) between I_o and I_p over a course of 10 seconds. At $\tau_p = 15$ s, the transition completes and the power optimized video is played

henceforth. That is,

$$I(\tau) = \text{lerp}\left(I_o, I_p, \frac{\tau - \tau_o}{\tau_p - \tau_o}\right) \quad (15)$$

The process is illustrated in Figure 7d. Note that the temporal insertion also implicitly compares the original frame with each of the two power-saved conditions.

Tasks. Our experiment consisted of 24 trials (6 scenes \times 2 condition \times 2 repetitions), lasting approximately 15 minutes for each participant. Before the experiment started, we first displayed 2 trial runs to familiarize the participant with the setup. Afterwards, six 20-second video sequences (with representative frames displayed in Figure 8) were shown to the participant in a counter-balanced randomized order.

During each trial, the participant was instructed to perform a scene-specific task, such as “count the number of chairs” to ensure they were actively viewing the scene. After each trial, the participant was instructed to answer both the scene-specific task and a two-alternative forced choice (2AFC, similar to [Cohen et al. 2020]) question “did you notice any visual artifacts?”. Before beginning the experiment, we show the participants static frames from the “skyline” video as visual examples of “artifact” stimuli, including one original frame and the two corresponding conditions, OUR, and LUM. The example images were displayed on a computer monitor (as opposed to the VR headset); thus, the participants’ retinal image was significantly different from the stimuli shown during the study. This is to ensure that the participants are not biased when shown artifacts.

Metrics and results. We use the percentage of trials where participants noticed artifacts as the metric of perceptual quality. Lower values indicate better quality, i.e., less noticeable visual modulation. Figure 7e plots the user-reported values of each scene and each condition. As visualized in Figure 7e, the average percentage of observed artifacts in LUM is $63.5 \pm 9.4\%$ (STE) and in OUR is $16.7 \pm 7.3\%$. The lowest percentage of observed artifacts in scenes with OUR applied occurred in the monkeys scene, a scene with large amounts of green, whereas the highest percentage occurred in the office scene, a very bright and uniformly colored scene relative to the other scenes. A one-way repeated ANOVA analysis showed that the shading condition (OUR vs. LUM) has a significant effect on the perceptual quality ($F_{(1,24)} = 18.42, p = .00025$).

As plotted in Figure 7f, we also measured the display power consumption of each power-saved shading condition for each scene. The average savings between OUR and LUM are similar ($20.8 \pm 1.2\%$ vs. $18.6 \pm 1.4\%$ (95% confidence)). OUR exhibits the highest power saving in the skyline scene, due to higher relative uniform distribution of blue colors.

Discussion. The results reveal our method OUR’s significant out-performance on preserving perceptual quality over a gaze-contingent luminance-reduction-based approach (LUM), even though both conditions achieve a comparably similar power-saving scale. Note that, under the same power, there are infinite ways of constructing LUM, including smoothing the edge but darkening the farther periphery. Our implementation of LUM is partially inspired by Pöppel and

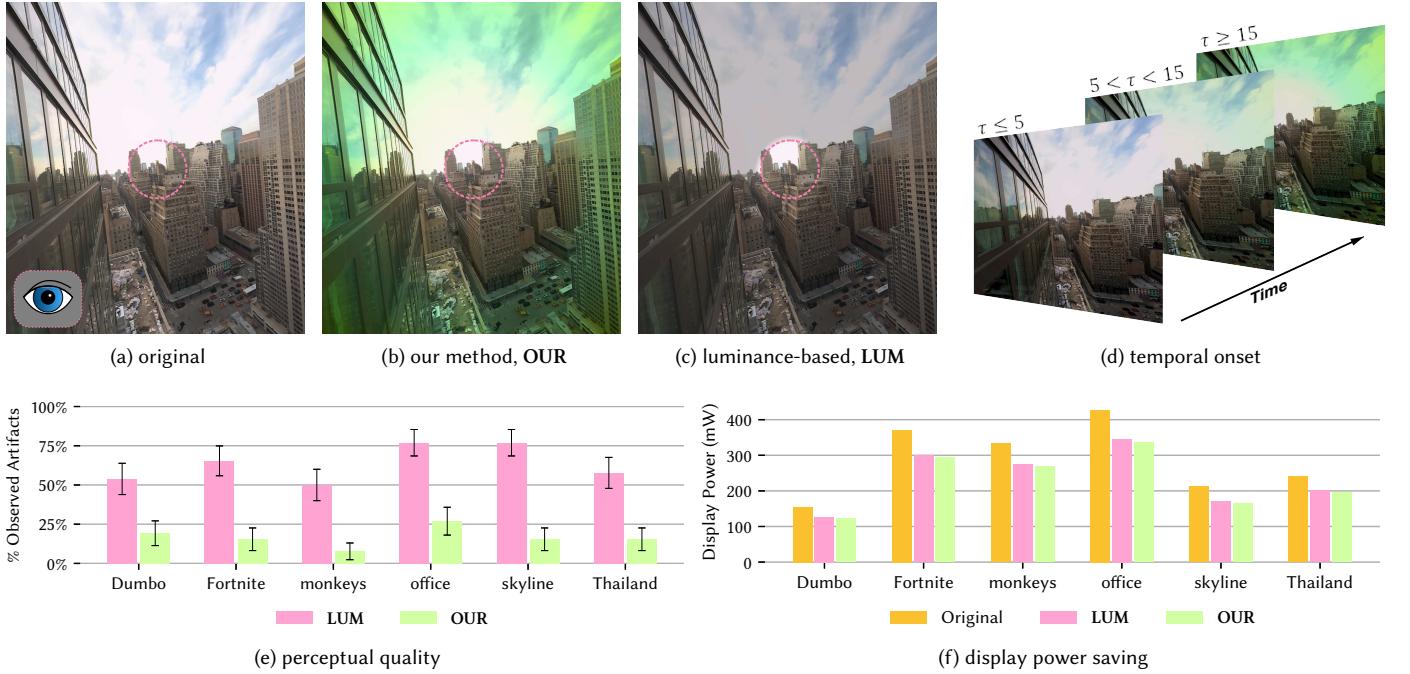


Fig. 7. *User study stimuli and results.* (b) and (c) show the results of applying the two gaze-contingent shading conditions OUR and LUM to an example frame (a) in the video sequence stimuli, with the dashed circles indicating the user’s gaze. (d) We increase the intensity of the gaze-contingent filter in a temporal fashion, while keeping the foveal region unchanged. From 0-5 seconds, the video is unfiltered. Between 5-15 seconds we gradually insert one of the two conditions, with the remaining time between 15-20 seconds the full filter is applied. Please refer to our supplementary video for a dynamic visualization of this temporal process. (e) shows the percentage of trials where users identified “artifacts”. Across the 6 scenes, OUR exhibits significantly lower values than LUM, evidencing our method’s benefit in preserving perceptual fidelity. The error bars indicate standard error across users. (f) Using the hardware setup in Section 4, we physically measure the scene-dependent display power (measured power excluding the $\sim 1000\text{mW}$ peripheral circuit power) consumption of each 20-second video clip among three conditions: the original, LUM, and OUR. OUR achieves similar power saving capability to LUM. Unmodified image credits to Humaneyes Technologies.

Harvey [1973], which suggests that human luminance change detection thresholds remain relatively constant beyond 10° eccentricity. The design, however, may not be perceptually optimal. Therefore, studying and modeling the luminance-induced effects may not only provide a stronger baseline condition, but improve our model that is currently restricted to colors only.

Our perceptual fidelity and power-saving capabilities are also content-based. For example, we notice high power savings in the “office” scene, but the average % of observed artifacts is higher than other scenes. This is hypothetically because the scene has a significantly higher brightness compared to the others. On the other hand, the “monkey” scene has relatively high density of green colors, and thus has the lowest perceived average % of artifacts. The observations motivate us to investigate, in the future, the chromaticity-luminance joint effect (Section 8) beyond the first model that guides color-perception-aware VR power optimization.

While detection tasks are commonly leveraged in the foveated rendering literature [Patney et al. 2016; Sun et al. 2017], we opted to validate our gaze contingent filter with active and natural viewing, similar to [Cohen et al. 2020]. The design choice is two-fold. First, we attempt to simulate the conditions of real-world VR applications

where the users, with various tasks in mind, make head and eye movements to explore the environment. Examples include gaming and video-watching. Second, prior literature suggests unequal color detection and discrimination thresholds. Vingrys and Mahon [1998] discovered that chromatic sensitivity for detection is significantly greater than for discrimination. However, by leveraging our model and shader in this experiment, we verify our hypothesis that our color sensitivity during natural and active vision is, in fact, lower than that of discrimination, and thus enable the method’s applicability for broad VR scenarios.

7.2 Measuring Power-Saving Capability for Broad Content

In our psychophysical study Section 7.1, we observe that the possible power savings are dependent on the displayed content. For example, colors that are highly saturated have little room in their equi-luminant plane that is within the bounds of the sRGB cube. Therefore, they have less power saving potential as any potential power-saving chromaticity shifts are clipped by the sRGB bounds. To study how much power can be saved in practical applications where users may observe arbitrary imagery, we conduct an objective evaluation by applying our method to a large sample of the



Fig. 8. *Panoramic Video Scenes.* Representative panoramic frames captured from 360 degree monoscopic video scenes used in the evaluation user study described in Section 7.1. Images of “Dumbo”, “monkeys”, and “skyline” scenes by Humaneyes Technologies, rendering of historical Two Embarcadero Center in San Francisco, CA (“office”) by Rene Rabbitt of Rabbitt Design, “Thailand” by VR Gorilla, and “Fortnite” by AmiramiX.

ImageNet dataset [Russakovsky et al. 2015]. We then measure the distribution of power savings using our power model.

Setup. We simulate how an image would be observed in a VR setting by resizing the image to be displayed at 90° field-of-view, and randomly sample a location within the image and select that as the gaze location. The randomization of the gaze location was applied to prevent bias in the power estimation. Specifically, we observed that many images in ImageNet have a foreground object centered on the image; selecting random gaze locations allows us to include images where the foreground objects will sometimes have the filter applied to them. We repeat this process for 10% (randomly sampled) of the ImageNet dataset, totaling in over $\sim 120k$ images, to collect original and power-optimized image pairs.

Metrics. For each image pair, we measure the estimated power consumption using our model from Section 5.2, and compute the relative decrease in power consumption by applying our filter with respect to the ground-truth condition.

Results. We observe that the mean display power saving recorded across the entire dataset is 13.9%, and guarantee 9.1 – 23.5% savings with $P95$ confidence. Please refer to Figure 9c for the detailed histogram of the estimated power savings. We visualize a small sample of the images we applied the filter to in Figures 9a and 9b.

Discussion. Sample images from different percentiles of power-saving as shown in Figure 9 show that images with the highest power savings are commonly bright and/or blueish scenes, and vice versa. Intuitively, bright scenes provide larger *percentage* changes in LED luminance, and thus unlock larger space for power-saving. Since

blue colors on the LED consume the most power, as demonstrated in Section 4, images rich in blue/green colors can be optimized most effectively. Meanwhile, images which are already saturated with red colors cannot be optimized for higher power-saving because the space of power-wise “cheaper” colors is narrower.

8 LIMITATIONS AND FUTURE WORK

Active vision vs. discrimination vs. detection. While our evaluation on active/natural viewing tasks in Section 7.1 is representative of real-world VR scenarios [Cohen et al. 2020], our initial perceptual data are collected using a more conservative *discrimination* task. It is also common in the foveated rendering literature to evaluate using *detection* tasks [Patney et al. 2016; Sun et al. 2017]. Our conducted preliminary detection-based experiments, which showed that sensitivity to color changes in detection tasks is significantly greater than that in discrimination tasks, are consistent with prior work [Vingrys and Mahon 1998].

An exciting future direction is, thus, to investigate an adaptive model that accommodates for color sensitivity under all three tasks (detection, discrimination, active/natural viewing). That way, our color modulation algorithm can be dynamically configured according to the specific viewing scenario of a VR user.

Perceptual model. Our current perceptual model is constructed with respect to per-pixel colors. An interesting future extension is to consider inter-pixel, potentially higher-dimensional, features such as spatial frequency and local contrast. Performing these analyses (e.g., frequency domain analyses as in [Tursun et al. 2019]), however, increases the computational overhead. How to best balance the level of details in perceptual analysis and display power saving is an open question we leave to future work.

Luminance adjustment. In our work, we model and modulate pixel *chromaticity* to reduce display power consumption while preserving *luminance*. This design choice reduces the dimension of our perceptual model and, thus, yields a convex constrained optimization problem with a closed-form solution. Investigating the luminance-chromaticity joint modulation is an interesting future research direction that would conceivably lead to higher power savings [Vingrys and Mahon 1998].

Jointly adjusting luminance and chromaticity, nevertheless, comes with a few challenges. First, it would require sampling a new dimension in constructing the perceptual model. Second, prior literature suggests the weak eccentricity-dependent effect [Metha et al. 1994] in detecting and discriminating absolute luminance. Finally, the perceptual level sets, when considering the luminance dimension, might not be convex, which might complicate the optimization, cause false local minima, and reduce the shading speed.

Color Temperature Adaptation. Another interesting direction for future research is to leverage chromatic adaptation [Fairchild 2013] to reduce display power by adjusting the color temperature of the display white point. The advantage of this approach is that it is not gaze-contingent: it can potentially reduce the power of the entire display without requiring eye tracking. Adaptation to display color has been long investigated [Fairchild and Reniff 1995; Peng et al. 2021], but such studies in VR displays are relatively new and

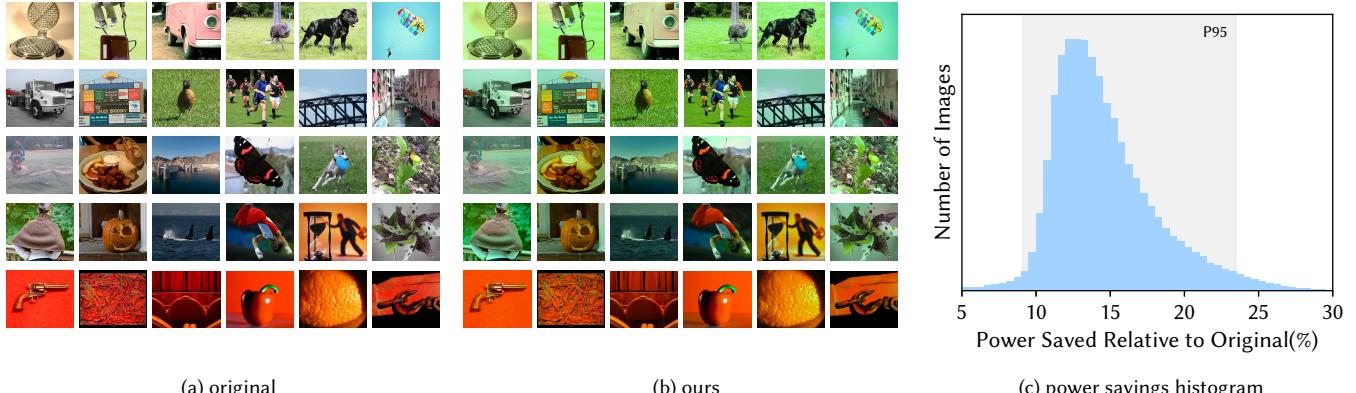


Fig. 9. *ImageNet Offline Power Savings Estimation.* We measure the statistics of power saved when our model is applied to the large database of natural images, ImageNet. (a) The original images are sampled randomly to show exemplar images which exhibit different amounts of power-saving when our model is applied to them. Each row of images are sampled from within different quantile bands of energy savings across the dataset (top to bottom each row saves better than 99.9%, 55%, 45%, 0.1% of the entire dataset considered, and the bottom row shows the worst 0.1% performers). (b) Our shaders are applied to each image have randomized gaze locations to prevent bias against the edges of the image being disproportionately affected. (c) We visualize the entire distribution of the potential power savings in this evaluation.

rare [Chinazzo et al. 2021] and lack a comprehensive computational model. Note that the chromatic adaptation benefits are additive: our model can be seen as a sub-space initial attempt (by exploiting spatial color perception) under a *given* adaptation state.

Display Persistence. VR displays usually have low persistence to reduce motion blur [Hainich and Bimber 2016]. A common solution is to hold a frame for only a short period of time during each display refresh [Google 2019]. As a consequence, the display power is relatively low to begin with (compared to displaying a frame throughout a refresh cycle). Nevertheless, our work demonstrates significant display power saving opportunities even with reduced displayed times. In addition, reducing the display period leads to low average luminance, which limits the applicability of a luminance-based approach to reduce power – another reason we choose to maintain the luminance.

Implementation. Our implementation of *LMS* to *DKL* color space transformation does not strictly follow the canonical implementation [Derrington et al. 1984; Westland et al. 2012]. Notably, we judiciously choose to use a per-luminance adaptation color rather than one single adaptation color. Effectively, this allows us to eliminate the luminance dimension from all colors. As a result, all subsequent color adjustments naturally preserve the luminance – a goal we set out to achieve. This design choice also simplifies the color space transformation and contributes to the real-time speed of the shader.

There is room for improving the speed of our shader, which currently is bottlenecked by the atomic *for*-loop. Deferred shading techniques may shed light on alleviating the bottleneck. One promising solution is to evaluate the optimization problem (Equation (12)) offline (e.g., sampling colors and eccentricities) and save the results as a 3D texture, which the shader simply looks up at rendering time.

Due to the tight integration of the display, computation module, and battery in commercial AR/VR devices, our display power measurement has to be done on a 3rd party display module that has the identical aspect ratio of the VR device we use for perceptual studies. Investigating physical means to measure the exact display power as in an AR/VR device would reveal the real-world energy savings concerning the battery equipped with the device. It would also be interesting to see how our perception-conserving color modulation idea can be applied to smartphone displays, which have much narrower field-of-views.

9 CONCLUSION

Understanding the nature of color generation and perception is a long-standing pillar for computer graphics. In this research, we build a bridge from color to the emerging demand for power-friendly graphics with the rapid growth of wireless platforms. In particular, we present a perceptually motivated computational model for optimizing the power usage of VR displays based on the limited discriminative ability of color by the human visual system. We validate the model’s application with various real-world immersive content viewing use-cases and physically-measured power consumptions. While the research is now benefiting untethered VR applications, we envision potential in exploiting more properties of the “human” aspects of the study of color and how to incorporate them into generic computer graphics systems. We hope that our work to develop the first building block in this direction will pave the way for more sophisticated applications, such as cloud rendering.

ACKNOWLEDGMENTS

The work was supported, in part, by the National Science Foundation (NSF) under grants #2225861 and #2044963.

REFERENCES

Mark Bohr. 2007. A 30 year retrospective on Dennard's MOSFET scaling paper. *IEEE Solid-State Circuits Society Newsletter* 12, 1 (2007), 11–13.

ML Boroson, JE Ludwicki, and MJ Murdoch. US Patent 7,586,497, Sep. 8, 2009. OLED display with improved power performance.

James K Bowmaker and HJ Dartnall. 1980. Visual pigments of rods and cones in a human retina. *The Journal of physiology* 298, 1 (1980), 501–511.

Doris I Braun, Alexander C Schütz, and Karl R Gegenfurtner. 2017. Visual sensitivity for luminance and chromatic stimuli during the execution of smooth pursuit and saccadic eye movements. *Vision Research* 136 (2017), 57–69.

Giorgia Chinazzo, Kynthia Chamlothori, Jan Wienold, and Marilynne Andersen. 2021. Temperature–color interaction: subjective indoor environmental perception and physiological responses in virtual reality. *Human Factors* 63, 3 (2021), 474–502.

Michael A. Cohen, Thomas L. Botch, and Caroline E. Robertson. 2020. The limits of color awareness during active, real-world vision. *Proceedings of the National Academy of Sciences* 117, 24 (2020), 13821–13827. <https://doi.org/10.1073/pnas.1922294117> arXiv:<https://www.pnas.org/content/117/24/13821.full.pdf>

Bevil R Conway, Rhea T Eskew Jr, Paul R Martin, and Andrew Stockman. 2018. A tour of contemporary color vision research. *Vision research* 151 (2018), 2–6.

Herbert JA Dartnall, James K Bowmaker, and John Dixon Mollon. 1983. Human visual pigments: microspectrophotometric results from the eyes of seven persons. *Proceedings of the Royal society of London. Series B. Biological sciences* 220, 1218 (1983), 115–130.

Pranab Dash and Y. Charlie Hu. 2021. How Much Battery Does Dark Mode Save? An Accurate OLED Display Power Profiler for Modern Smartphones. In *Proceedings of the 19th Annual International Conference on Mobile Systems, Applications, and Services (Virtual Event, Wisconsin) (MobiSys '21)*. Association for Computing Machinery, New York, NY, USA, 323–335. <https://doi.org/10.1145/3458864.3467682>

Russell L De Valois, Israel Abramov, and Gerald H Jacobs. 1966. Analysis of response patterns of LGN cells. *JOSA* 56, 7 (1966), 966–977.

Kurt Debattista, Keith Bugeja, Sandro Spina, Thomas Bashford-Rogers, and Vedad Hulusic. 2018. Frame rate vs resolution: A subjective evaluation of spatiotemporal perceived quality under varying computational budgets. In *Computer Graphics Forum*, Vol. 37. Wiley Online Library, 363–374.

Andrew M Derrington, John Krauskopf, and Peter Lennie. 1984. Chromatic mechanisms in lateral geniculate nucleus of macaque. *The Journal of physiology* 357, 1 (1984), 241–265.

Mian Dong, Yung-Seok Kevin Choi, and Lin Zhong. 2009. Power modeling of graphical user interfaces on OLED displays. In *2009 46th ACM/IEEE Design Automation Conference*. IEEE, 652–657.

Mian Dong and Lin Zhong. 2011. Chameleon: A color-adaptive web browser for mobile OLED displays. In *Proceedings of the 9th international conference on Mobile systems, applications, and services*, 85–98.

Andrew Duchowski, Nathan Courania, and Hunter Murphy. 2005. Gaze-Contingent Displays: A Review. *Cyberpsychology & behavior : the impact of the Internet, multimedia and virtual reality on behavior and society* 7 (01 2005), 621–34. <https://doi.org/10.1089/cpb.2004.7.621>

Mark D Fairchild. 2013. *Color appearance models*. John Wiley & Sons.

Mark D Fairchild and Lisa Reniff. 1995. Time course of chromatic adaptation for color-appearance judgments. *JOSA A* 12, 5 (1995), 824–833.

Hugh S Fairman, Michael H Brill, and Henry Hemmendinger. 1997. How the CIE 1931 color-matching functions were derived from Wright-Guild data. *Color Research & Application* 22, 1 (1997), 11–23.

Google. 2019. Fundamental Concepts of Google VR. https://developers.google.com/vr/discover/fundamentals#display_persistence.

Brian Guenter, Mark Finch, Steven Drucker, Desney Tan, and John Snyder. 2012. Foveated 3D graphics. *ACM Transactions on Graphics (TOG)* 31, 6 (2012), 1–10.

John Guild. 1931. The colorimetric properties of the spectrum. *Philosophical Transactions of the Royal Society of London. Series A, Containing Papers of a Mathematical or Physical Character* 230, 681–693 (1931), 149–187.

Rolf R Hainich and Oliver Bimber. 2016. *Displays: fundamentals & applications*. AK Peters/CRC Press.

Matthew Halpern, Yuhao Zhu, and Vijay Janapa Reddi. 2016. Mobile CPU's rise to power: Quantifying the impact of generational mobile CPU design trends on performance, energy, and user satisfaction. In *2016 IEEE International Symposium on High Performance Computer Architecture (HPCA)*. IEEE, 64–76.

Thorsten Hansen, Martin Giesel, and Karl R Gegenfurtner. 2008. Chromatic discrimination of natural objects. *Journal of Vision* 8, 1 (2008), 2–2.

Thorsten Hansen, Lars Pracejus, and Karl R Gegenfurtner. 2009. Color perception in the intermediate periphery of the visual field. *Journal of vision* 9, 4 (2009), 26–26.

Nargess Hassani and Michael J Murdoch. 2016. Color appearance modeling in augmented reality. In *Proceedings of the ACM Symposium on Applied Perception*. 132–132.

Katherine L Hermann, Shridhar R Singh, Isabelle A Rosenthal, Dimitrios Pantazis, and Bevil R Conway. 2021. Temporal dynamics of the neural representation of hue and luminance contrast. *BioRxiv* (2021), 2020–06.

Yuge Huang, En-Lin Hsiang, Ming-Yang Deng, and Shin-Tson Wu. 2020. Mini-LED, Micro-LED and OLED displays: Present status and future perspectives. *Light: Science & Applications* 9, 1 (2020), 1–16.

Anton S Kaplanyan, Anton Sochenov, Thomas Leimkühler, Mikhail Okunev, Todd Goodall, and Gizem Rufo. 2019. DeepFovea: neural reconstruction for foveated rendering and video compression using learned statistics of natural videos. *ACM Transactions on Graphics (TOG)* 38, 6 (2019), 1–13.

Minjung Kim, Maryam Azimi, and Rafal K Mantiuk. 2021. Color threshold functions: Application of contrast sensitivity functions in standard and high dynamic range color spaces. *Electronic Imaging* 2021, 11 (2021), 153–1.

Brooke Rajcianich, Petr Kellnhofer, and Gordon Wetzstein. 2021. A Perceptual Model for Eccentricity-dependent Spatio-temporal Flicker Fusion and its Applications to Foveated Graphics. *ACM Trans. Graph.* 40 (2021). Issue 4.

John Krauskopf and Gegenfurtner Karl. 1992. Color discrimination and adaptation. *Vision research* 32, 11 (1992), 2165–2175.

James Larimer, Carol M Cicerone, et al. 1974. Opponent-process additivity—I: Red/green equilibria. *Vision Research* 14, 11 (1974), 1127–1140.

James Larimer, David H Krantz, and Carol M Cicerone. 1975. Opponent process additivity—II: Yellow/blue equilibria and nonlinear models. *Vision research* 15, 6 (1975), 723–731.

Yue Leng, Chi-Chun Chen, Qiuyue Sun, Jian Huang, and Yuhao Zhu. 2019. Energy-efficient video processing for virtual reality. In *Proceedings of the 46th International Symposium on Computer Architecture*. 91–103.

Yang Li, Alexandre Chapiro, Mehmet N. Agaoglu, Nicolas Pierre Marie Frederic Bonnier, Yi-Pai Huang, Chaohao Wang, Andrew B. Watson, and Pretesh A. Mascarenhas. 2022. PERIPHERAL LUMINANCE OR COLOR REMAPPING FOR POWER SAVING.

David L MacAdam. 1942. Visual sensitivities to color differences in daylight. *JOSA* 32, 5 (1942), 247–274.

Rafal K Mantiuk, Gyorgy Denes, Alexandre Chapiro, Anton Kaplanyan, Gizem Rufo, Romain Bachy, Trisha Lian, and Anjul Patney. 2021. FovVideoDVP: A visible difference predictor for wide field-of-view video. *ACM Transactions on Graphics (TOG)* 40, 4 (2021), 1–19.

Michael Mauderer, David R Flatla, and Miguel A Nacenta. 2016. Gaze-contingent manipulation of color perception. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*. 5191–5202.

Andrew B Metha, Algis J Vingrys, and David R Badcock. 1994. Detection and discrimination of moving stimuli: the effects of color, luminance, and eccentricity. *JOSA A* 11, 6 (1994), 1697–1709.

Microsoft. 2021. About HoloLens 2. <https://docs.microsoft.com/en-us/hololens/hololens2-hardware>.

ME Miller, RS Cok, AD Arnold, and MJ Murdoch. US Patent 7,230,594, Jun. 12, 2007. Color OLED display with improved power efficiency.

ME Miller, MJ Murdoch, RS Cok, and AD Arnold. US Patent 7,333,080, Feb. 19, 2008. Color OLED display with improved power efficiency.

Michael E Miller, Michael J Murdoch, John E Ludwicki, and Andrew D Arnold. 2006. P-73: Determining Power Consumption for Emissive Displays. In *SID Symposium Digest of Technical Papers*, Vol. 37. Wiley Online Library, 482–485.

Michael J Murdoch, Mariska GM Stokkermans, and Marc Lambooij. 2015. Towards perceptual accuracy in 3D visualizations of illuminated indoor environments. *Journal of Solid State Lighting* 2, 1 (2015), 1–19.

Dirk V Norren and Johannes J Vos. 1974. Spectral transmission of the human ocular media. *Vision research* 14, 11 (1974), 1237–1244.

Anjul Patney, Marco Salvi, Joohwan Kim, Anton Kaplanyan, Chris Wyman, Nir Benty, David Luebke, and Aaron Lefohn. 2016. Towards Foveated Rendering for Gaze-Tracked Virtual Reality. *ACM Trans. Graph.* 35, 6, Article 179 (Nov. 2016), 12 pages. <https://doi.org/10.1145/2980179.2980246>

Rui Peng, Mingkai Cao, Qiyan Zhai, and Ming Ronnier Luo. 2021. White appearance and chromatic adaptation on a display under different ambient lighting conditions. *Color Research & Application* 46, 5 (2021), 1034–1045.

Ernst Pöppel and Lewis O Harvey. 1973. Light-difference threshold and subjective brightness in the periphery of the visual field. *Psychologische Forschung* 36, 2 (1973), 145–161.

Parthasarathy Ranganathan, Erik Geelhoed, Meera Manahan, and Ken Nicholas. 2006. Energy-aware user interfaces and energy-adaptive displays. *Computer* 39, 3 (2006), 31–38.

Theresa-Marie Rhyne, Nicholas Bazarian, Jose Echevarria, Michael J Murdoch, and Danielle Feinberg. 2018. Color mavens advise on digital media creation and tools: SIGGRAPH 2018 panel. In *ACM SIGGRAPH 2018 Panels*. 1–2.

Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, et al. 2015. Imagenet large scale visual recognition challenge. *International journal of computer vision* 115, 3 (2015), 211–252.

Peter H Schiller and Nikos K Logothetis. 1990. The color-opponent and broad-band channels of the primate visual system. *Trends in neurosciences* 13, 10 (1990), 392–398.

Fred Schlachter. 2013. No Moore's Law for batteries. *Proceedings of the National Academy of Sciences* 110, 14 (2013), 5273–5273.

Donghwa Shin, Younghyun Kim, Naehyuck Chang, and Massoud Pedram. 2013. Dynamic driver supply voltage scaling for organic light emitting diode displays. *IEEE Transactions on Computer-Aided Design of integrated circuits and systems* 32, 7 (2013), 1017–1030.

Alex Shye, Benjamin Scholbrock, and Gokhan Memik. 2009. Into the wild: studying real user activity patterns to guide power optimizations for mobile architectures. In *Proceedings of the 42nd annual IEEE/ACM international symposium on microarchitecture*. 168–178.

Vivianne C Smith and Joel Pokorny. 1975. Spectral sensitivity of the foveal cone photopigments between 400 and 500 nm. *Vision research* 15, 2 (1975), 161–171.

Hongxin Song, Toco Yuen Ping Chui, Zhangyi Zhong, Ann E Elsner, and Stephen A Burns. 2011. Variation of cone photoreceptor packing density with retinal eccentricity and age. *Investigative ophthalmology & visual science* 52, 10 (2011), 7376–7384.

Andrew Stockman and Lindsay T Sharpe. 2000. The spectral sensitivities of the middle- and long-wavelength-sensitive cones derived from measurements in observers of known genotype. *Vision research* 40, 13 (2000), 1711–1737.

Qi Sun, Fu-Chung Huang, Joohwan Kim, Li-Yi Wei, David Luebke, and Arie Kaufman. 2017. Perceptually-Guided Foveation for Light Field Displays. *ACM Trans. Graph.* 36, 6, Article 192 (Nov. 2017), 13 pages. <https://doi.org/10.1145/3130800.3130807>

Qi Sun, Fu-Chung Huang, Li-Yi Wei, David Luebke, Arie Kaufman, and Joohwan Kim. 2020. Eccentricity effects on blur and depth perception. *Optics express* 28, 5 (2020), 6734–6739.

Takatoshi Tsujimura. 2017. *OLED display fundamentals and applications*. John Wiley & Sons.

Okan Tarhan Tursun, Elena Arabadzhyska-Koleva, Marek Wernikowski, Radosław Mantiuk, Hans-Peter Seidel, Karol Myszkowski, and Piotr Didyk. 2019. Luminance-contrast-aware foveated rendering. *ACM Transactions on Graphics (TOG)* 38, 4 (2019), 1–14.

Algis J Vingrys and Luke E Mahon. 1998. Color and luminance detection and discrimination asymmetries and interactions. *Vision research* 38, 8 (1998), 1085–1095.

David R Walton, Rafael Kuffner Dos Anjos, Sebastian Friston, David Swapp, Kaan Akşit, Anthony Steed, and Tobias Ritschel. 2021. Beyond blur: Real-time ventral metamers for foveated rendering. *ACM Transactions on Graphics* 40, 4 (2021), 1–14.

Rui Wang, Bowen Yu, Julio Marco, Tianlei Hu, Diego Gutierrez, and Hujun Bao. 2016. Real-time rendering on a power budget. *ACM Transactions on Graphics (TOG)* 35, 4 (2016), 1–11.

Stephen Westland, Caterina Ripamonti, and Vien Cheung. 2012. *Computational colour science using MATLAB*. John Wiley & Sons.

Wisecoco. 2022. 3.81 inch OLED AMOLED Display . <https://www.aliexpress.com/item/33015630007.html>.

William David Wright. 1929. A re-determination of the trichromatic coefficients of the spectral colours. *Transactions of the Optical Society* 30, 4 (1929), 141.

Hao Xie, Susan P Farnand, and Michael J Murdoch. 2020. Observer metamericism in commercial displays. *JOSA A* 37, 4 (2020), A61–A69.

Zhisheng Yan, Chen Song, Feng Lin, and Wenyao Xu. 2018. Exploring Eye Adaptation in Head-Mounted Display for Energy Efficient Smartphone Virtual Reality. In *Proceedings of the 19th International Workshop on Mobile Computing Systems & Applications* (Tempe, Arizona, USA) (*HotMobile '18*). Association for Computing Machinery, New York, NY, USA, 13–18. <https://doi.org/10.1145/3177102.3177121>

Lili Zhang, Michael J Murdoch, and Romain Bachy. 2021a. Color appearance shift in augmented reality metameric matching. *JOSA A* 38, 5 (2021), 701–710.

Yunjin Zhang, Marta Ortín, Victor Arellano, Rui Wang, Diego Gutierrez, and Hujun Bao. 2018. On-the-Fly Power-Aware Rendering. In *Computer Graphics Forum*, Vol. 37. Wiley Online Library, 155–166.

Yunjin Zhang, Rui Wang, Yuchi Huo, Wei Hua, and Hujun Bao. 2021b. PowerNet: Learning-based Real-time Power-budget Rendering. *IEEE Transactions on Visualization and Computer Graphics* (2021).