

Noise Prism: A Novel Multispectral Visualization Technique

Trevor D. Canham,¹ Javier Vazquez-Corral,^{2,3} David L. Long,⁴ Richard F. Murray,¹ and Michael S. Brown¹

¹York University, Canada; ²Universitat Autònoma de Barcelona, Spain; ³Computer Vision Center, Spain

⁴Rochester Institute of Technology, USA

Abstract

A novel technique for visualizing multispectral images is proposed. Inspired by how prisms work, our method spreads spectral information over a chromatic noise pattern. This is accomplished by populating the pattern with pixels representing each measurement band at a count proportional to its measured intensity. The method is advantageous because it allows for lightweight encoding and visualization of spectral information while maintaining the color appearance of the stimulus. A four alternative forced choice (4AFC) experiment was conducted to validate the method's information-carrying capacity in displaying metameric stimuli of varying colors and spectral basis functions. The scores ranged from 100% to 20% (less than chance given the 4AFC task), with many conditions falling somewhere in between at statistically significant intervals. Using this data, color and texture difference metrics can be evaluated and optimized to predict the legibility of the visualization technique.

Introduction

Billions of years before Newton's insight into the nature of color via the prism, our single-celled ancestors experienced sensitivity to light. In the eons that followed the visual system developed according to survival needs, resulting in spatial and color sensitivity. By virtue of its construction, color vision begins with a limited probe of the spectral power distribution (SPD) of stimuli, where energy is integrated over three spectral bands and encoded as cone excitation values. It follows that identical sets of three excitation values can be triggered by different SPDs, a phenomenon known as metamerism. While this phenomenon allows for imaging systems to reproduce the appearance of color stimuli without exactly matching spectral distributions, it leaves the naked eye blind to potentially relevant physical information.

However, objects like prisms and diffraction gratings allow light decomposition into its spectral components. Taking inspiration from these objects, we propose a method to visually decompose multispectral readings by populating a noise pattern with pixels representing the individual measurement bands such that its probability distribution is proportional to the SPD (Figure 1). The method can be extended to the context of images by randomly selecting an element from this pattern at each pixel location, offering the aesthetic benefit of color-dependent texture (like film grain) and covering up capture and compression artifacts (like dithering) while maintaining color appearance. Users can overlay the patterns at an opacity of their choice for a balance between image legibility and spectral information visibility.

Psychophysicists have shown that we have some ability to perceive probability distribution characteristics of noise patterns, but there is no perfect perceptual metric for predicting the detectability of distribution differences in any domain of visual stimuli. To aid in this effort, a four alternative forced choice (4AFC) psychophysical experiment was conducted in which the legibility of the proposed visualization technique is tested. In this experiment observers were asked to distinguish between




Figure 1. The spectral distributions of metameric stimuli can be differentiated using the Noise Prism technique.

metamers rendered with the noise prism technique.

In summary, our contributions are the following:

- A novel method for visualizing multispectral image data which balances the original input color appearance with the enriched information
- A novel noise generation technique for adding texture to images and dithering artifacts
- A methodology for encoding low resolution multispectral images in a single channel image array
- Psychophysical data on the discriminability of color distributions.

In the sections that follow, related work on the use of color and texture in visualization, multispectral image visualization, and signal distribution discriminability will be discussed. Next, the implementation details of the proposed method will be outlined. Finally, the psychophysical discrimination experiment will be described, followed by its results and a discussion on the usability of the method and future work.

Related Work

Color and Texture in Visualization

If a feature of visual appearance can be synthesized, chances are it has been utilized in a visualization scheme. In particular, color and texture properties have proven useful for extending dimensionality in data visualization [10, 14, 16, 30, 32]. However, due to the non-uniform sensitivity of the visual system in these domains, their mapping remains an open problem and a number of psychophysical experiments have been conducted to characterize the information-carrying capacity of visualization techniques [26, 27]. In most cases, the visualization community simply trusts that color models like CIELAB [6] will offer a degree of perceptual uniformity in data mapping that will be tolerable for their particular application. Urness et al. [29] introduce color weaving, where instead of blending primary colors which are mapped to multivariate components, they are randomly distributed throughout a noise pattern in different frequencies which are proportional to their magnitude. Later, Hagh-Shenas et al. [11] compare weaving to blending, and find the former to have a greater information-carrying capacity.

Multispectral Image Visualization

While the majority of work in multispectral imaging focuses on the computational analysis of captured material, a number of works have addressed the topic of visualization. The most direct method for reproducing multispectral data is via the multi-primary display, for which an extensive review is provided in the work of Long [20]. These devices aim to allow larger gamuts to be displayed while avoiding metamerism failures caused by current wide-color gamut displays with narrow-band primaries. That said, this technology is in its infancy and its potential to advance color rendering is open-ended.

Working within the current display ecosystem, many methodologies for remapping multispectral images for visualization have been proposed [13, 18, 25]. In each method, three spectral bands of interest are identified either manually or automatically from contrast information and mapped to primary colors to render a final image. All of the cited works employ a visualization strategy which would be described by Hagh-Shenas as color blending, and necessarily stray from the color appearance of the captured scenes. The authors are not aware of any previous work applying a color weaving technique to multispectral visualization.

In summary, the multi-primary display has the strength of maintaining fidelity in color rendering, while the channel re-mapping strategy sacrifices color rendering to highlight application-relevant spectral details through re-appropriation of the display primaries and color blending. The proposed method aims to balance this tradeoff by showing spectral details while maintaining color appearance by using a color-weaving visualization technique.

Signal Distribution Discriminability

The foundational works of Attneave [1] and Barlow [2] suggest that neural structures on the visual pathway are optimized to accommodate the statistical properties of stimuli to which we are most commonly exposed (i.e., natural scenes). A thorough review of this concept is provided by Simoncelli & Olshausen [24]. It stands to reason, then, that observers should be sensitive to specific properties of this distribution. In his work which served as the foundation for the field of texture synthesis, Julesz [15] shows that textures can be well approximated simply by matching statistical properties and identifies limits for the discriminability of random Markov patterns along these axes. In particular, his findings show that changes in higher-order statistical moments like skewness and kurtosis are not discriminable in this context, so long as the mean and variance of the distribution are held constant. However, Pratt et al. [22] showed in later experiments that with alternative statistical techniques noise patterns could be created which are discriminable but differ in skewness alone. This finding was recently backed by the discovery of neural structures which are sensitive to changes in skewness [33]. Finally, Canham et al. demonstrate that shifts in a decorrelated “ortho-kurtosis” property are visible in the context of luminance distributions of natural images when all lower order moments are kept the same [4]. In summary, visual sensitivity to the specific shape of stimulus distributions has been underestimated.

The authors are not aware of a general or domain-specific metric which predicts signal distribution discriminability. Such a metric could be useful for optimizing the visibility of relevant comparisons in the proposed technique. While a number of metrics have been proposed to predict color differences, they are not ideal for use with noise patterns as they make comparisons at the pixel level. Image quality metrics like VIF [23], SSIM [31], and

FSIM [34] operate over larger spatial contexts, but are heavily weighted towards deviations in spatial structure, making them inappropriate for this application.

There exist, however, various information theory based divergence measures like L2 and Kullback-Leibler which have proven useful for the analysis of signal distributions [7, 3] and as loss functions for the optimization of image processing methods. However, they do not predict psychophysical data and their scores often conflict with one another [28]. A potential alternative may be the metrics found in the field of texture analysis, where all metrics are necessarily tuned to evaluate global differences in sample distributions [17]. For example, Chubb et al. [5] suggest taking a weighted sum of the differences in statistical moments between test and reference stimuli.

Computing Noise Prisms

We begin with the stimulus SPD and its tristimulus value given an observer spectral sensitivity or color matching function (CMF). First, the SPD or CMF should be sampled such that both tables represent congruent wavelength bands. The SPD is normalized such that it sums to one and is scaled by the desired number of pixels in the noise array. For every wavelength λ , the noise pattern is populated with N pixels of the CMF tristimulus at λ , where N is the value of the normalized SPD at λ , and the array is randomly shuffled.

Once populated, the noise pattern must be mapped to be reproducible by common displays, as the near-monochromatic tristimulus values are all out of gamut for traditional 3-primary displays, according to additive color matching mathematics of an abridged primary system (demonstrated in Figure 2). To accomplish this, the cone excitation responses to the stimulus are converted to *IPT* [8] where hue and chroma are well decorrelated and accessible by converting chroma channels P and T to polar coordinates. In this representation, the chroma channel of the noise pattern is globally scaled until all RGB values are in the 0-1 range. It can be observed that the compressed gamut shown in Figure 2 maintains the general shape of the spectral locus but at a reduced scale. In an alternative strategy, users can set the chroma, intensity, or both channels to equal the input stimulus, leaving only hue to vary as a function of wavelength λ .

Finally, the mean value of the pattern is shifted to the RGB value of the stimulus in the intended display space using a channel-wise exponent to ensure all pixels stay in range (as opposed to an additive offset or a multiplicative scaling which could push values out of range). Figure 5 shows a comparison of the proposed power law matching versus an alternative where the mean value is additively adjusted. It can be observed that the former strategy better preserves the color appearance of the Macbeth chart as it requires less gamut compression. The results of the default strategy described above are applied in Figures 3 and 8, where the method is extended to full-resolution images by selecting a random pattern element for each pixel location.

A simplified version of the method can be used to generate single-channel transmission patterns. In this configuration, the input spectral bands are linearly mapped to code values in the 0-1 range, such that the histogram of the pattern returns the area normalized spectral power distribution directly. The array can be encoded in the alpha channel of PNG or TIFF images and allows for a low-resolution multispectral image to be transmitted and decoded. Figure 4 shows that averaging over the Macbeth reflectance set and a range of illumination sources, the Mean Relative Absolute Error (MRAE [19]) stabilizes at a spatial resolution of 10×10 pixels for a spectral resolution of 78 bands. Since




Figure 2. Displayed in CIEXYZ space with an sRGB gamut volume, with red dots, noise pattern elements before gamut mapping, and their resulting position after gamut mapping and matching the mean RGB value of the stimulus (gray Macbeth patch) with blue dots. It can be seen that the shape of the spectral locus is roughly preserved, but is compressed to fit within the monitor gamut.

observers' contrast sensitivity will be highest for luminance distributions, this rendering also has the potential to be useful for visualization. However, it does not maintain the appearance of the input stimulus. For example, it will have the effect of brightening red image regions and darkening blue regions.

Methodology

Participants

25 students and staff between the ages of 21 and 50 from York University participated in the experiment. Each observer was verified to have normal color vision with an Ishihara color deficiency test. They received a small gift as an incentive for their participation.

Apparatus

The experiment employed a 2014 iMac as a computing platform and display, running software coded in MATLAB using the Psychtoolbox [21]. The software managed the cadence of the experiment, displaying stimuli, querying the input device, and saving participant performance data.

The display output was measured for primary colors and white at 16 drive values with a PhotoResearch PR-655 spectrophotometer to verify additivity and consistency across color regions. The display's gamma function and primary values were derived from the measurements to accurately represent noise pattern colorimetry. The experiment was conducted in a dark environment for ease of repeatability. Observers' input responses were submitted on a standard QWERTY keyboard.

Stimuli

Observers were presented with four noise prism patterns corresponding to the same patch reflectance. Each pattern subtended two degrees of visual angle at the display's native resolution. Three out of the four patterns were generated with the same illumination source, while the fourth was generated with a metameric alternative. The patch reflectances were a subset of X-Rite Macbeth color checker patches representing a range of hue, saturation, and brightness levels. While the number of patches




Figure 3. The MetaCow set [9] is a multispectral image where the front and back half of each cow have differing reflectance SPDs. (a) Shows the set rendered under a D65 light source, where both sides are metamers, while (b) is the noise prism output. It can be observed that distributions subtly but visibly differ between the front and back ends of each cow.

used in the experiment was limited due to time constraints, the method results on the full set of Macbeth reflectances are shown in Figure 5. The light sources included the CIE D65 "daylight" standard and measured SPDs from CRT, OLED, and two LED devices which are shown in Figure 6. Display spectra were chosen since their primaries could be conveniently re-scaled on a patch-by-patch basis to produce the same CIEXYZ values as the D65 illumination case, rendering all stimuli for the same patch reflectance as metameric matches. Given these SPDs, the stimuli are computed using the noise prism methodology described in the previous section in the configuration where the chroma channel is fixed to equal the stimulus.




Figure 4. Spectral reconstruction Mean Relative Absolute Error (MRAE), averaged over Macbeth reflectances illuminated by an assortment of sources, as a function of transmission pattern size ($N \times N$) for various spectral resolution levels.




Figure 5. (a) Noise prism rendering of Macbeth reflectances with the proposed method, (b) noise prism rendering with additive appearance matching, (c) single channel transmission pattern, where spectral samples are mapped to luminance levels, (d) intensity as a function of wavelength reflectance plots.




Figure 6. Illuminant spectral power distributions measured from 380-780 nanometers. From left to right, the illuminants are D65, CRT, LED, OLED, and an alternative LED.

Procedure

At the start of each session the experimental instructions were read aloud, detailing the cadence of the experiment and task (i.e., to select the patch that is different from the other three). Before starting the body of the experiment, the observers practiced using the interface with 5 random trials. For the main experiment, the conditions were presented to observers in a randomized order. Patches were flashed for 1 s, then disappeared to prevent delays due to overthinking. The next trial began immediately after participants submitted their responses. Most observers completed the experiment in 15 minutes or less.

Design

The experiment employed an 8×10 within-subjects design with the following independent variables and levels:

- 8 patch reflectances (dark skin, light skin, blue sky, foliage, purplish blue, moderate red, green, and blue)
- 10 illumination metamer pairs (all unique combinations of Figure 6)

The dependent variable was discrimination accuracy. The total number of trials was 8000 ($= 25 \text{ participants} \times 8 \times 10 \times 4 \text{ repeats}$)

Results

The overall mean for discrimination accuracy was 69%. The effects of both patch reflectance ($F_{7,79} = 25.7, p < .001$) and illumination metamer pair ($F_{9,79} = 382.82, p < .001$) were significant. Across conditions, the average observer results ranged from 20% to 100% accuracy with a standard deviation of roughly 8%. The mean values for different patch reflectance and illumination pair conditions are plotted in Figure 7. Intervals represent the standard error for each measurement. It can be ob-




Figure 7. Average accuracy scores for independent variables (a) patch reflectance and (b) illumination metamer pair. Error bars represent standard error.

served that these scarcely overlap between adjacent ranked conditions. The plots also show that the illumination pair condition had a greater effect on discrimination accuracy when compared to patch reflectance, with some conditions at almost 100% accuracy across all observers and patch reflectances and others approaching chance (25% accuracy). Table 1 breaks down the results per condition.

Discussion

A multispectral visualization technique is proposed and its capability to render metamers such that they can be discriminated is demonstrated. In particular, broadband SPDs like D65 appear different from multi-modal narrow band distributions. Figure 3 shows the method applied to a full-resolution metameric scene.

The application of multispectral architectural renderings is proposed and demonstrated in Figure 8. While the interpretation of these visualizations may take practice, the resulting color noise and transmission patterns provide a continuous map of the light spectrum at different regions of the scene, which would not be possible with one-dimensional graphs of the spectral readings or a standard RGB rendering. The disparities between the noise prism and RGB renderings indicate that the visualization may aid in understanding color appearance in scenes with complex illumination conditions.

As a first step towards validating the visualization technique, a 4AFC experiment was conducted to gauge discriminability between different color distributions. The distributions were generated in the application-relevant scenario of differentiating between identical patches which are illuminated by metameric light sources. The accuracy scores ranged from clearly visible (100%) to clearly invisible (20%) with many test conditions in between. This property and the confidence intervals between conditions make this experimental data useful for

	D65-OLED	D65-LED1	D65-LED2	D65-CRT	CRT-OLED	CRT-LED1	CRT-LED2	OLED-LED2	LED1-LED2	LED1-OLED
dark skin	100	100	100	94	80	89	73	51	41	30
light skin	100	100	99	95	85	83	50	51	55	44
moderate red	99	100	98	98	81	75	32	56	60	35
foliage	100	100	100	100	41	49	77	69	55	25
blue sky	96	99	94	90	80	58	47	39	31	58
green	100	100	100	99	25	32	60	54	61	20
purplish blue	100	98	100	95	58	36	30	46	22	45
blue	100	98	98	99	46	34	31	26	22	52

Table 1. Accuracy scores were averaged across 25 observers for each individual test condition.



Figure 8. (a) Original RGB rendering of the bathroom scene from Hao & Funt [12] along with (b) color and (c) transmission pattern renderings of multispectral EXR. The visualizations allow for a spatially continuous assessment of the spectral content of the scene. The noise prism renderings reveal the spatial distribution of the mixed illumination in a way that is not obviously apparent in the RGB version (e.g., the reflection of the floor molding on the base of the sink can be seen in the middle image, and the shadow boundaries of the sink can be better seen on the far right.)

evaluating or optimizing metrics to predict distribution difference visibility. This is a promising direction for future work. Also, further psychophysical experimentation, varying presentation conditions like the power spectrum of the patterns, temporal presentation conditions, and testing with image data would be of interest.

Acknowledgments

The authors would like to thank the observers who volunteered their time to participate in the experiment, as well as Yuta Asano for providing the display measurements. JVC was supported by Grant PID2021-128178OB-I00 funded by MCIN/AEI/10.13039/501100011033, ERDF “A way of making Europe”, the Departament de Recerca i Universitats from Generalitat de Catalunya with reference 2021SGR01499, and the “Ayudas para la recualificación del sistema universitario español” financed by the European Union-NextGenerationEU. This work was also supported by the Canada First Research Excellence Fund for the Vision: Science to Applications (VISTA) program, a NSERC Discovery Grant, and the Canada Research Chair program (MSB).

References

- [1] Fred Attneave. Some informational aspects of visual perception. *Psychological review*, 61(3):183, 1954.
- [2] Horace B Barlow. Possible principles underlying the transformation of sensory messages. *Sensory communication*, 1(01):217–233, 1961.
- [3] Ramiro RA Barreira and Lee Luan Ling. Kullback–leibler divergence and sample skewness for pathological voice quality assessment. *Biomedical Signal Processing and Control*, 57:101697, 2020.
- [4] Trevor D. Canham, Adrián Martín, Marcelo Bertalmío, and Javier Portilla. Using decoupled features for photo-realistic style transfer. *SIAM SIIMS*, 16(3), 2023.
- [5] Charles Chubb, Michael S. Landy, and John Economouly. A visual mechanism tuned to black. *Vis. Res.*, 44(27):3223–3232, 2004.
- [6] CIE Colorimetry. Publication No. 15. Technical report, Commission Internationale de l’Éclairage (CIE), 2004.
- [7] JJ Joshua Davis, Florian Schübeler, Sungchul Ji, and Robert Kozma. Discrimination between brain cognitive states using shannon entropy and skewness information measure. In *IEEE SMC*, pages 4026–4031, 2020.
- [8] Fritz Ebner and Mark Fairchild. Finding constant hue surfaces in color space. In *Photonics West '98 Electronic Imaging*, pages 4207–4216, 1998.
- [9] Mark D. Fairchild and Garrett M. Johnson. Metacow: A public-domain, high-resolution, fully-digital, noise-free, metameric, extended-dynamic-range, spectral test target for imaging system analysis and simulation. In *International Conference on Communications in Computing*, 2004.
- [10] Patricia Gilmartin and Elisabeth Shelton. Choropleth maps on high resolution crts/the effects of number of classes and hue on communication. *Cartographica: The International Journal for Geographic Information and Geovisualization*, 26(2):40–52, 1989.
- [11] Haleh Hagh-Shenas, Sunghee Kim, Victoria Interrante, and Christopher Healey. Weaving versus blending: a quantitative assessment of the information carrying capacities of two alternative methods for conveying multivariate data with color. *IEEE TVCG*, 13(6):1270–1277, 2007.
- [12] Xiangpeng Hao and Brian Funt. A multi-illuminant synthetic image test set. *Color Research & Application*, 45(6):1055–1066, 2020.
- [13] Noriaki Hashimoto, Yuri Murakami, Pinky A Bautista, Masahiro Yamaguchi, Takashi Obi, Nagaaki Ohyama, Kuniaki Uto, and Yukio Kosugi. Multispectral image enhancement for effective visualization. *Optics express*,

- 19(10):9315–9329, 2011.
- [14] Christopher G Healey. Choosing effective colours for data visualization. In *IEEE VIS*, pages 263–270, 1996.
- [15] Bela Julesz. Visual pattern discrimination. *IRE transactions on Information Theory*, 8(2):84–92, 1962.
- [16] Robert M Kirby, Haralambos Marmanis, and David H Laidlaw. Visualizing multivalued data from 2d incompressible flows using concepts from painting. In *IEEE VIS*, 1999.
- [17] M. Kokare, B.N. Chatterji, and P.K. Biswas. Comparison of similarity metrics for texture image retrieval. In *TENCON*, 2003.
- [18] Steven Le Moan, Alamin Mansouri, Yvon Voisin, and Jon Y Hardeberg. A constrained band selection method based on information measures for spectral image color visualization. *IEEE TGRS*, 49(12):5104–5115, 2011.
- [19] Yi-Tun Lin and Graham D Finlayson. Physically plausible spectral reconstruction from rgb images. In *CVPR Workshops*, 2020.
- [20] David Long and Mark D Fairchild. Optimizing spectral color reproduction in multiprimary digital projection. In *Color and Imaging Conference*, volume 19, pages 290–297, 2011.
- [21] Dennis G. Pelli. The VideoToolbox software for visual psychophysics: transforming numbers into movies. *Spatial Vision*, 10:437–442, 1997.
- [22] William K Pratt, Olivier D Faugeras, and Andre Gagalowicz. Visual discrimination of stochastic texture fields. *IEEE TSMC*, 8(11):796–804, 1978.
- [23] Hamid R. Sheikh and Alan C. Bovik. Image information and visual quality. *IEEE TIP*, 15(2), 2006.
- [24] Eero P Simoncelli and Bruno A Olshausen. Natural image statistics and neural representation. *Annual review of neuroscience*, 24(1):1193–1216, 2001.
- [25] Diego A Socolinsky and Lawrence B Wolff. Multispectral image visualization through first-order fusion. *IEEE TIP*, 11(8):923–931, 2002.
- [26] Maureen Stone, Danielle Albers Szafrir, and Vidya Setlur. An engineering model for color difference as a function of size. In *Color and Imaging Conference*, 2014.
- [27] Danielle Albers Szafrir. Modeling color difference for visualization design. *IEEE TVCG*, 24(1):392–401, 2017.
- [28] Eric E Thomson and William B Kristan. Quantifying stimulus discriminability: a comparison of information theory and ideal observer analysis. *Neural Computation*, 17(4):741–778, 2005.
- [29] Timothy Urness, Victoria Interrante, Ivan Marusic, Ellen Longmire, and Bharathram Ganapathisubramani. Effectively visualizing multi-valued flow data using color and texture. In *IEEE VIS*, 2003.
- [30] Jarke J Van Wijk. Spot noise texture synthesis for data visualization. In *ACM SIGGRAPH*, 1991.
- [31] Zhou Wang, Alan C. Bovik, Hamid R. Sheikh, and Eero P. Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE TIP*, 13(4):600–612, 2004.
- [32] Chris Weigel, William Emigh, Geniva Liu, Russel M Taylor II, James T Enns, and Christopher G. Healey. Oriented sliver textures: A technique for local value estimation of multiple scalar fields. In *GI*, 1999.
- [33] Matthew Yedutenko, Marcus HC Howlett, and Maarten Kamermans. Enhancing the dark side: asymmetric gain of cone photoreceptors underpins their discrimination of visual scenes based on skewness. *The Journal of physiology*, 600(1):123–142, 2022.
- [34] Lin Zhang, Lei Zhang, Xuanqin Mou, and David Zhang. FSIM: A feature similarity index for image quality assessment. *IEEE TIP*, 20(8):2378–2386, 2011.

Author Biography

Trevor Canham is studying color imaging under the supervision of Michael Brown at York University in Toronto. He received a BSc in Motion Picture Science from the Rochester Institute of Technology, and spent several years working in Marcelo Bertalmio's Image Processing for Enhanced Cinematography lab in Barcelona. His interests lie in the interaction between color phenomenology and imaging systems.

Javier Vazquez-Corral is an Associate Professor at Universitat Autònoma de Barcelona and a researcher at the Computer Vision Center in Barcelona. His research interests lie in the topics of computational color imaging and low-level vision. He is also interested in bridging the gap between color in the human brain and its use in computer-vision applications.

David Long is Director of the RIT Center for Media, Arts, Games, Interaction & Creativity (MAGIC) and MAGIC Spell Studios, a university-wide research and development laboratory, and a commercial media production studio. His research interests at RIT include designing virtual production workflows and color management, engineering multispectral video capture and display systems, studying variability in human color vision, and modeling sensory perception of motion media.

Richard Murray is a faculty member in the Department of Psychology and Centre for Vision Research at York University. His laboratory develops computational models of human visual perception, recently with a focus on achromatic surface color.

Michael S. Brown is a professor and Canada Research Chair at York University. His research focuses on in-camera color processing algorithms for photographic and scientific applications.