32nd Color and Imaging Conference 2024 (CIC32)

Color Science and Engineering Systems, Technologies, and Applications

Montreal, Canada
28 October - 1 November 2024

ISBN: 979-8-3313-1392-0

Printed from e-media with permission by:

Curran Associates, Inc.
57 Morehouse Lane
Red Hook, NY 12571

Some format issues inherent in the e-media version may also appear in this print version.

Copyright© (2024) by Society for Imaging Science & Technology
All rights reserved.

Printed with permission by Curran Associates, Inc. (2025)

For permission requests, please contact Society for Imaging Science & Technology
at the address below.

Society for Imaging Science & Technology
7003 Kilworth Lane
Springfield, Virginia 22151
USA

Phone: 703-642-9090
Fax: 703-642-9094

info@imaging.org

Additional copies of this publication are available from:

Curran Associates, Inc.
57 Morehouse Lane
Red Hook, NY 12571 USA
Phone: 845-758-0400
Fax: 845-758-2633
Email: curran@proceedings.com
Web: www.proceedings.com

TECHNICAL PROGRAM: SCHEDULE AND ABSTRACTS

WEDNESDAY 30 OCTOBER

WELCOME AND OPENING KEYNOTE

Session Chair: Carol Payne, Apple Inc. (US)
09:00 – 10:00

Welcome Remarks
Suzanne Grinnan, IS&T executive director (US), and Javier Vázquez Corral, CIC32 general chair, Universitat Autònoma de Barcelona (Spain)

How Television Determined the Look of Modern Cinema, *Joseph Goldstone, Lilliputian Pictures LLC (US)* [Abstract only]
A term rarely heard today is "videographer"; the IMDB entry for "Emily in Paris" lists its image creators under "Series Cinematography by". Broadcast, streamed, and cinematic content is now (mostly) created and edited with the same tools and techniques—but that wasn't always so. For a long time, improvements in television acted as a goad, compelling corresponding innovations in cinema; each time television "cleared the bar", cinema tried to raise it—but this too has changed. Nowadays, within the cadre of those creating moving imagery, little stigma remains to the act of creating for "the small screen" rather than for a movie theater.

This talk is about the historic antipathy of the television and movie industries; the innovative push by cinema to one-up advances in television's "look"; and how this antipathy is fading as their technologies of acquisition and content manipulation converge. The "film look" being applied to television production now is really a reversal of this traditional creative tension. This talk discusses landmarks in the motion picture industry's evaluation of and adaptation to digital projection, digital capture, and digital delivery, as manifested in equipment, standards, and workflow, and concludes with observations as to what other industries (such as gaming) take—and reject—from the looks of television and film.

COLOR THEORY

10:00 – 10:40
Session Chair: Alex Forsythe, Academy of Motion Picture Arts and Sciences (US)

10:00 **Colour Meets Geometry in Colorimetric Filter Design,** *Graham D. Finlayson, University of East Anglia (UK), and Ivar Farup, Norwegian University of Science and Technology (Norway)* 1
Panoramas are formed by stitching together two or more images of a scene viewed from different positions. Part of the solution to this stitching problem is 'solving for the homography': where the detail in one image is geometrically warped so it appears in the coordinate frame of another. In this paper, we view the spectral loci for a given camera and the human visual system (i.e. their respective chromaticity diagrams) as two pictures of the same 'scene' and warp one to the other by finding the best homography. When this geometric distortion renders the two loci to be identical then there exists a unique colour filter (that falls gracefully from the derivation without further calculation) which makes the camera colorimetric (the camera+filter measures RGBs that are exactly linearly related to XYZs). When the best homography is not exact the filter derived by this method still makes cameras approximately colorimetric. Experiments validate our method.

10:20 **CIE Colour Matching Functions and Cone Fundamentals: Problems and Modifications,** *Cheng Gao[1], Kaida Xiao[2], Michael R. Pointer[2], and Changjun Li[1]; [1]University of Science and Technology Liaoning (China) and [2]University of Leeds (UK)* 8
CIE colorimetry based on colour matching functions has been successfully applied in various industrial applications. In the past it was generally accepted that the chromaticity diagram based on either 1931 or 1964 colour matching functions (CMFs), contains all the chromaticity coordinates of stimuli, which means the domain (Ω) enclosed by the spectrum locus and the purple line is convex. Also based on the chromaticity diagram, the Helmholtz coordinates (dominant wavelength and excitation purity) can be defined. In this paper, these properties of chromaticity coordinates based on CIE 1931, CIE 1964, CIE 2006 2- and 10-degree cone fundamentals (CFs) and cone fundamentals-based CMFs are evaluated. It is found that the domain Ω does not contain the chromaticity coordinates of all stimuli, and spectral chromaticity coordinates do not distribute in the wavelength order along the spectrum locus, which results in no way to determine the Helmholtz coordinates for certain stimuli. Finally modified CIE 1931 CMFs, CIE 1964 CMFs, and CFs and CMFs based on CFs, are derived.

COFFEE BREAK AND EXHIBIT

10:40 – 11:20

APPLICATIONS, LIGHTING, AND MOTION PICTURE

11:20 – 12:00
Session Chair: Jérémie Gerhardt, CAE (Canada)

11:20 **Assessing Color-rendering Differences in Cinematic Lighting with a New Metamer Mismatch Metric,** *Alexander Forsythe[1,2], and Brian Funt[1]; [1]Simon Fraser University (Canada) and [2]Academy of Motion Picture Arts and Sciences (US)* 13
Consistent color-rendering in motion pictures is critical for creating natural scenes that enhance storytelling and don't distract the audience's attention. In today's production environments, it is common to use a wide variety of light sources. Traditional tungsten-halogen sources—red, green, and blue light emitting diode (RGB LED) sources and white light LEDs—are often mixed, leading to color-rendering issues. This paper introduces a new metric, the Camera Lighting Metamer Index (CLMI), rooted in the concept of metamer mismatching. The CLMI is for assessing the color-rendering differences of disparate sources when a single camera is used and the camera's spectral sensitivities are known. By leveraging the known spectral sensitivities of the camera and the spectral power distributions (SPDs) of the light sources, CLMI quantifies the potential for color discrepancies between objects lit by the different sources. We propose that this metric can serve as a useful tool for cinematographers and visual effects artists, providing more predictable and precise control over color fidelity. The metric could also be used to supplement existing generalized metrics, such as Spectral Similarity Index (SSI), when detailed camera and light source spectral characteristics are available.

11:40 **A Color Characterization Model for APL Dependent OLED Displays,** *Pooshpanjan Roy Biswas, Thibault Cabana, and Adrien Carmone, DXOMARK (France)* **17**

Several studies in the past have proposed models to characterize the colorimetry of displays, most of which have poor performance for OLED displays. This is primarily due to the dependency of the colorimetry of OLED panels on the Average Pixel Level (APL) of the content displayed on them. In this study, a workflow is proposed to characterize the colorimetry of APL dependent OLED panels based on the power consumption of actual pixel-content of the displayed scene. The method performed well with a mean of mean CIEDE2000 of 19 natural images as 2.18 units.

3-MINUTE EXHIBITOR PREVIEWS

12:00 – 12:10

Session Chair: Carol Payne, Apple Inc. (US)

Learn more about our exhibitors, Image Engineering and ThousLite.

INTERACTIVE PAPER PREVIEWS I

12:10 – 12:25

Session Chair: Raquel Gil-Rodriguez, Justus-Liebig-Universität Giessen (Germany)

Evaluating the Impact of Display Light Settings on Circadian Rhythms, Visual Fatigue, and Cognitive Performance: A Comparative Study of Static and Dynamic Backgrounds, *Zhenzhen Li, Ming Ronnier Luo, Yuechen Zhu, and Qichen Ye, Zhejiang University (China)* **23**

As displays become ubiquitous and increasingly integrated into daily life, their impact on human health is a major concern for academia and industry. The aim of this study was to investigate the effect of display backgrounds with different Correlated Colour Temperature (CCT) and Circadian Stimulus (CS) settings on human circadian rhythms and visual fatigue. Twelve participants underwent four 9-hour display lighting interventions over a 10-day period, including S1 (CCT at 4000K; CS from 0.29 to 0.15), S2 (CS at 0.2; CCT from 6500K to 4000K), S3 (CCT from 6500K to 4000K; CS from 0.30 to 0.15), and static S4 (CCT at 4000K; CS at 0.2). Participants' melatonin levels, visual fatigue, cognitive performance, sleep quality and 24-hour core body temperature were monitored. The results showed that S4 was the most circadian-friendly condition, with the least visual fatigue and the best sleep quality. In addition, the S3 intervention resulted in the lowest nighttime alertness. Therefore, static display backgrounds with low CCT and CS appear to be more beneficial for circadian health than dynamic display backgrounds. Furthermore, the results of several statistical tests showed that CS has a greater effect on rhythm than CCT.

Estimating Spectral BRDF Parameters Using Handheld Devices, *Bita Panahi, Aditya Sole, and Ivar Farup, Norwegian University of Science and Technology (Norway)* **29**

The Bidirectional Reflectance Distribution Function (BRDF) is one of the tools for characterising the appearance of real-world materials. However, bidirectional reflectance measurements and data processing can be time-consuming and challenging. This paper aims to estimate the BRDF values of eight matt samples using two portable, handheld devices, one for diffuse reflectance and one for specular reflectance measurements. The data is fitted to the Cook-Torrance BRDF model in the spectral domain to get the optimised parameters and the estimated spectral BRDF values using three different cost functions. The estimated BRDF is evaluated using a colour-difference metric. The results show that it would be possible to estimate spectral BRDF of a sample using measurements from two simple measurement devices having fewer angle combinations for both the diffuse and specular measurements. This results in a shorter measurement and processing time, lower storage usage, and estimations of spectral BRDF values. Moreover, the cube root cosine-weighted RMSE cost function shows more consistency in the colour reproduction estimated by the fitted BRDF model.

Color Constancy: Color Interaction Between Local Surround and Illumination in Virtual Reality Scenarios, *Raquel Gil Rodríguez[1], Laysa Hedjar[1], Matteo Toscani[2], Dar'ya Guarnera[3], Claudio Giuseppe Guarnera[3], and Karl Reiner Gegenfurtner[1]; [1]Justus-Liebig University Giessen (Germany), [2]Bournemouth University (UK), and [3]University of York (UK)* . . . **Appendix A-1**

In our study, we used virtual reality to examine how the color of an object's surroundings affects color constancy. We manipulated lighting and object properties in computer-generated scenes using Unreal Engine with five different light sources, together with the virtual reality HTC Vive Pro Eye headset for visualization. Participants assessed color constancy by matching the target to a neutral reference among five differently colored objects in the scene. Our findings showed that color constancy performance drops significantly when considering illuminant colors in the opposite direction of the local surround color, revealing the interactive effects of surround color and illumination.

CNP2D: The Chromatic Noise Pattern Discrimination Dataset, *Trevor Dalton Canham[1], Peter Morovic[2], Richard Murray[1], Ján Morovic[3], and Michael Brown[1]; [1]York University (Canada), [2]HP Inc. (Spain), and [3]HP Inc. (UK)* **35**

Pixel-wise color difference metrics like ΔE_{00} have long been used in image analysis, but it remains unclear how scores should be integrated over space. To highlight this, a psychophysical experiment was conducted to characterize visual sensitivity to differences in chromatic noise patterns in different color and pattern contexts. The results demonstrated that observers were more sensitive to chromatic noise pattern (CNP) differences when similar colors were spatially dispersed over the pattern as opposed to clustered. Further analysis with common image color and texture difference metrics showed that none were sensitive to this effect. This finding highlights the need for metrics which capture the perceptual interaction between color and texture.

Corresponding Colors in Virtual Reality: A Proof of Concept, *Nicoletta Prencipe, Mikki Hiltunen, Lauri Klemettilä, Evan G. Center, Matti Pouke, Anna Yershova, Timo Ojala, and Steven M. LaValle, University of Oulu (Finland)* **41**

Corresponding colors data are necessary to evaluate and improve chromatic adaptation models. Virtual reality (VR) technologies offer a promising and flexible tool for studying chromatic adaptation. This work revisits the classic technique of haploscopic matching to collect corresponding colors data using VR. A virtual environment was designed in which the same visual scene was lit with a light of a different chromaticity for each eye, and an exploratory study was conducted, collecting corresponding colors data from 10 participants (5 color stimuli, 3 trials for each). The averaged CIEDE2000 standard deviation for individual observers, across 5 colors, was 2.94, showing adequate precision despite the low amount of repetitions. This work represents a proof of concept for an efficient and realistic VR haploscopic matching paradigm, which may be extended by future studies employing display color characterization and greater numbers of participants and trials.

©2024 Society for Imaging Science and Technology

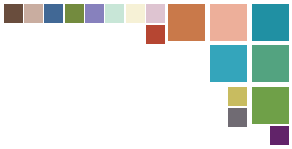

Neural Image Compression and the Degradation of Color: An Analysis, *Ian MacPherson, Trevor Canham, and Michael S. Brown, York University (Canada)* 47

Neural image compression employs deep neural networks and generative models to achieve impressive compression rates and reconstruction qualities compared to traditional signalprocessing- based compression algorithms such as JPEG. However, color artifacts that arise in an image as the amount of compression increases have not been formally analyzed for neuralbased compression algorithms. This paper provides an initial investigation into the degradation of color when images are compressed at comparable bit rates using lossy neural image compression and variants of JPEG. Our findings indicate that neural image compression degrades color more gracefully than JPEG, JPEG 2000, and JPEG XL.

GROUP LUNCH

12:25 – 14:00

Join other attendees for a buffet lunch in Salon Pierre de Coubertin.

HDR: A CHALLENGING NEW DIMENSION IN PHOTOGRAPHY

14:00 – 15:30

Session Chair: Nicolas Bonnier, Apple Inc. (US)

14:00 **HDR Image Review,** *Paul Hubel, Apple Inc. (US)* [Abstract only]

High dynamic range images allow a greater range of shadow and highlight detail to be conveyed, including sufficient separation of diffuse white and specular highlights. Walking through a portfolio of images, this talk demonstrates how HDR images, produced for high dynamic range displays, provide a better user experience than standard dynamic range images.

14:15 **Building the Foundations of a Standardized HDR Photography Ecosystem,** *Nicolas Bonnier, Apple Inc. (US)* [Abstract only]

After a quick introduction of key HDR concepts and existing HDR video standards, this talk discusses how it the evolution began with building standards for storing HDR photos, then pivoted to building a headroom management workflow, to ensure proper representation of these HDR photos on any display, in any viewing environment.

14:45 **A New HDR Workflow for Creators,** *Eric Chan, Adobe Inc., and Paul Hubel, Apple Inc. (US)* [Abstract only]

This talk presents the theory and practice of ISO 21496-1 HDR Gain Maps with a complete end-to-end HDR workflow demonstration, including capture, edit, export, and sharing, as well as Adobe Lightroom's HDR editing tools and its use of ISO 21496-1 Gain Maps.

15:15 **Research Challenges in HDR Imaging,** *Paul Hubel, Apple Inc. (US)* [Abstract only]

This talk reviews key HDR imaging concepts and raises several open questions for the color imaging scientific community to consider for future research and development.

COFFEE BREAK / EXHIBIT / HDR SYSTEM DEMONSTRATIONS

15:30 – 16:10

During the break, those interested can view various HDR demonstrations related to the HDR: A Challenging New Dimension in Photography session.

USING MODERN DISPLAY TECHNOLOGIES IN SCIENCE AND R&D

16:10 – 17:30

Session Chair: Timo Kunkel, Dolby Laboratories, Inc. (US)

16:10 **Understanding and Harnessing the Potential of Modern Display Technologies,** *Timo Kunkel, Dolby Laboratories, Inc. (US)* [Abstract only]

Over the past decade, displays have evolved significantly. They offer HDR with widened color gamut and dynamic range, higher peak luminance levels and deeper blacks, as well as increased frame rates and spatial resolution. While mainly designed for content consumption, these display improvements are also enticing to scientists and engineers that are interested in running psychovisual studies, using these extended capabilities to exhibit stimuli that previously were not easily achievable. The challenge in using modern displays in the context of scientific experimentation is that these displays are designed to show entertainment content where variability in aspects such as peak luminance are of lesser concern. However, if experiment test targets are not exactly rendered as intended, results might be biased. How to alleviate such challenges, avoid pitfalls and harness the potential of modern displays will be discussed in this session.

16:30 **Can I Trust what I See? Possibilities and Potential Pitfalls when Using Modern Displays with Research and Development Projects: A Panel Discussion**

Moderator: Timo Kunkel, Dolby Laboratories, Inc. (US)

Panelists:

Nate McFarlin, Dolby (US)
Richard Murray, York University (Canada)
Christopher Reidy, Meta (US)
Lili Zhang, Meta (US)

DINNER ON OWN

17:30 – 20:00

EVENING TALK

20:00 – 21:00

Session Chair: Paul Hubel, Apple Inc. (US)

The Intersection of Creativity and Technology: HDR in Fashion and Beauty Photography

Sarah Silver, photographer [Abstract only]

Photographer Sarah Silver is known for pushing boundaries with her fashion and beauty work, but the SDR range often undermines her ability to reach her full artistic vision. Using her recent iPhone fashion and beauty HDR photoshoot as the backdrop, Silver will illustrate how HDR technology expands creative possibilities, solves traditional photographic limitations, and opens up a bright and vivid new world within the genre of fashion and beauty image making.

THURSDAY 31 OCTOBER

THURSDAY KEYNOTE

09:00 – 10:00

Session Chair: Javier Vázquez Corral, Universitat Autònoma de Barcelona (Spain)

Speeding Up Creativity through Generative Imaging,
Yannick Hold-Geoffroy, Adobe (US) [Abstract only]

Recent advancements in computational imaging, particularly in generative imaging, have had a significant impact on the work of multimedia artists, introducing new and unexpected workflows to accelerate artistic processes. We are currently experiencing a time where both technical professionals and artists, as well as the general public, are adapting to a new reality marked by the availability of on-demand content and generation of masterpieces at our fingertips. This presentation delves into the recent developments and breakthroughs in image generation, editing, and lighting capture, along with their practical applications. New possibilities for creating realistic extrapolations of images and their role in image editing are discussed. A goal of the talk is for listeners to gain insights into lighting and image generation, as seen through the eyes of users.

COLOR IN COMPUTER VISION

10:00 – 10:40

Session Chair: Ian MacPherson, York University (Canada)

10:00 **Data-Driven Light Source Selection for Camera Colorimetric Calibration,** *Yuyang Liu and Minchen Wei, The Hong Kong Polytechnic University, and Xinchao Qu and Tao Hu, Dajiang Innovations Technology Co., Ltd. (Hong Kong)* **51**

The colorimetric calibration of cameras are critical in imaging systems, with the sources used in light booths being widely used in practice. These sources, however, may not good presentations of the sources in real life, which possibly results in poor colors. In this study, we adopted a genetic algorithm and a large dataset of real light sources to identify an optimal set of sources that can better represent the sources in real life. The experiment results suggested that the identified set of sources can result in better color performance. Moreover, the selection of the sources was much less complicated in comparison to manual selections, which can be considered and implemented in practice.

10:20 **Intrinsic-GS: Multi-view Intrinsic Image Decomposition Using Gaussian Splatting and Color-Invariant Priors,** *Xiaoyan Xing[1], Konrad Groh[2], Sezer Karaoglu[1], and Theo Gevers[1]; [1]University of Amsterdam (the Netherlands) and [2]Bosch (Germany)* **56**

Despite significant advancements in single-view intrinsic image decomposition, a domain disparity exists due to the limited information that can be obtained from a single-view image andthe ill-posed nature of the problem of intrinsic image decomposition. Multi-view images offer an alternative method to circumvent the ambiguity present in 2D intrinsic image decomposition. Building on the concepts of multi-view intrinsic images and recent neural rendering techniques, we propose Intrinsic-GS, a Multiview intrinsic image decomposition method utilizing Gaussiansplatting. To achieve this, we first augment each Gaussian ellipsoid with additional attributes (i.e., albedo, shading, and a residual term) to model the intrinsic radiance field. Next, we use several color-invariants and physics-based priors to jointly regularize the optimization of the intrinsic and composited radiance fields. Finally, we conduct experiments on both synthetic and real-world datasets, demonstrating stable intrinsic decomposition results across various (including non-Lambertian) objects and scenes.

COFFEE BREAK AND EXHIBIT

10:40 – 11:20

COLOR PERCEPTION AND COGNITION I

11:20 – 12:20

Session Chair: Romain Bachy, Meta Reality Labs (US)

11:20 **JPI-first: Visualizing Uncertainty with Simulated Chromatic Aberration,** *Rashidul Islam and Stephen Brooks, Dalhousie University (Canada)* . . . . . **see JPI 7/**
DOI: 0.2352/J.Percept.Imaging.2024.7.000403

The area of uncertainty visualization attempts to determine the impact of alternative representations and evaluate their effectiveness in decision-making. Uncertainties are often an integral part of data, and model predictions often contain a significant amount of uncertain information. In this study, we explore a novel idea for a visualization to present data uncertainty using simulated chromatic aberration (CA). To produce uncertain data to visualize, we first utilized existing machine learning models to generate predictive results using public health data. We then visualize the data itself and the associated uncertainties with artificially spatially separated color channels, and the user perception of this CA representation is evaluated in a comparative user study. From quantitative analysis, it is observed that users are able to identify targets with the CA method more accurately than the comparator state-of-the-art approach. In addition, the speed of target identification was significantly faster in CA as compared to the alternative, but the subjective preferences of users do not vary significantly between the two.

11:40 **Modelling Contrast Matching Across Luminance Levels,** *Maliha Ashraf and Rafal K. Mantiuk, University of Cambridge (UK)* **64**

The study investigates the modelling of contrast matching functions (SCMF) across various luminance levels, addressing the nonlinear behaviour of the human visual system in perceiving suprathreshold contrasts. Using a comprehensive dataset of contrast matching experiments involving younger and older observers, the research tests existing models and proposes a new hybrid model. The additive model by Kulikowski (1976) and the multiplicative model inspired by Peli *et al.* (1996) are evaluated against the dataset, revealing their limitations in predicting contrast matching across a broad luminance range. A novel model combining additive and multiplicative elements is introduced, accounting for threshold ratios and differences, and optimised using regression analysis. The proposed model demonstrates superior prediction accuracy, particularly for achromatic contrasts at extreme luminance levels, and holds potential applications in image processing, particularly for high dynamic range (HDR) content adaptation across different luminance conditions.

12:00 **Individual Color Matching Functions and Application in Cross-media Color Reproduction,** *Siyuan Song[1], Ming Ronnier Luo[1], Tingwei Huang[2], Andrew Rider[3], and Andrew Stockman[1,3]; [1]Zhejiang University (China), [2]THOUSLITE Light Lighting (Changzhou) Ltd. (China), and [3]University College London Institute of Ophthalmology (UK)* **70**

In this study, two experiments were conducted to assess the importance of using individual color matching functions (CMFs) for cross-media color

©2024 Society for Imaging Science and Technology

reproduction. Firstly, five observers used a visual trichromator to perform a series of color matches. Using the extended CIEPO06 model, these matches were used to derive estimates of each individual's CMFs. Each observer repeated the experiment five times to assess intra-observer variability. Subsequently, the same observers performed cross-media color reproduction experiments using a cross-media color matching system. The findings indicated that the individual CMFs can provide much more accurate predictions of the visual results than the standard CIE CMFs. Thus, individual CMFs can greatly enhance the accuracy of color reproduction.

INTERACTIVE PAPER PREVIEWS II

12:00 – 12:35

Session Chair: Raquel Gil-Rodriguez, Justus-Liebig-Universität Giessen (Germany)

White point chromaticities play an important role in display specifications and calibrations, with the D65 chromaticities widely used in various standards. Such a D65 chromaticities are calculated and specified with the CIE 1931 2° color matching functions (CMFs). Though it is well known that wide color gamut displays always introduce color shifts (i.e., the same D65 chromaticities result in shifted color appearance) and other CMFs were found to have better performance, there is no guidelines about how the D65 white point chromaticities should be shifted. A color matching experiment was carried out in this study, in which 12 displays were used to match the color appearance of a broadband D65 reference. The results can be used as a reference when defining the D65 white point chromaticities using different CMFs.

In modern image signal processors (ISPs), many modules have adopted deep neural networks (DNNs). This study explores whether a single DNN can effectively replace both the auto white balance (AWB) and denoising modules, or if they should be processed separately. Our experiment results suggest that performing AWB and denoising individually can produce better performance than an end-to-end approach. Moreover, processing denoising before AWB leads to a significant improvement, with an increase of nearly 6 dB in PSNR and 30% reduction in mean angular error (MAE). These findings suggest that careful consideration of the processing order in ISP pipelines can lead to substantial enhancements in image quality.

Stable diffusion is a generative algorithm for creating images from text prompts. This paper explores prompts with color terms and proposes a process to generate, visualize and assess these results. Automated prompts are used to generate and render a color term, an object and a context. The results are then evaluated using two dashboard views of the underlying images. First is a sampling based on a collection of frequently used color terms. Second is a sampling by object prompts, such as apples and boxes. This paper considers the following questions: how effectively are the colored objects generated? how do the colors generated by stable diffusion compare to human color naming? How might color terms be useful in visualizing properties and features of generative algorithms? The dashboard view of color terms suggests that less frequently used color terms may be generated less consistently. In addition, even the most common color terms can fail to be correctly generated. Likewise, objects with more frequent color associations, such as apples or pumpkins, will result in less accurate color generation.

Lighting is a crucial component of an image, especially during the task of virtual object insertion in a photograph, e.g. for AR/VR/MR applications. The human brain is quite attuned to changes in lighting for a given object. Thus, it is imperative to estimate the scene's lighting in order to produce a realistic image. However, this task remains complex, as disentangling lighting from the 3D geometry and the material properties of the objects in the image is an ill-posed problem. For this reason, the community has been developing lighting estimation methods for over the past decades, by using handcrafted priors and more recently leveraging the power of deep learning. The great variety of lighting estimation methods currently available must be compared to each order in order to quantify their progress with regard to their accuracy and realism of their estimations.

For this task, it is standard to use popular image quality assessment (IQA) metrics to compare a rendered virtual object, lit using the predicted lighting from different methods, compared to the ground truth render. However, standard IQA metrics are not designed to quantify differences in lighting, since they are usually developed for other specific tasks (such as noise perception for compression). Thus, it is unclear if standard IQA metrics are appropriate to use when judging the perceptual quality of renders generated with estimated lighting.

In this work, we evaluate whether IQA metrics and human perception align. To do so, we perform a calibrated user study, which allows us to compare the preferences of humans with standard IQA metrics. We demonstrate that they are not in agreement; hence we propose our new IQA metric for lighting estimation, which is in agreement with the perceptual data. Our new perceptual IQA metric shows great generalisation to other lighting estimation methods not included in our dataset, meaning that it will be helpful for the development of new lighting estimation methods. To encourage future research, all (anonymised) perceptual data and code are available at https://lvsn.github.io/PerceptionMetric/ .

Branded athletic apparel is a multi-billion dollar business and it continues to expand year over year. Sports teams and their fans expect brand accuracy and durability across all types of textile-based products from clothing to home decor. Color is a vital part of a brand's identity and printing consistent, durable, and accurate brand colors across different materials can be challenging for manufacturers. This study examines the accuracy and durability of two brand colors, specified by the brand as Pantone and CMYK values and measured as LAB values using a spectrophotometer, and printed on three textiles commonly used for clothing. The study examined the three most common printing processes for manufacturing apparel–dye sublimation, direct-to-garment (DTG), and screen printing. Study results show that the screen printing process is able to provide the most accurate color replication on any

of the three textiles when compared to the target brand color. The DTG process in combination with a polyester or poly-blend material provides the best durability of the colour after the first cycle of laundering has been completed. However, dye sublimation on blend or polyester textiles provides the most color stability after the full five treatment cycles.

Large Size of Color Constancy: Enhancing Pure Color Image Illuminant Estimation with Kolmogorov-Arnold Networks, *LiangWei Chen and Ming Ronnier Luo, Zhejiang University (China), and Minchen Wei, The Hong Kong Polytechnic University (Hong Kong)* **95**

Large efforts have been made to perform illuminant estimation, resulting in the development of various statistical- and learning-based methods. However, there have been challenges for some types of images, such as a single color, referred to as pure color images, which is the focus of the present research.. In this study, the neural network approach is used. It was found the Kolmogorov-Arnold Networks (KAN) model, a novel approach that diverges from traditional Multi-Layer Perceptron (MLP) architectures gave the accurate predictions. Our method, "Large Size Colour Constancy" (LSCC), characterized by its unique neural network structure, achieves high accuracy in illuminant estimation with significantly fewer parameters and enhanced interpretability. Additionally, three new pure color image datasets—"ZJU Color Fabric", "ZJU 0.8 Real Scene", and "ZJU 1.0 Real Scene" were produced—covering a wide range of conditions, including indoor and outdoor environments, as well as natural and artificial light sources. The results showed LSCC method to outperform existing methods across not only the pure colour datasets but also the traditional datasets, including classical normal images. It should offers practical deployment potential due to its efficiency and reduced computational requirements.

Color in Visual-Language Models: CLIP Deficiencies, *Guillem Arias, Ramon Baldrich, and Maria Vanrell, Computer Vision Center / Universitat Autònoma de Barcelona (Spain)* **101**

This work explores how color is encoded in CLIP (Contrastive Language-Image Pre-training) which is currently the most influential VML (Visual Language model) in Artificial Intelligence. After performing different experiments on synthetic datasets created for this task, we conclude that CLIP is able to attribute correct color labels to colored visual stimulus, but, we come across two main deficiencies: (a) a clear bias on achromatic stimuli that are poorly related to the color concept, thus white, gray and black are rarely assigned as color labels; and (b) the tendency to prioritize text over other visual information. Here we prove it is highly significant in color labelling through an exhaustive Stroop-effect test. With the aim to find the causes of these color deficiencies, we analyse the internal representation at the neuron level. We conclude that CLIP presents an important amount of neurons selective to text, specially in deepest layers of the network, and a smaller amount of multimodal color neurons which could be the key of understanding the concept of color properly. Our investigation underscores the necessity of refining color representation mechanisms in neural networks to foster a more comprehensive comprehension of colors as humans understand them, thereby advancing the efficacy and versatility of multimodal models like CLIP in real-world scenarios.

LUNCH ON OWN
12:35 – 14:00

CAPTURE AND REPRODUCTION
14:00 – 15:20

Session Chair: Trevor Canham, York University (Canada)

14:00 **Adaptive Chroma Correction of Tone Mapping Operators for Natural Image Appearance,** *Imran Mehmood, Muhammad Usman Khan, and Ming Ronnier Luo, Zhejiang University (China)* **107**

Preserving perceptual quality of the tone mapped images is one of the major challenges in tone mapping. Most traditional tone mapping operators (TMOs) compress the luminance of high dynamic range (HDR) images without taking account of image color information, resulting into less natural or preferable colors. Current color management algorithms require either manual fine-tuning or introduce lightness and hue shifts. An adaptive color correction model is proposed to address color distortions in tone mapping. It is based on the CIECAM16 to compute perceptual correlates, i.e., Lightness, Chroma and Hue. Regardless of the tone mapping technique, the proposed model recovers natural colors of tone mapped images for spatially invariant and variant operators, making it an effective post-processing technique for color reproduction. Unlike other models, it requires no gamut mapping correction, reproducing more accurate hue, chroma, and lightness. The algorithm was evaluated using objective and subjective methods, revealing that it produced significantly better color reproduction for tone mapped images in terms of naturalness of the colors.

14:20 **Mean Color, Probably,** *Ján Morovic, HP Inc. (Spain), and Peter Morovic, HP Inc. (UK)* **114**

Colors are characterized by their appearance attributes such as brightness, colorfulness and hue, or lightness, chroma and hue when considered relative to an adapted white. Consequently, they can be represented by coordinates in threedimensional spaces like those of CIELAB or CIECAM02 appearance predictors. Therefore, the color of a stimulus is customarily associated with a point in 3D. While this is appropriate when dealing with theoretical quantities, in practice the color space coordinates of a stimulus are subject to variations both in the stimulus itself and in its measurements. This tends to lead to multiple measurements being the basis of identifying the color of a stimulus by averaging them and in some cases excluding outliers. This, however, obscures the fundamental variability of color data. In this paper an alternative, probabilistic approach to dealing with color data will be introduced and applied to the computation of color difference, color gamuts and gamut inclusion, yielding distributions of answers instead of only single values.

14:40 **JIST-first: Monastic Color Reproduction: A Software Tool for Printing and Assessing the Monk Skin Tone Scale,** *Wei-Chung Cheng, US Food and Drug Administration (US)* . . . . . **see JIST 68(5)/ DOI: 10.2352/J.ImagingSci.Technol.2024.68.5.050404**

The Monk Skin Tone (MST) scale, comprising ten digitally defined colors, represents a diverse range of skin tones and is utilised by Google to promote social equity. The MST scale also has the potential to address health disparities in the medical field. However, these colors must be printed as physical charts for health practitioners and researchers to conduct visual comparisons. This study examines the colorimetric characteristics of the MST scale and establishes two per-patch acceptance criteria to determine the acceptability of a printed MST chart using a four-tier grading method. A software tool was developed to assist end users in printing accurate MST charts. The tool was tested with four printer/paper combinations, including inkjet, dye sublimation, and laser printers. Results show that a wide color gamut covering the brightest and darkest MST levels is crucial for producing an accurate MST chart, which can be achieved with a consumer-grade inkjet printer and glossy paper.

©2024 Society for Imaging Science and Technology

15:00 **Connecting Individual Color Matching Functions to Accuracy and Preference in Image Reproduction,** *Eddie Pei, Susan Farnand, and Michael J. Murdoch, Rochester Institute of Technology (US)* **122**

Human perception of color varies between individuals, raising the question of how well the standard color matching functions (CMFs) represent individual observers in image reproduction. The goal of this research is to explore the relationship between CMFs and both fidelity and preference in the image reproduction pipeline. Three experiments were conducted: an experiment to estimate approximately individualized CMFs, an image fidelity experiment, and an image preference experiment. The results show that the CMFs influence the accuracy of image reproduction, however, preferences are affected by factors in addition to CMFs. The findings offer insights into the limitations and potential implications of relying on using standard CMFs in image reproduction technologies.

15:20 **Practical RGB Measurement of Fluorescence and Blood Distributions in Skin,** *Emilie Nogué[1,3], Arvin Lin[1,3], Xiaohui Li[1,3], Giuseppe Claudio Guarnera[2,3], and Abhijeet Ghosh[1,3]; [1]Imperial College London, [2]University of York (UK), and [3]Lumirithmic Ltd. (UK)* **133**

Biophysical skin appearance modeling has previously focused on spectral absorption and scattering due to chromophores in various skin layers. In this work, we extend recent practical skin appearance measurement methods employing RGB illumination to provide a novel estimate of skin fluorescence, as well as direct measurements of two parameters related to blood distribution in skin – blood volume fraction, and blood oxygenation. The proposed method involves the acquisition of RGB facial skin reflectance responses under RGB illumination produced by regular desktop LCD screens. Unlike previous works that have employed hyperspectral imaging for this purpose, we demonstrate successful isolation of elastin-related fluorescence, as well as blood distributions in capillaries and veins using our practical RGB imaging procedure.

INTERACTIVE POSTER SESSION / EXHIBIT HAPPY HOUR

15:40 – 17:00

Meet with Interactive (poster) Paper authors to discuss their work over a beverage of your choice and help select the CIC32 Cactus Award winner. Exhibits also open.

CONFERENCE RECEPTION/DINNER AND DRINK THE ΔE CONTEST

18:30 – 21:30

Location: Printemps Room

Join other attendees for a night of fun. We'll start with cocktails and the "Drink the ΔE Contest", where you can test your color-matching skills by trying to create the closest ΔE using food-safe dyes. Afterward, we'll enjoy a buffet dinner together.

FRIDAY 1 NOVEMBER

KEYNOTE AND AWARDS

09:00 – 10:10

Session Chair: Minjung Kim, Meta Reality Labs (US)

Perceptual Color Space: Intrinsic Geometry, Mental Representation, Adaptation Invariance & Neural Decoding, *Qasim Zaidi, State University of New York College of Optometry* [Abstract only]

Similarity based color spaces are indispensable for applications that require equally different colors but could also reveal the mental representations that people use to accomplish many tasks, such as object identification across illuminations, and that make associations efficient for generalization and learning. Similarity judgements are widely used to construct Euclidean perceptual/cognitive spaces, but the intrinsic validity of the assumed geometry is almost never tested. For perceptual color space, we measured perceived midpoints between pairs of colors to estimate equal similarities and used Varignon's Theorem to test the intrinsic geometry. We demonstrate that geometrical structure depends on the mental representation used in judging similarity: No structure was discernible unless observers were instructed to use an opponent-color representation, which resulted in an affine geometry. We show that this affine space is invariant to changes in adaptation. An affine geometry implies that similarity can only be judged within straight lines and across parallel lines, and its neural coding could involve ratios of responses. In primate vision, we show that colors can be decoded from the responses of narrowly tuned Infero-Temporal Cortex cells by using a winner-take-all criterion. We suggest ways to reconcile the manifold of ITC neuronal responses with the affine perceptual geometry.

COLOR PERCEPTION AND COGNITION II

10:10 – 12:10

Session Chair: Maliha Ashraf, University of Cambridge (UK)

10:10 **An Experiment to Evaluate Observer Metamerism on Displays,** *Jinyi Lin, Ming Ronnier Luo, and Keyu Shi, Zhejiang University (China), and Tingwei Huang, THOUSLITE Lighting (Changzhou) Ltd. (China)* **142**

This article describes two color matching experiments using 6 displays and a visual trichromator based on spectral tunable LED technology respectively. Ten observers performed the experiment. The results were used to reveal observer metamerism between displays, to derive each individual's color matching functions (CMFs), and to test their performance against the available CMFs. Comparing different CMFs to estimate observer metamerism of displays, individual CMF (expressed by mean ΔE_{00} of 2.7) clearly outperformed the others. Amongst the CIE CMFs, CIE 2006 10° performed the best (5.1), followed by the other color matching functions, CIE 1931 2° gave the worst performance (11.1). The results indicates a large error could occur by adopting the current standard (CIE 1931 2°) to specify colors and this can be much improved using the individual CMFs. A simulation analysis of display primaries was also conducted to show peak wavelength to have a greater impact on observer metamerism than peak bandwidth.

COFFEE BREAK

10:30 – 11:10

11:10 **Two-dimensional Models to Predict a New Colour Appearance Dataset,** *Molin Li and Ming Ronnier Luo, Zhejiang University (China)* 148

An experiment was carried out to study the two-dimensional (2D) colour appearance scales including whiteness, blackness, vividness, and depth. Forty samples were used to be assessed by 10 observers under a 6000K illuminant at 10, 100, 1,000 to 10,000 cd/m^2. These data were used to develop new 2D scales and compared with the existing ones. It comes out the CAM6- UCS and sCAM based scales performed the best. The sCAM is a colour appearance model based on the Simple Uniform Colour Space (sUCS). Its structure will also be introduced.

11:30 **The Paris Effect: A Tale of Eccentricity,** *Peter Morovic, HP Inc. (Spain), and Ján Morovic, HP Inc. (UK)* 154

Differences in color perception as a function of eccentricity have been studied in the past, especially due to variations of macular pigment in the retina and of cone pigment density. However, such differences have been observed under laboratory conditions and this property of the human visual system is not pronounced under everyday conditions. This paper presents an example of a strong effect that was observed at the 31st Color Imaging Conference in Paris when viewing its stage layout and lighting under the overall environmental conditions of the conference hall. In particular, multiple chromatic purplish-violet focus lights were directed at the stage wall. When focusing on one of the lights' reflections on the stage wall, the other lights to the left or right appeared to have a significantly different, duller and cooler color appearance. The effect was immediate and robust to focusing on any one of the chromatic lights, with the others appearing different at the periphery of the field of view.

In order to understand the effect, an attempt was made first to reproduce it using common LED torches coupled with colored gel filters and a pilot psychophysical experiment was performed. The filtered LED light sources were then also measured for a variety of filter colors and simulations were performed, first with differences between 1931 and 1964 color matching functions and then using a cone sensitivity model that allows for predicting the effect of varying eccentricity. This led to the finding that the effect can be measured and expressed at least qualitatively using existing tools, with simulated colorimetries showing the correct direction of shift from warmer more vibrant to cooler, duller colors. We call this the "Paris effect" in honor of the location where we observed it and Paris being the "city of light".

11:50 **Compensating for Individual Differences in Color Perception in Color Reproduction,** *Camilla Simoncelli and Michael Webster, University of Nevada Reno (US)* 161

Individual differences in normal trichromatic color vision arise at many levels, from variations in the wavelength sensitivity of the eye to factors that influence how we see, categorize, and communicate about colors. There is growing recognition of the value of accounting for these differences in order to provide more consistent color percepts across observers. Most of these efforts have concentrated on correcting for differences in spectral sensitivity and associated effects on color matching. However spectral sensitivity differences have limited influence on color appearance, which can be large, and which may potentially be a more important source of variation in how people respond to and interpret color information. We describe a simple method for measuring observers' hue percepts along with a technique for rendering images to compensate for inter-observer differences in appearance. The approach is easy to implement and does not require specialized equipment and offers potential advantages for many color applications including data visualization and communication.

LUNCH ON OWN
12:10 – 13:30

MATERIAL APPEARANCE

13:30 – 15:30

Session Chair: Lili Zhang, Meta Reality Labs (US)

13:30 **CIC32 Best Paper: Evaluating Lightness Constancy in a High-definition VR Environment,** *Khushbu Y. Patel[1], Laurie M. Wilcox[1], Laurence T. Maloney[2], Krista A. Ehinger[3], Suyash Singh[1], and Richard F. Murray[1]; [1]York University and Center for Vision Research (Canada), [2]New York University (US), and [3]University of Melbourne (Australia)* 166

As VR technology has advanced, its use in performance-critical fields such as medical training and vision research has grown, driving a need for increasingly realistic VR environments. In previous work, we evaluated lightness constancy in a task where viewers matched the reflectance of surfaces at different 3D orientations, and we found substantially poorer lightness constancy in VR than in a physical apparatus. Poor constancy in VR may have been due to simplified rendering of scenes in that study, e.g., largely achromatic Lambertian surfaces. Motivated by these findings, here we evaluated lightness constancy in more realistic VR scenes, rendered with a broad array of materials, colors, textures, and specular highlights, as well as more realistic shadows. We tested two conditions: a Full-Context condition, where these lighting and material cues were available, and a Reduced-Context condition, where they were not. Participants had significantly better lightness constancy in the Full-Context condition than in the Reduced-Context condition, indicating that they exploited these additional cues. However, lightness constancy was still quite poor in absolute terms, despite the availability of rich lighting and material cues. The reasons for this failure of constancy are unclear from previous literature, and this finding suggests a promising research problem with both fundamental interest and practical applications.

13:50 **JPI-first: Adjusting Transparency Toward Optimizing Face Appearance,** *Sofie R. Herbeck, Michael J. Murdoch, and Christopher A. Thorstenson, Rochester Institute of Technology (US)* . . . . **see JPI 7/** DOI: 10.2352/J.Percept.Imaging.2024.7.000404

Augmented reality (AR) combines elements of the real world with additional virtual content, creating a blended viewing environment. Optical see-through AR (OST-AR) accomplishes this by using a transparent beam splitter to overlay virtual elements over a user's view of the real world. However, the inherent see-through nature of OST-AR carries challenges for color appearance, especially around the appearance of darker and less chromatic objects. When displaying human faces—a promising application of AR technology—these challenges disproportionately affect darker skin tones, making them appear more transparent than lighter skin tones. Still, some transparency in the rendered object may not be entirely negative; people's evaluations of transparency when interacting with other humans in AR-mediated modalities is not yet fully understood. In this work, two psychophysical experiments were conducted to assess how people evaluate OST-AR transparency across several characteristics including different skin tones, object types, lighting conditions, and display types. The results provided a scale of perceived transparency allowing comparisons to transparency for conventional emissive displays. The results also demonstrate how AR transparency impacts perceptions of object preference and fit within the environment. These results reveal several areas with need for further attention, particularly regarding darker

©2024 Society for Imaging Science and Technology

skin tones, lighter ambient lighting, and displaying human faces more generally. This work may be useful in guiding the development of OST-AR technology, and emphasizes the importance of AR design goals, perception of human faces, and optimizing visual appearance in extended reality systems.

14:10 **The Influence of Material Roughness on Perceived Gloss and Color Appearance of Graphical Generated Faces,** *Yuan Tian, Mekides Assefa Abebe, and Christopher A. Thorstenson, Rochester Institute of Technology (US)* **173**

The human face is an essential stimulus in our social life and occupies a large proportion of digital content. The perceptual appearance of faces is important in computer vision, psychology, digital media, and related fields. Various structural and color features of faces have been studied over the years. However, the study of perceptual glossiness of faces and its influence on facial skin color appearance is very limited. This study investigates the relationship between perceived glossiness of facial skin, skin roughness, and perceptual color attributes. Psychophysical studies were conducted to model perceptual gloss and its effects on the perceived color appearance of faces. The investigation was carried out across varying roughness levels and skin tones. The results indicated that facial gloss influenced the perceived lightness of the facial skin, a phenomenon not observed to the same extent in non-face objects included in the experiment. The effect on lightness could partially be explained by a strategy of discounting specular components for surface color perception. Observers tended to focus on the brightest regions of the objects while avoiding specular highlights to infer color attributes. The current findings provide insights into understanding visual appearance characteristics of face and non-face objects and will be useful for accurate gloss and color reproduction of graphical generated faces.

COFFEE BREAK
14:30 – 14:50

14:50 **Surface Roughness Estimation for Reproducing Appearance of Glossy Object Surfaces,** *Shoji Tominaga, Norwegian University of Science and Technology (Norway) and Nagano University (Japan), and Motonori Doi, Osaka Electro-Communication University (Japan)* **179**

We consider the estimation of surface roughness to reproduce the appearance of objects with smooth and glossy surfaces such as lacquerware and plastic objects. Two methods comprising measurement-based and image-based roughness are used for estimation. First, a laser scanning system is used to measure the microscopic surface height of the target object and calculate the surface normal vectors at every grid point from the height information. The surface roughness is then calculated as the 2D deviation of the surface normal vectors. Next, the Beckmann roughness parameter is then estimated using a high-dynamic-range (HDR) image captured from a flat surface of the target material. The specular lobe is approximated using the Beckmann distribution function with a surface roughness parameter. Furthermore, images are rendered to reproduce the surface appearance and confirm the reliability of the estimated roughness parameters. We study the relationship between the measurement- and image-based roughness and find a linear relationship between them. The Beckmann roughness parameter required for image rendering is predicted from the measured roughness.

15:10 **Searching for Colors of 3D Translucent Objects,** *Davit Gigilashvili, Dipayan Chowdhury, and Jon Yngve Hardeberg, Norwegian University of Science and Technology (Norway)* **185**

3-dimensional translucent objects exhibit large spatial variations of color, and it remains unclear what colors humans choose cognitively to represent monomaterial 3D translucent objects. Previous works asked human observers to match translucent object with a most representative uniform color patch. Matched colors were compared with the average color of the object, which turned out to be a poor predictor of the match. In this work, we conducted a more thorough analysis of the color matching experiments and investigated whether matched color is systematically overestimated or underestimated relative to the average color in any particular axis in the CIELAB space. Afterward, we conducted K-means clustering of the pixel colors to extract a palette of dominant colors of the translucent object and compare the palette with observers' match. And finally, we went through a psychophysical experiment to quantify perceived color differences between pairs of 3D translucent objects. We explore whether perceptual color differences can be predicted with traditional color difference formulae using pixel-wise averages or dominant colors from the K-means clustering.

CLOSING REMARKS AND BEST STUDENT PAPER AWARD
15:30 – 15:40

CIC32 General Chair Javier Vázquez Corral, Universitat Autònoma de Barcelona (Spain), and Program Chairs Carol Payne, Apple Inc. and Minjung Kim, Meta Reality Labs (US)

The day ends with some closing remarks from the chairs and the presentation of the Best Student Paper Award.