

Spatial Color Appearance and Image Difference Models

(iCAM and the like)

VIPLab Webinar

Mark Fairchild, Rochester Institute of Technology Munsell Color Science Laboratory

March 11, 2021





MAGIC

Outline

6 Papers

- **MOM** (S.N. Pattanaik, J.A. Ferwerda, M.D. Fairchild, and D.P. Greenberg, A multiscale model of adaptation and spatial vision for image display, *Proceedings of SIGGRAPH 98*, 287-298 (1998).)
- **iCAM** (M.D. Fairchild and G.M. Johnson, The iCAM framework for image appearance, differences, and quality, *Journal of Electronic Imaging* **13**, 126-138 (2004).)
- **HDR Survey** (M.D. Fairchild, The HDR photographic survey, *IS&T/SID 15th Color Imaging Conference*, Albuquerque, 233-238 (2007).)
- **iCAM06** (J. Kuang, G.M. Johnson, and M.D. Fairchild, iCAM06: A refined image appearance model for HDR image rendering, *Journal of Visual Communication and Image Representation* **18**, 406-414 (2007).)
- **Large Color Differences** (S. Abasi, M.A. Tehran, and M.D. Fairchild, Distance metrics for very large color differences, *Color Research and Application* **44**, 10.1002/col.22451 208-223(2019).)
- **FLIP** (P. Andersson, J. Nilsson, M. Oskarsson, K. Åström and M.D. Fairchild, FLIP: A difference evaluator for alternating images, *ACM SIGGRAPH / Eurographics High Performance Graphics 2020*, ONLINE (2020).)

MOM (*Multiscale Observer Model*)

A Multiscale Model of Adaptation and Spatial Vision for Realistic Image Display

Sumanta N. Pattanaik James A. Ferwerda Mark D. Fairchild* Donald P. Greenberg
Program of Computer Graphics†, Cornell University

Abstract

In this paper we develop a computational model of adaptation and spatial vision for realistic tone reproduction. The model is based on a multiscale representation of pattern, luminance, and color processing in the human visual system. We incorporate the model into a tone reproduction operator that maps the vast ranges of radiances found in real and synthetic scenes into the small fixed ranges available on conventional display devices such as CRT's and printers. The model allows the operator to address the two major problems in realistic tone reproduction: wide absolute range and high dynamic range scenes can be displayed; and the displayed images match our perceptions of the scenes at both threshold and suprathreshold levels to the degree possible given a particular display device. Although in this paper we apply our visual model to the tone reproduction problem, the model is general and can be usefully applied to image quality metrics, image compression methods, and perceptually-based image synthesis algorithms.

CR Categories: I.3.0 [Computer Graphics]: General;

Keywords: realistic imaging, visual perception, tone reproduction, adaptation, spatial vision

electronic and print-based media which have only moderate output levels and typical dynamic ranges of less than 100 to 1.

Recently graphics researchers have started to address this issue by developing *tone reproduction operators* that map scene radiances to display outputs with the goal of producing a visual match between the scene and the display. There are two major problems to be solved in realistic tone reproduction:

- to find an operator that maps the vast ranges of radiances found in scenes into the range that can be produced by a given display device.
- to be certain that this operator produces images that match our perceptions of the scenes.

The critical element that links these two problems is the *visual model* used in the tone reproduction operator. Visual models are used to relate the perceptual responses of a scene observer to the responses of the display observer in order to specify a mapping that produces a visual match between the scene and the display. A central issue is that different tone reproduction operators have made use of different visual models to determine what constitutes a match.

Tumblin and Rushmeier's [1993] operator is based on Stevens'

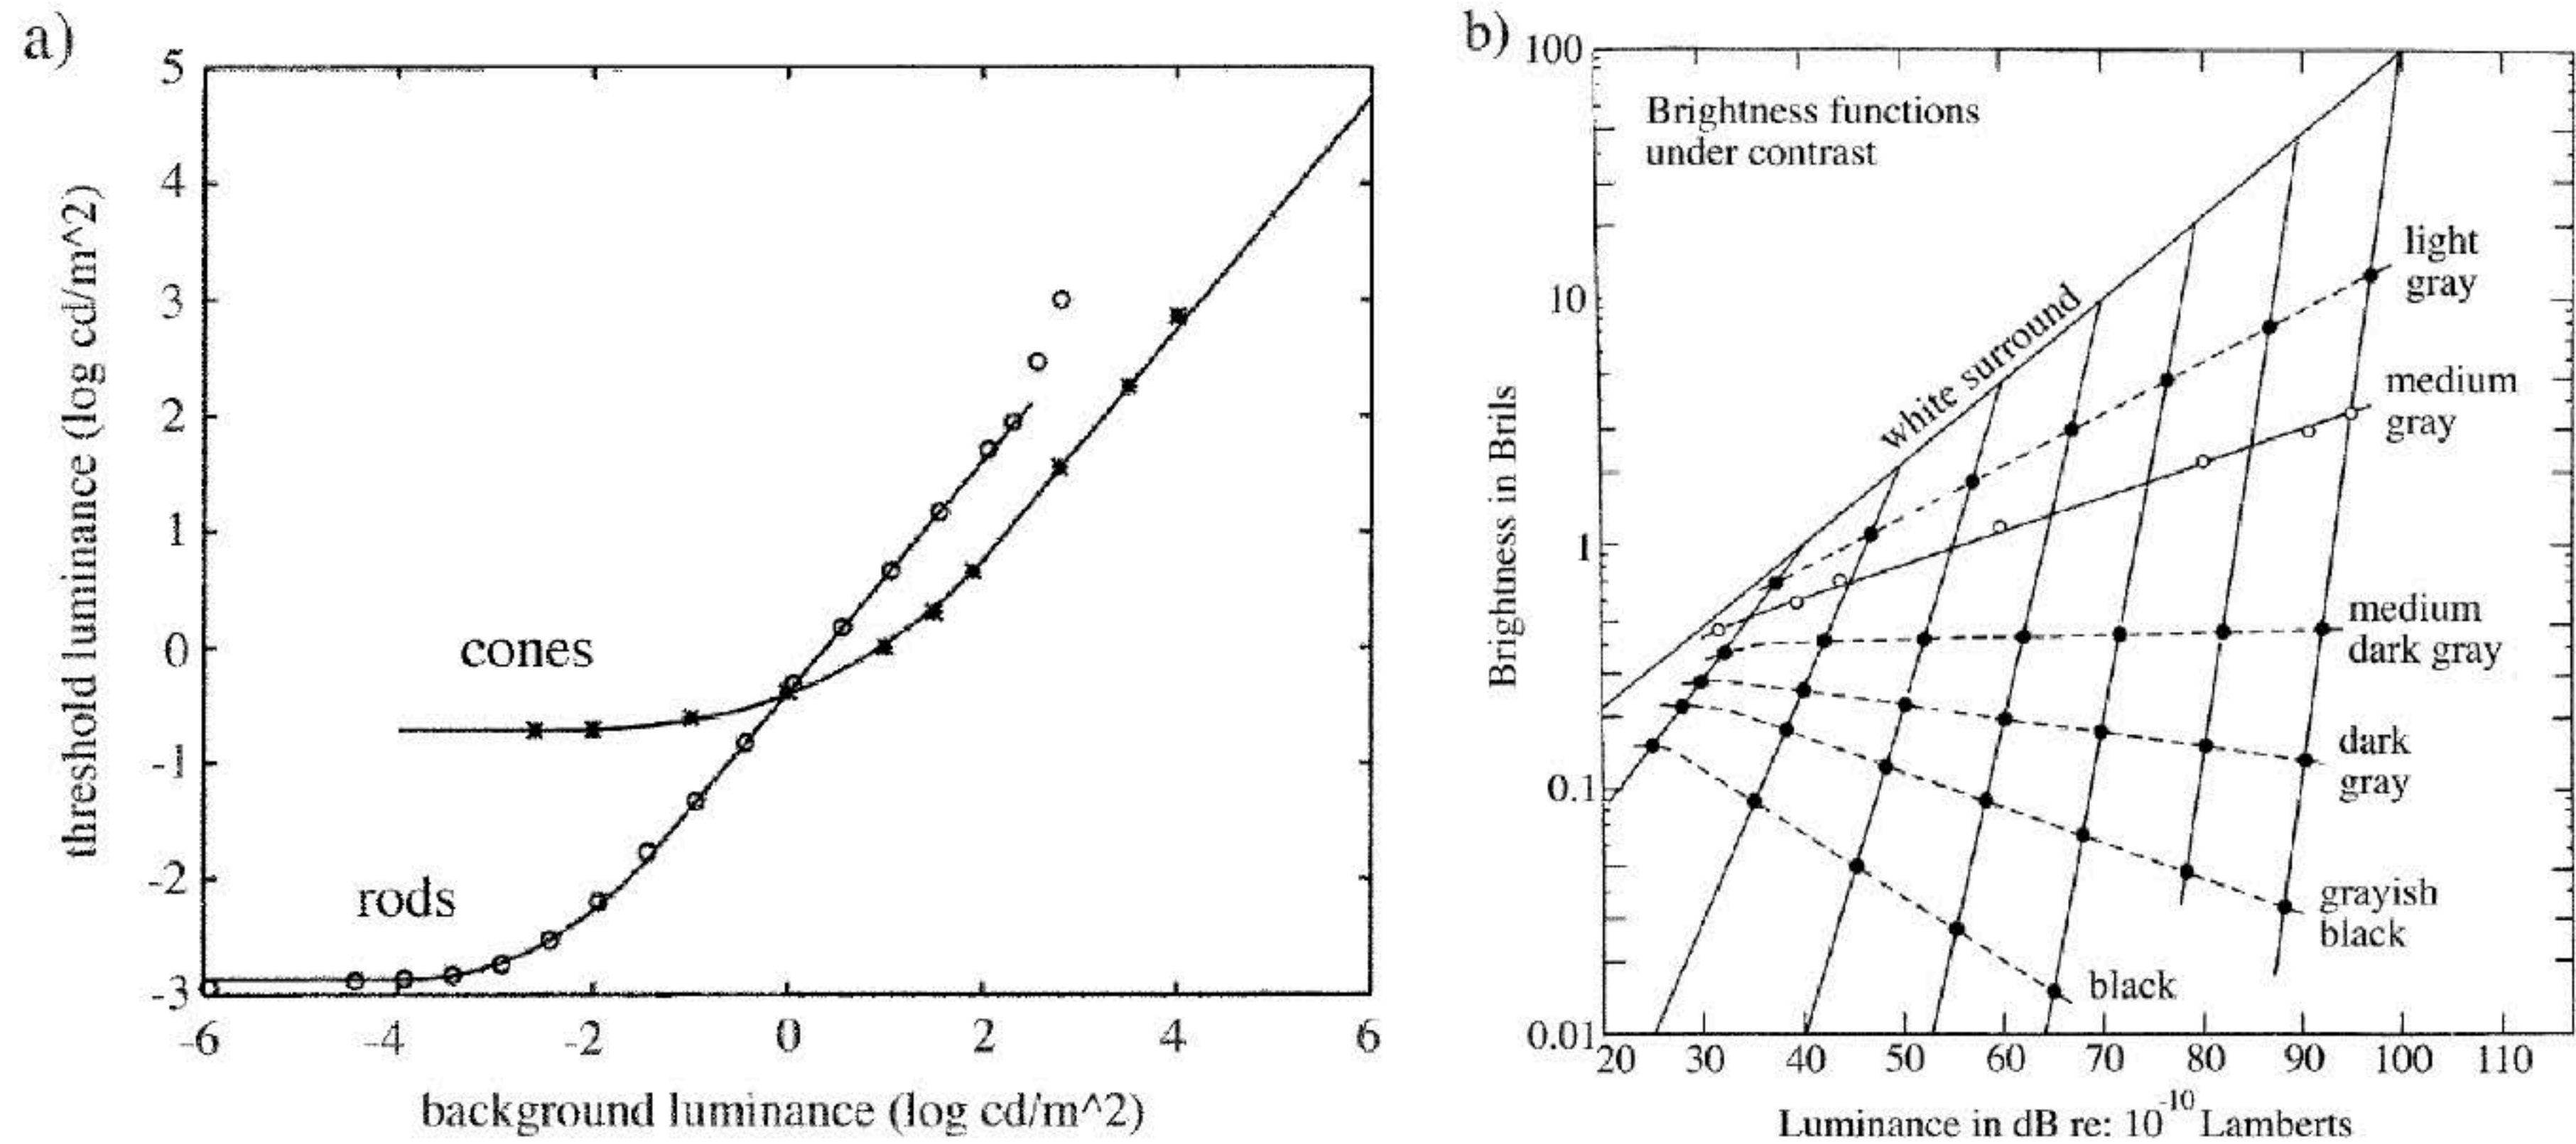


Figure 1: Threshold and suprathreshold models of vision: a) Threshold vs. intensity (TVI) functions for the rod and cone systems. The curves plot the smallest threshold increment ΔL necessary to see a spot against a uniform background with luminance L . b) Stevens' model of suprathreshold brightness and apparent contrast. The curves plot the changes in brightness and apparent contrast of gray targets and a white surround as the level of illumination rises (1 Brill = apparent brightness of a target with a luminance of $1 \mu\text{Lambert}$). Adapted from [Ferwerda96, Stevens61].

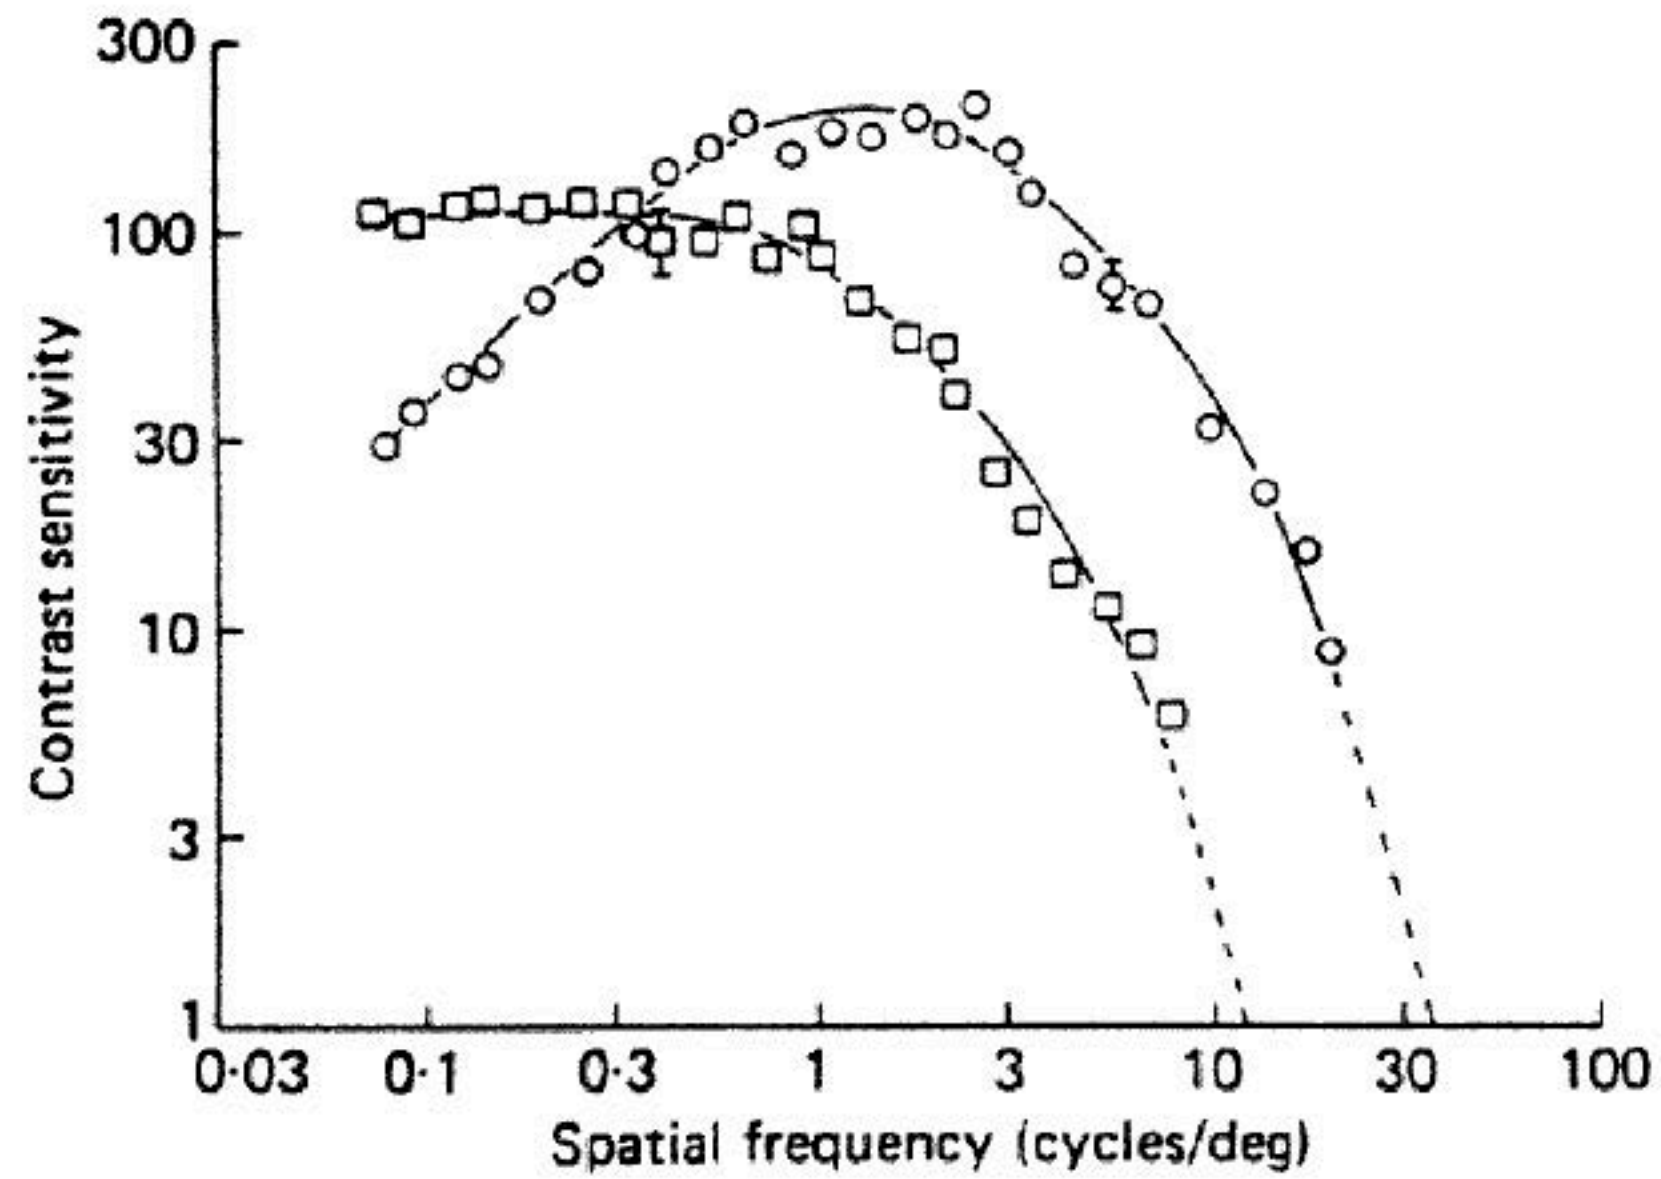


Figure 2: Threshold contrast sensitivity as a function of spatial frequency for a monochromatic luminance grating (\circ ; *green*; 526nm) and a isoluminant chromatic grating (\square ; *red/green*; $602, 526\text{nm}$). Adapted from [Mullen85].




Figure 3: Contrast sensitivity functions for sinusoidal gratings illuminated at different mean luminance levels. Levels are specified in Troland (Td) units of retinal illuminance (Trolands = luminance in $\text{cd}/\text{m}^2 \times$ pupil area). Adapted from [vanNes67].




Figure 4: Multiscale bandpass mechanisms underlying the contrast sensitivity functions. Adapted from [Lubin95].

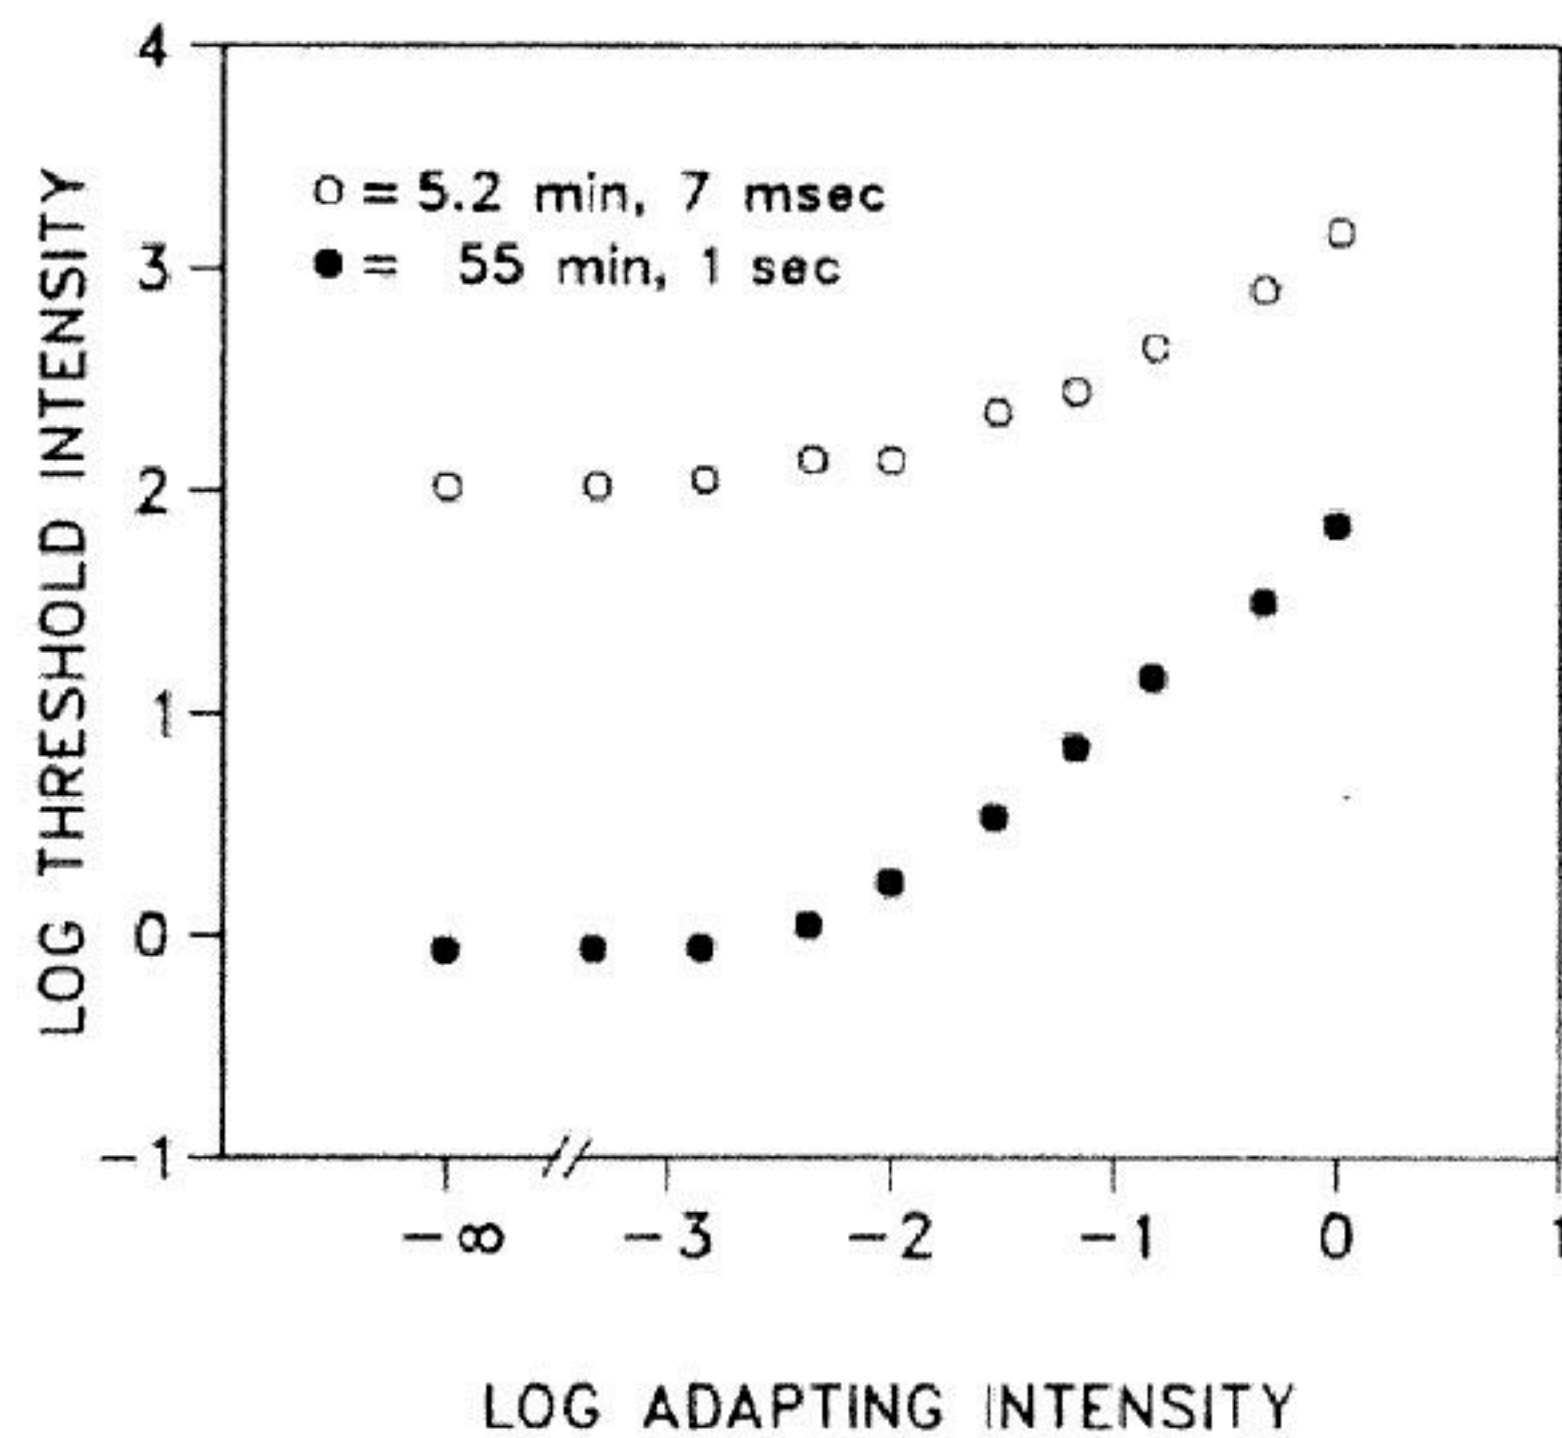


Figure 5: Threshold-vs.-intensity functions for spot patterns with different spatial and temporal parameters. Adapted from [Barlow72].





Figure 6: Suprathreshold contrast constancy and non-linear contrast transducers in human vision. Adapted from [Georgeson75, Watson97b].

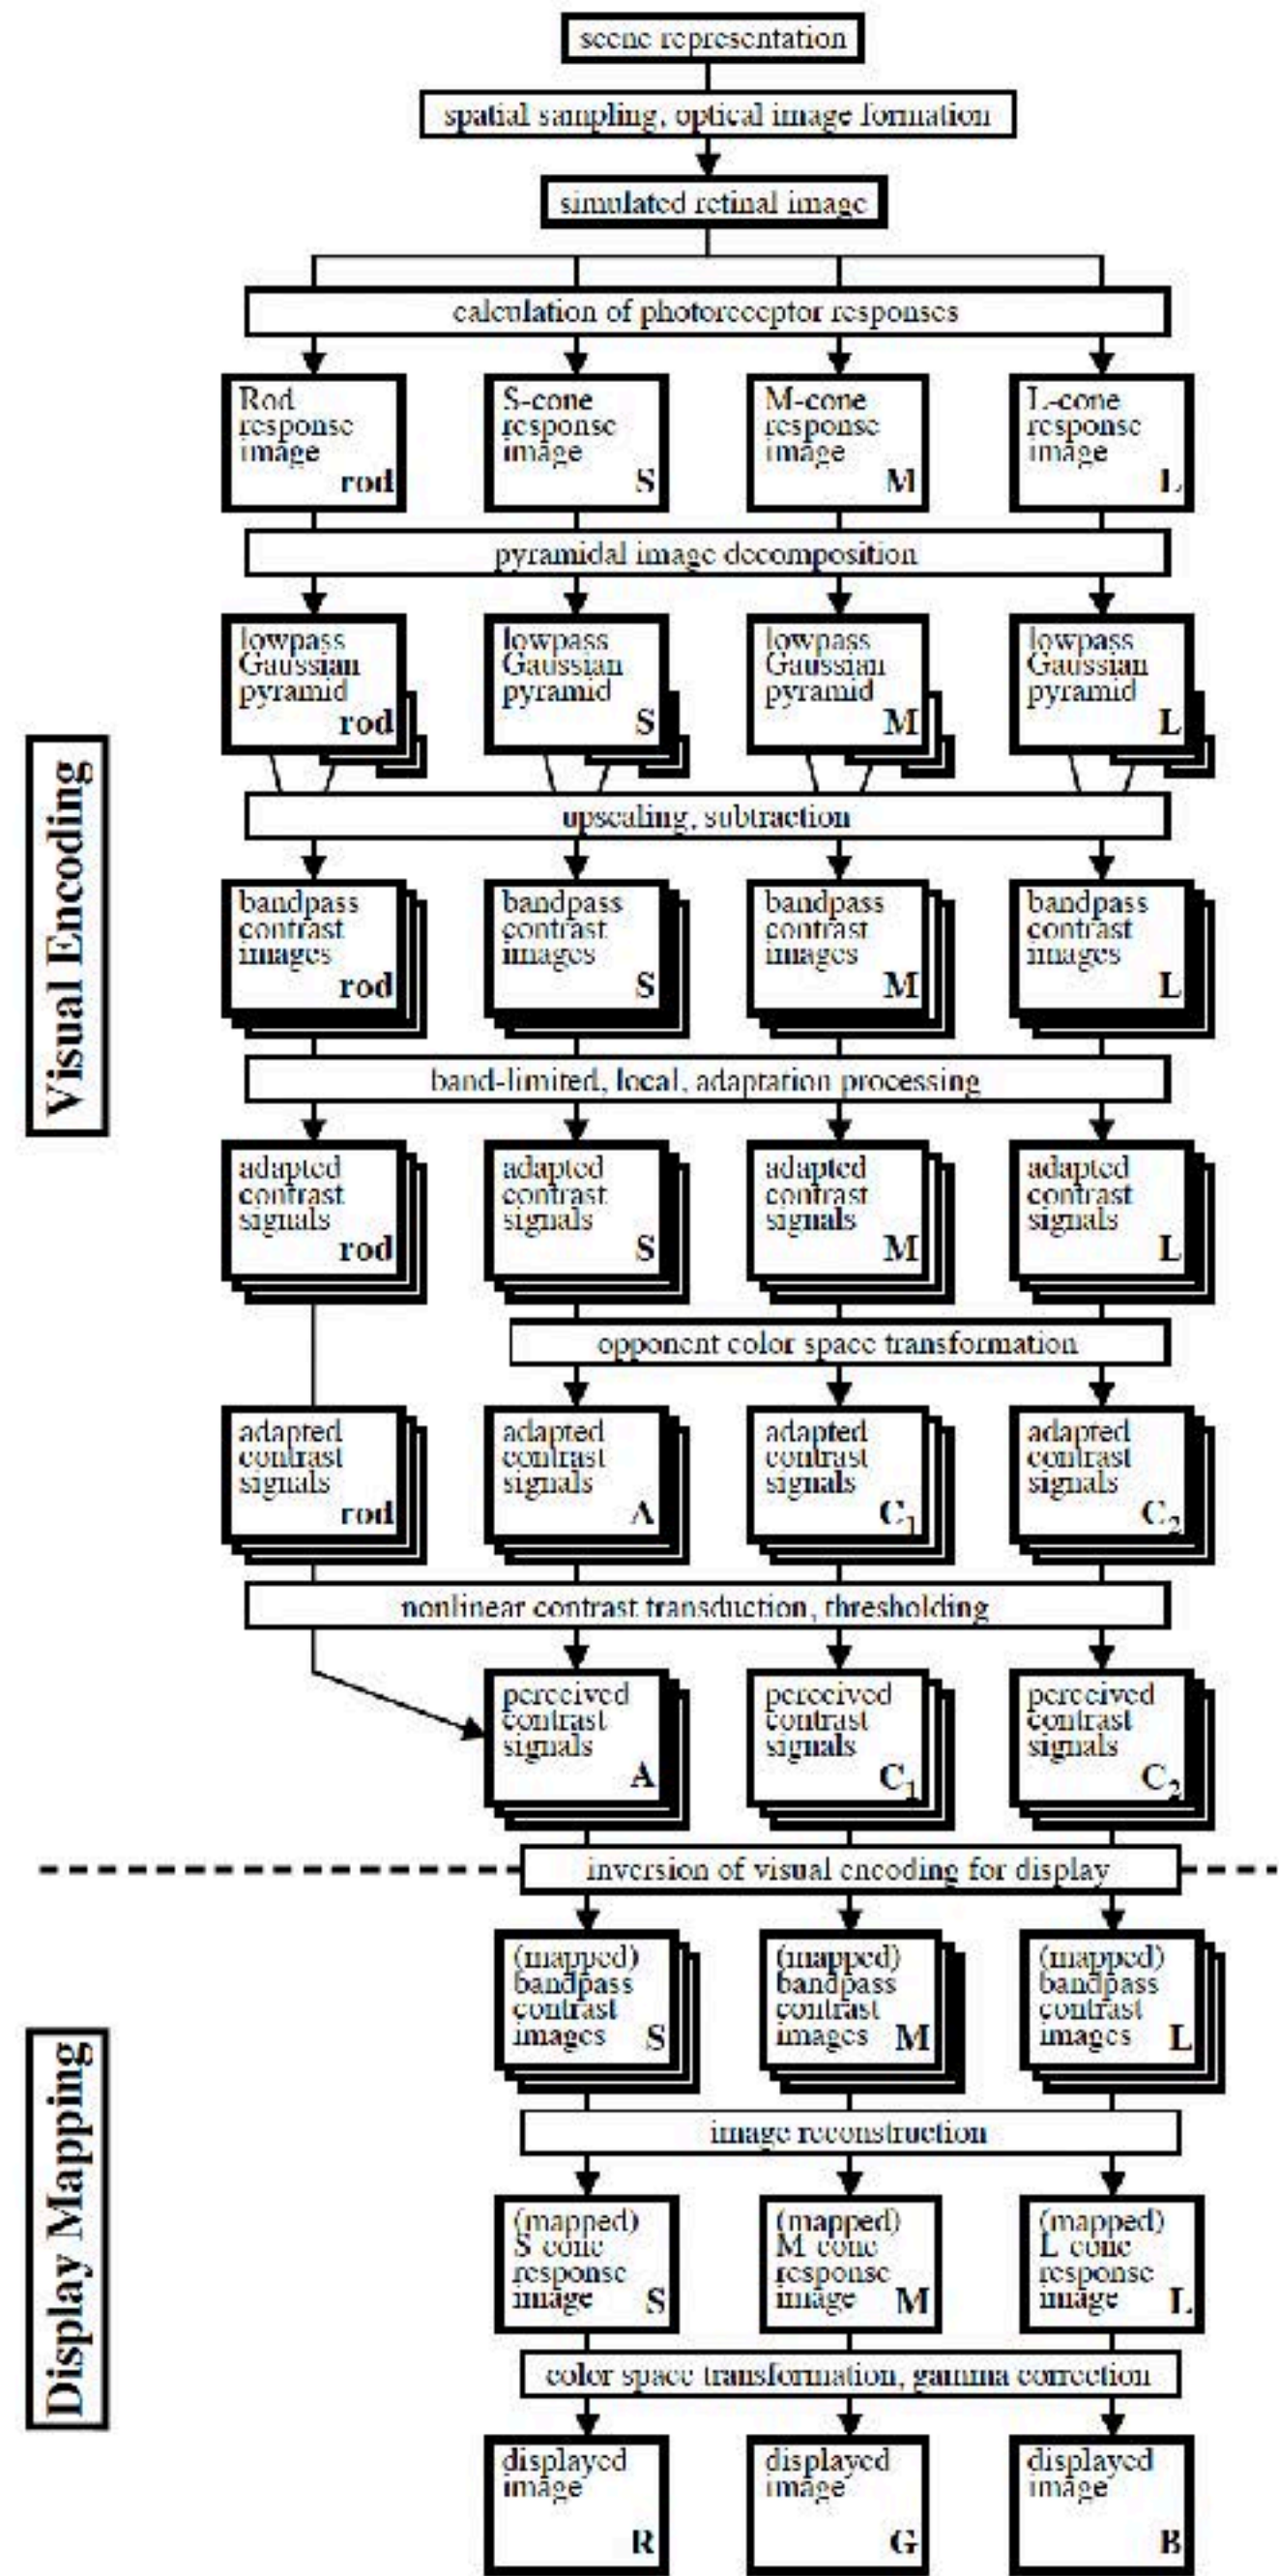


Figure 7: Flow chart of the computational model of adaptation and spatial vision for realistic tone reproduction.

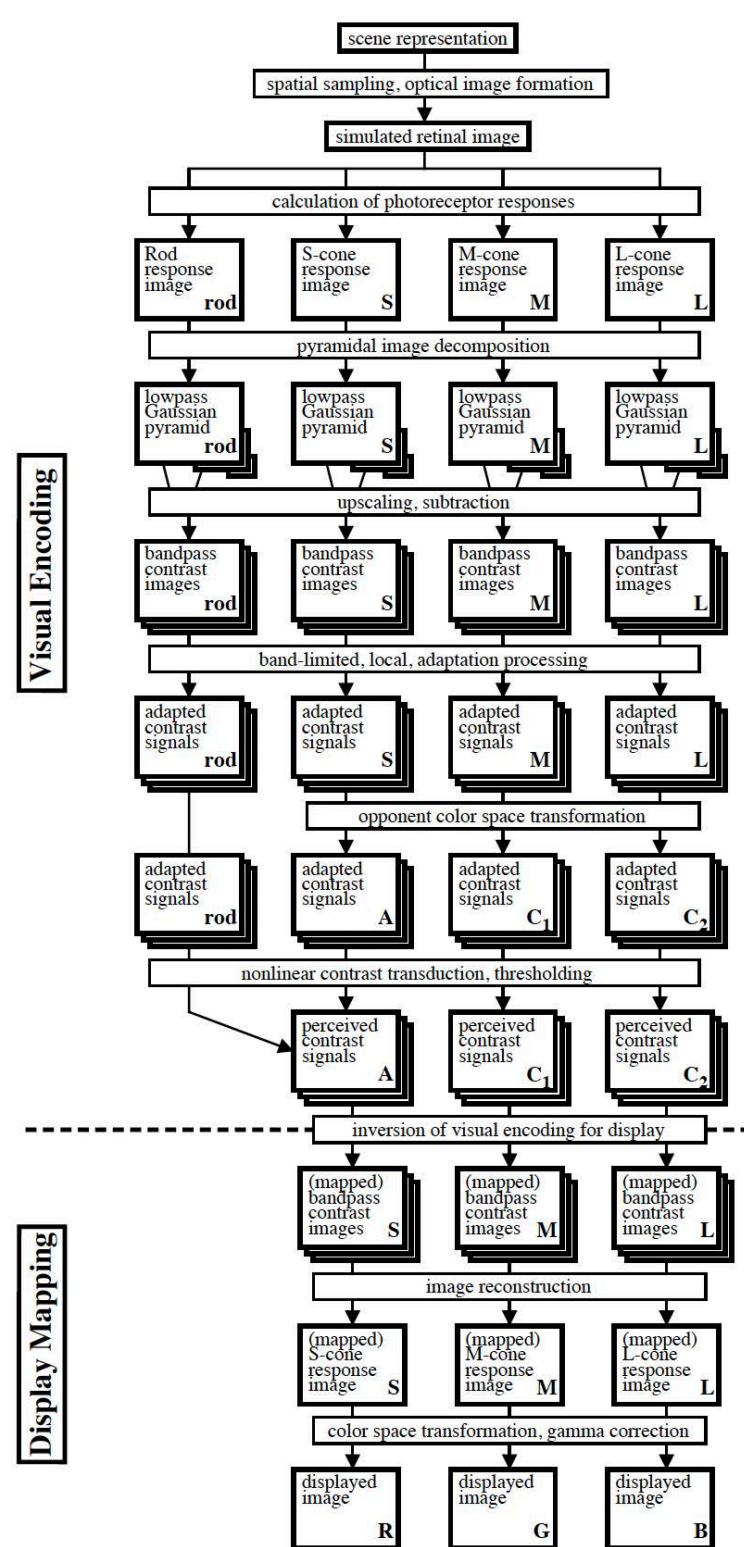
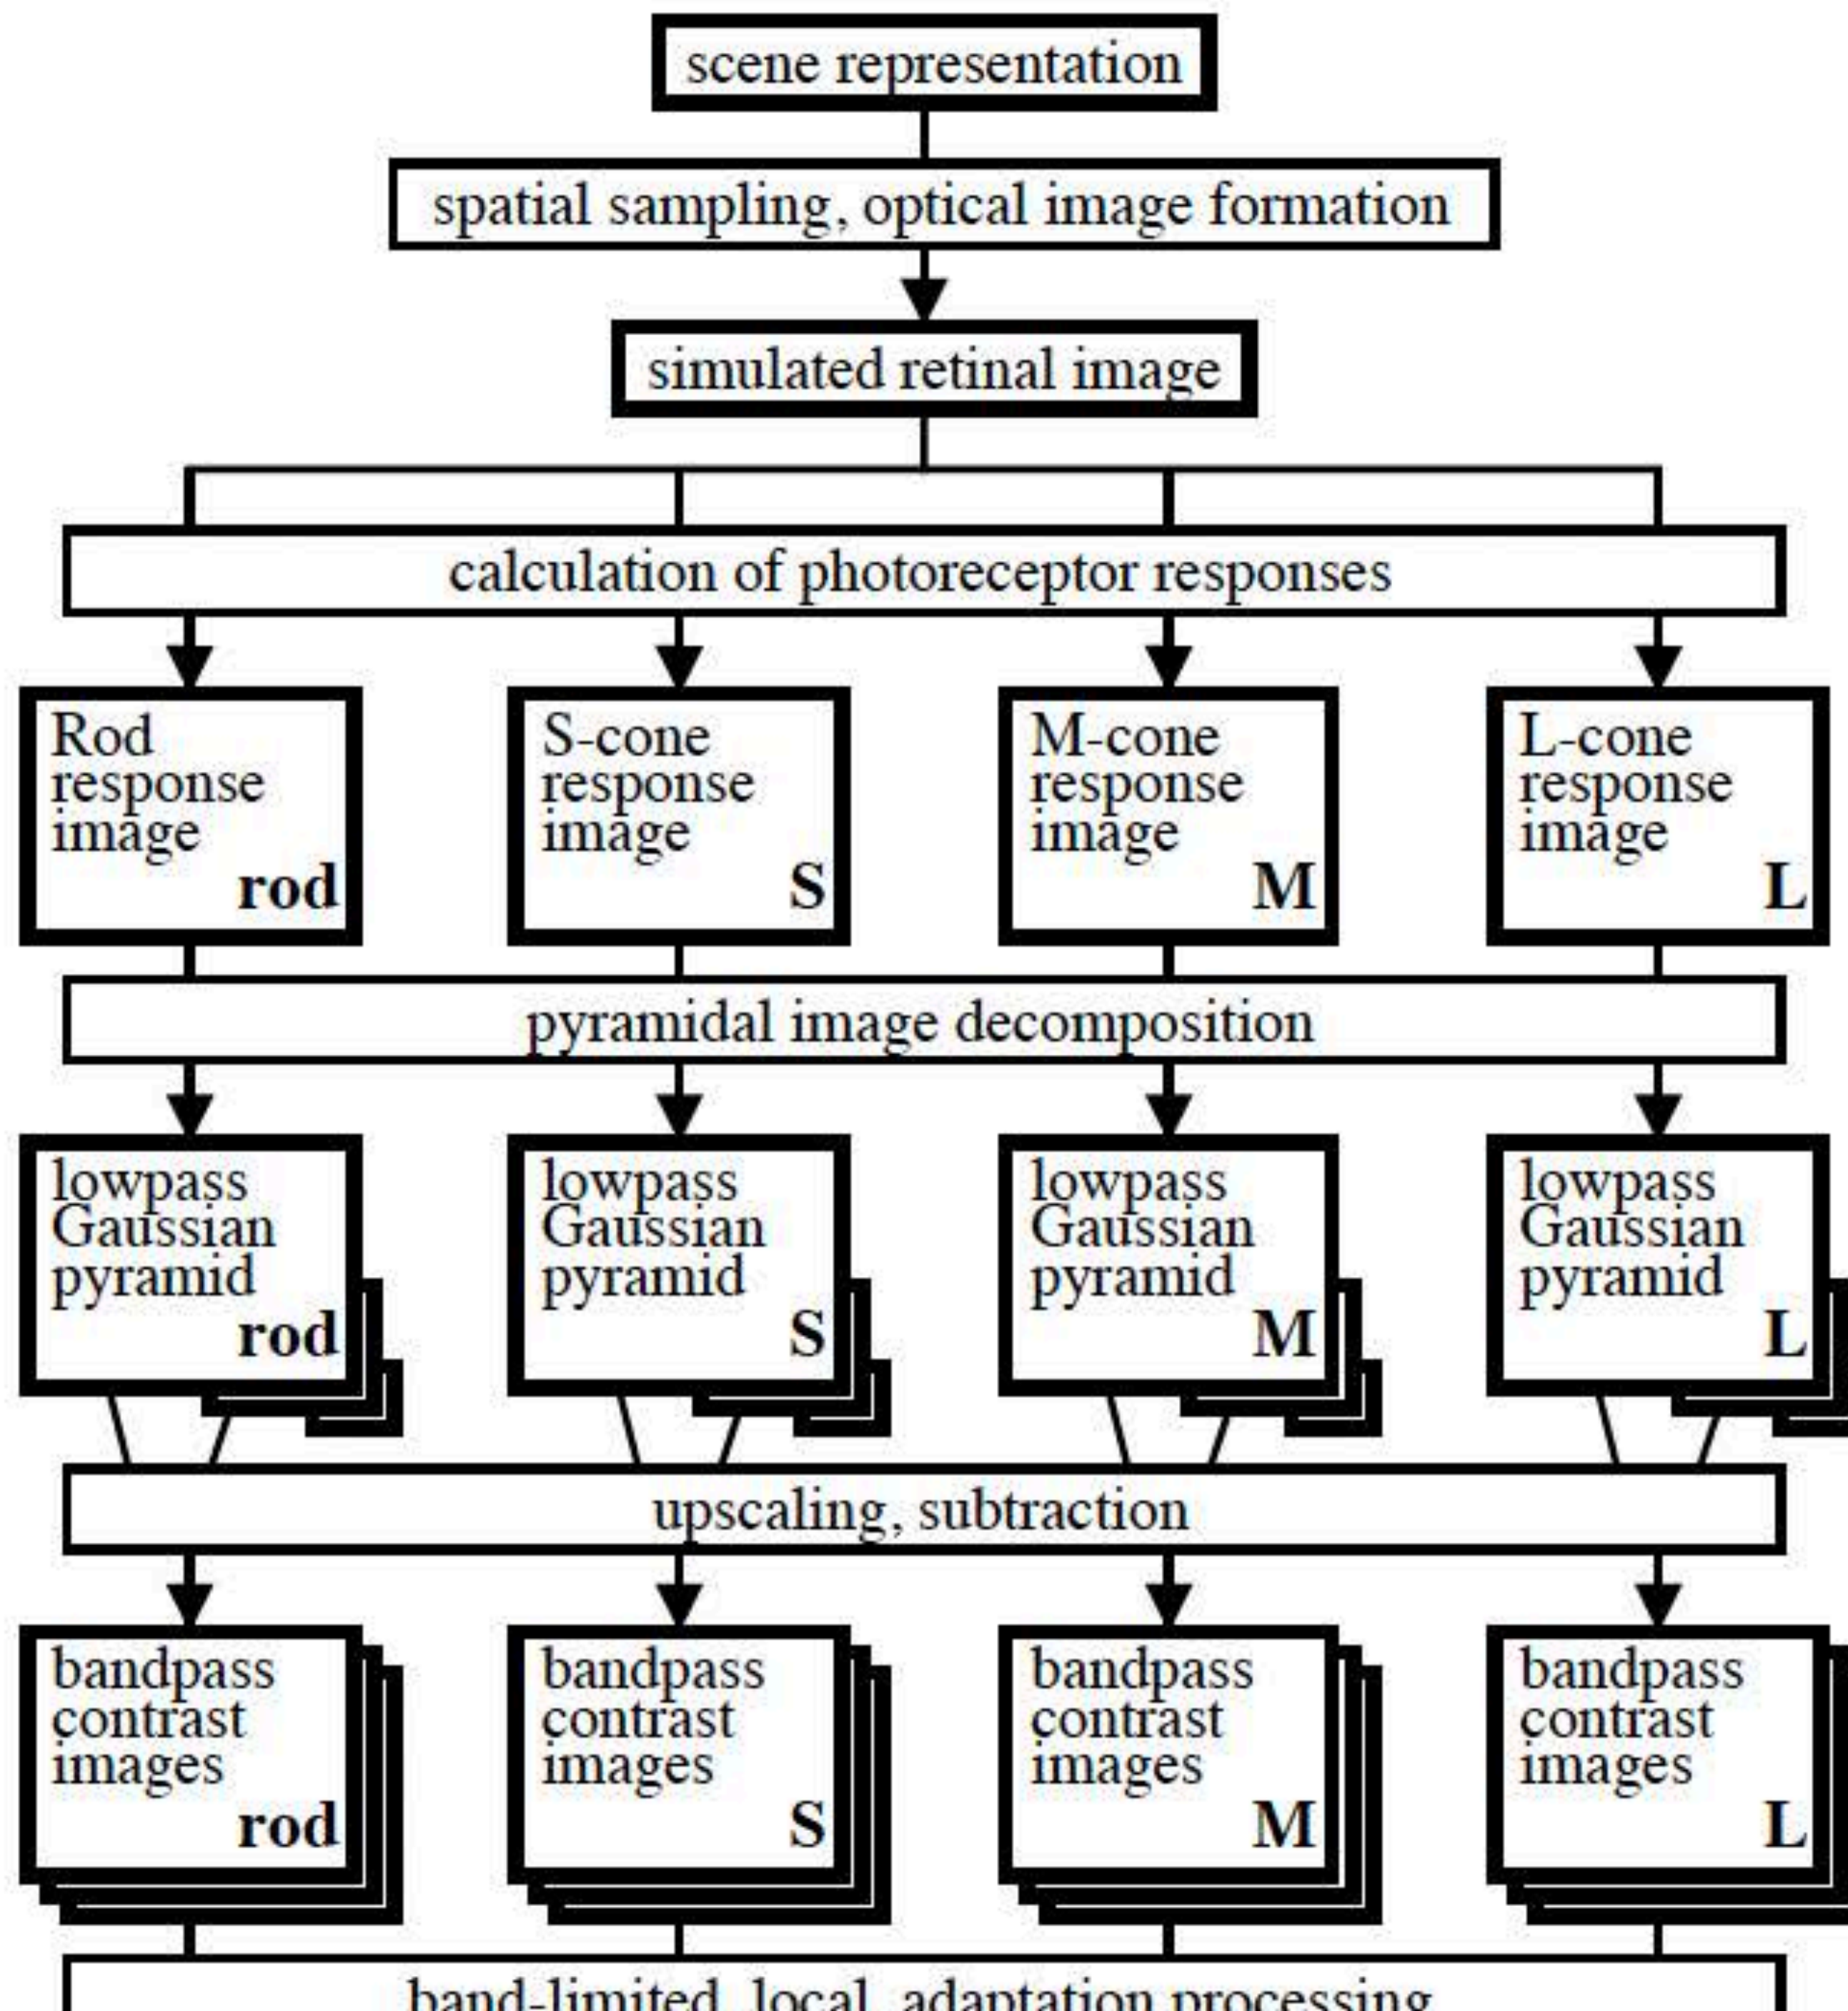


Figure 7: Flow chart of the computational model of adaptation and spatial vision for realistic tone reproduction.

Visual Encoding



Visual Encoding

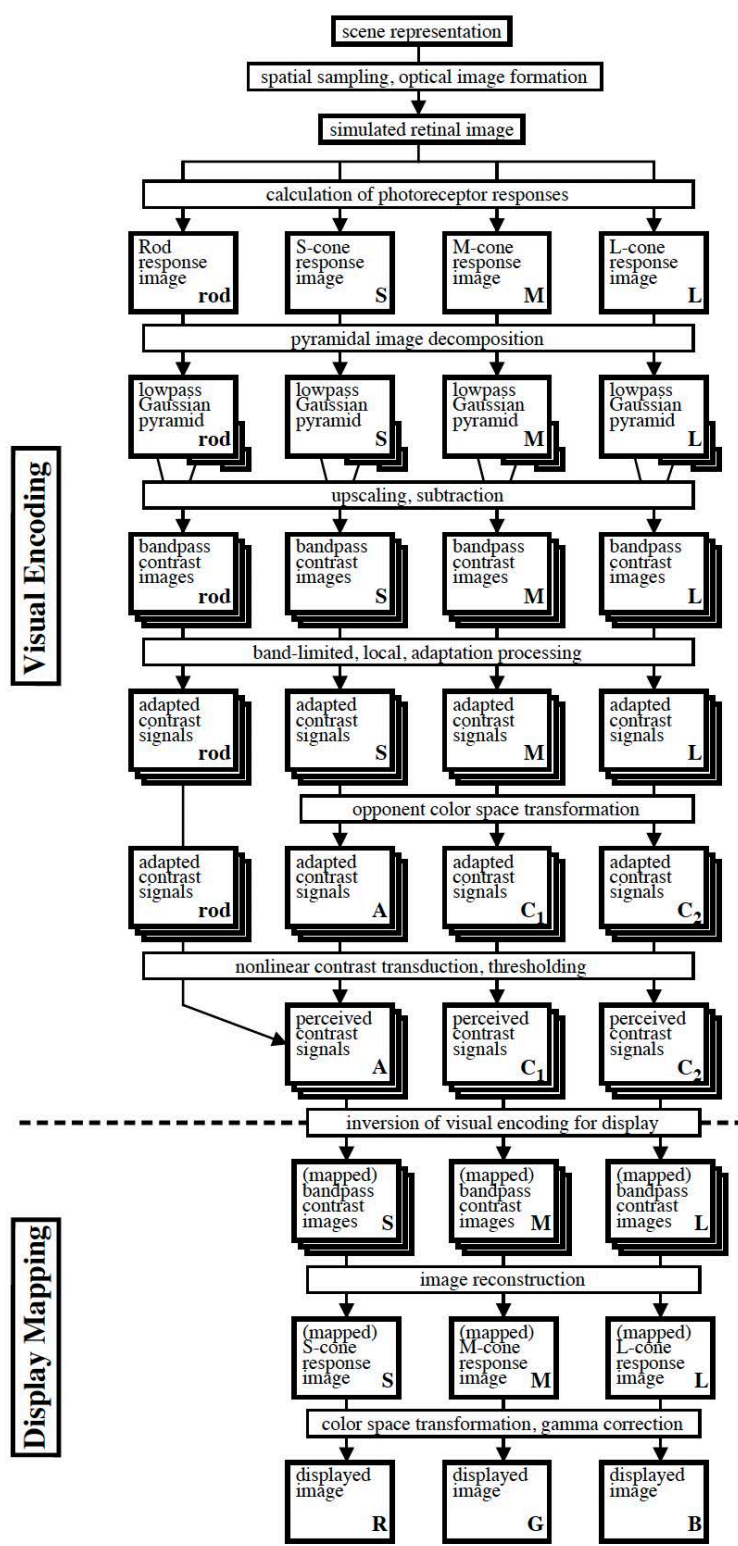
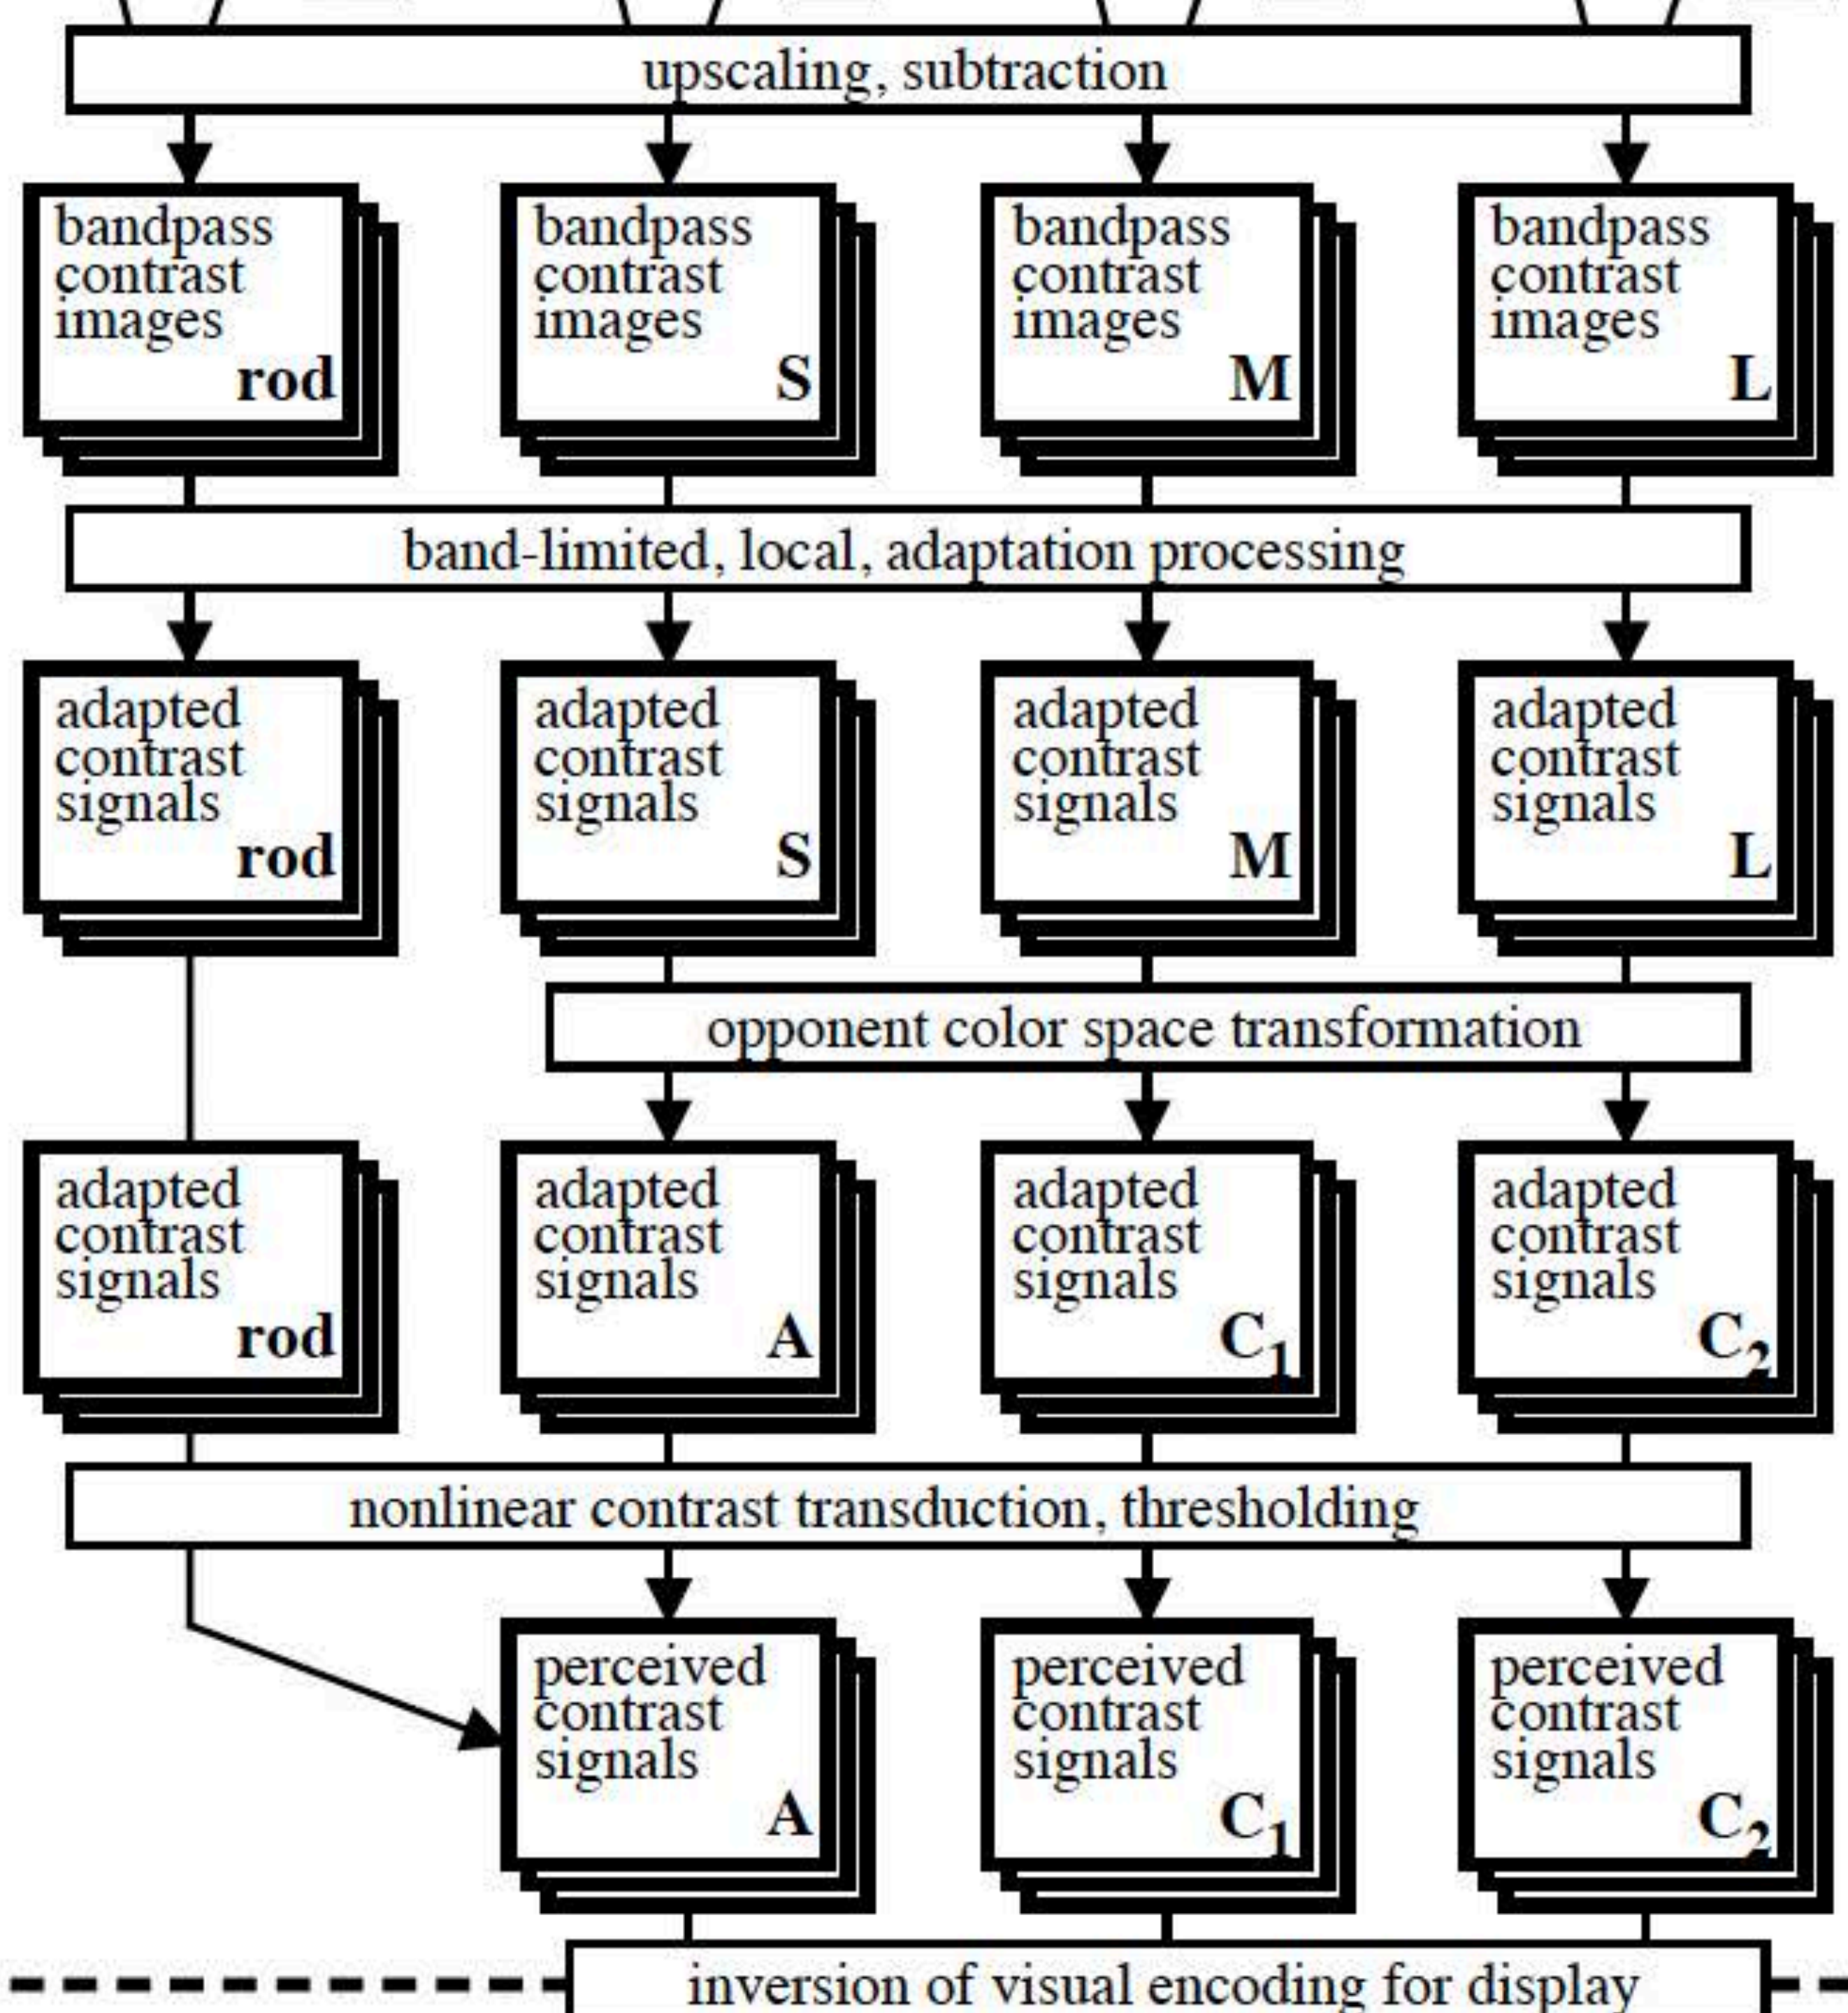


Figure 7: Flow chart of the computational model of adaptation and spatial vision for realistic tone reproduction.



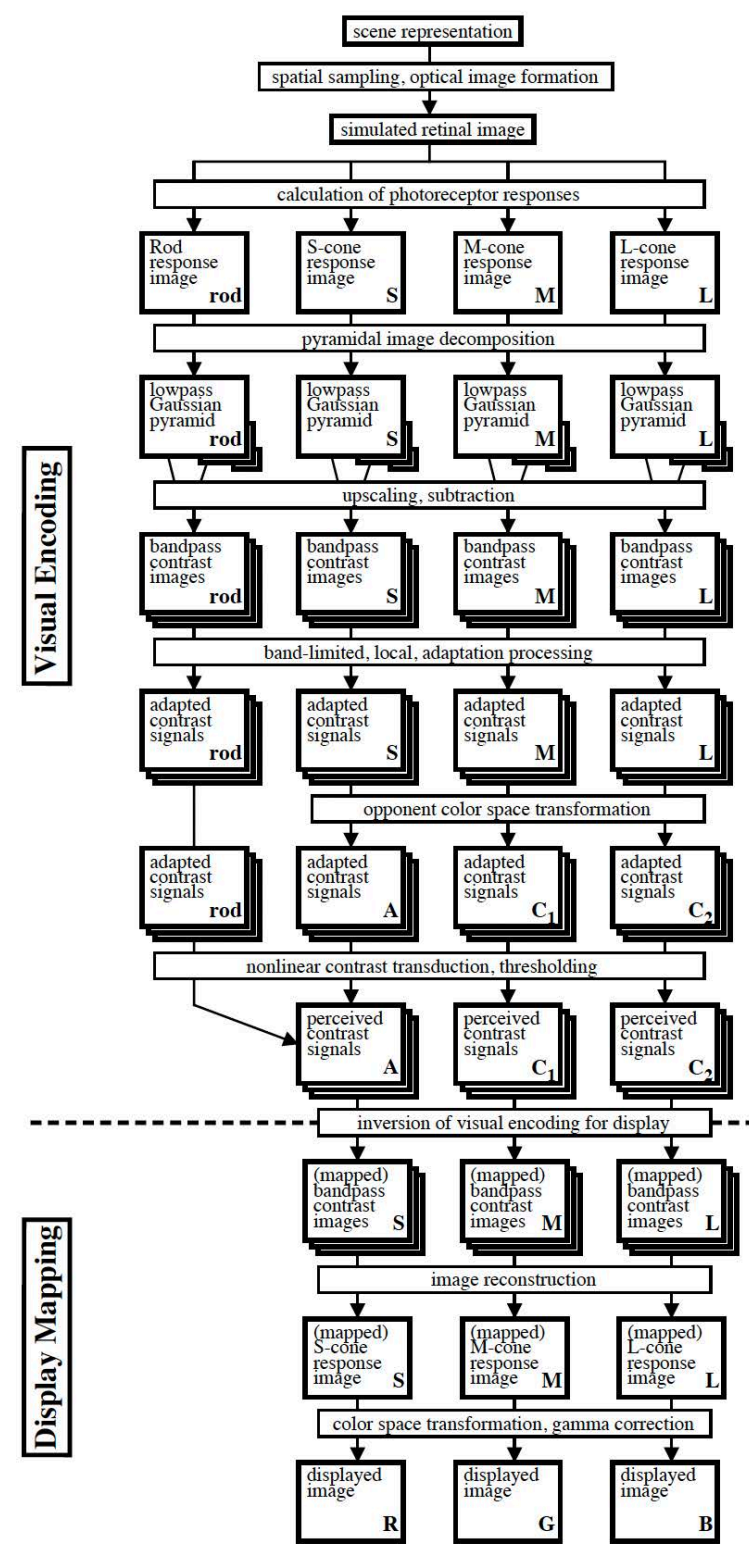


Figure 7: Flow chart of the computational model of adaptation and spatial vision for realistic tone reproduction.

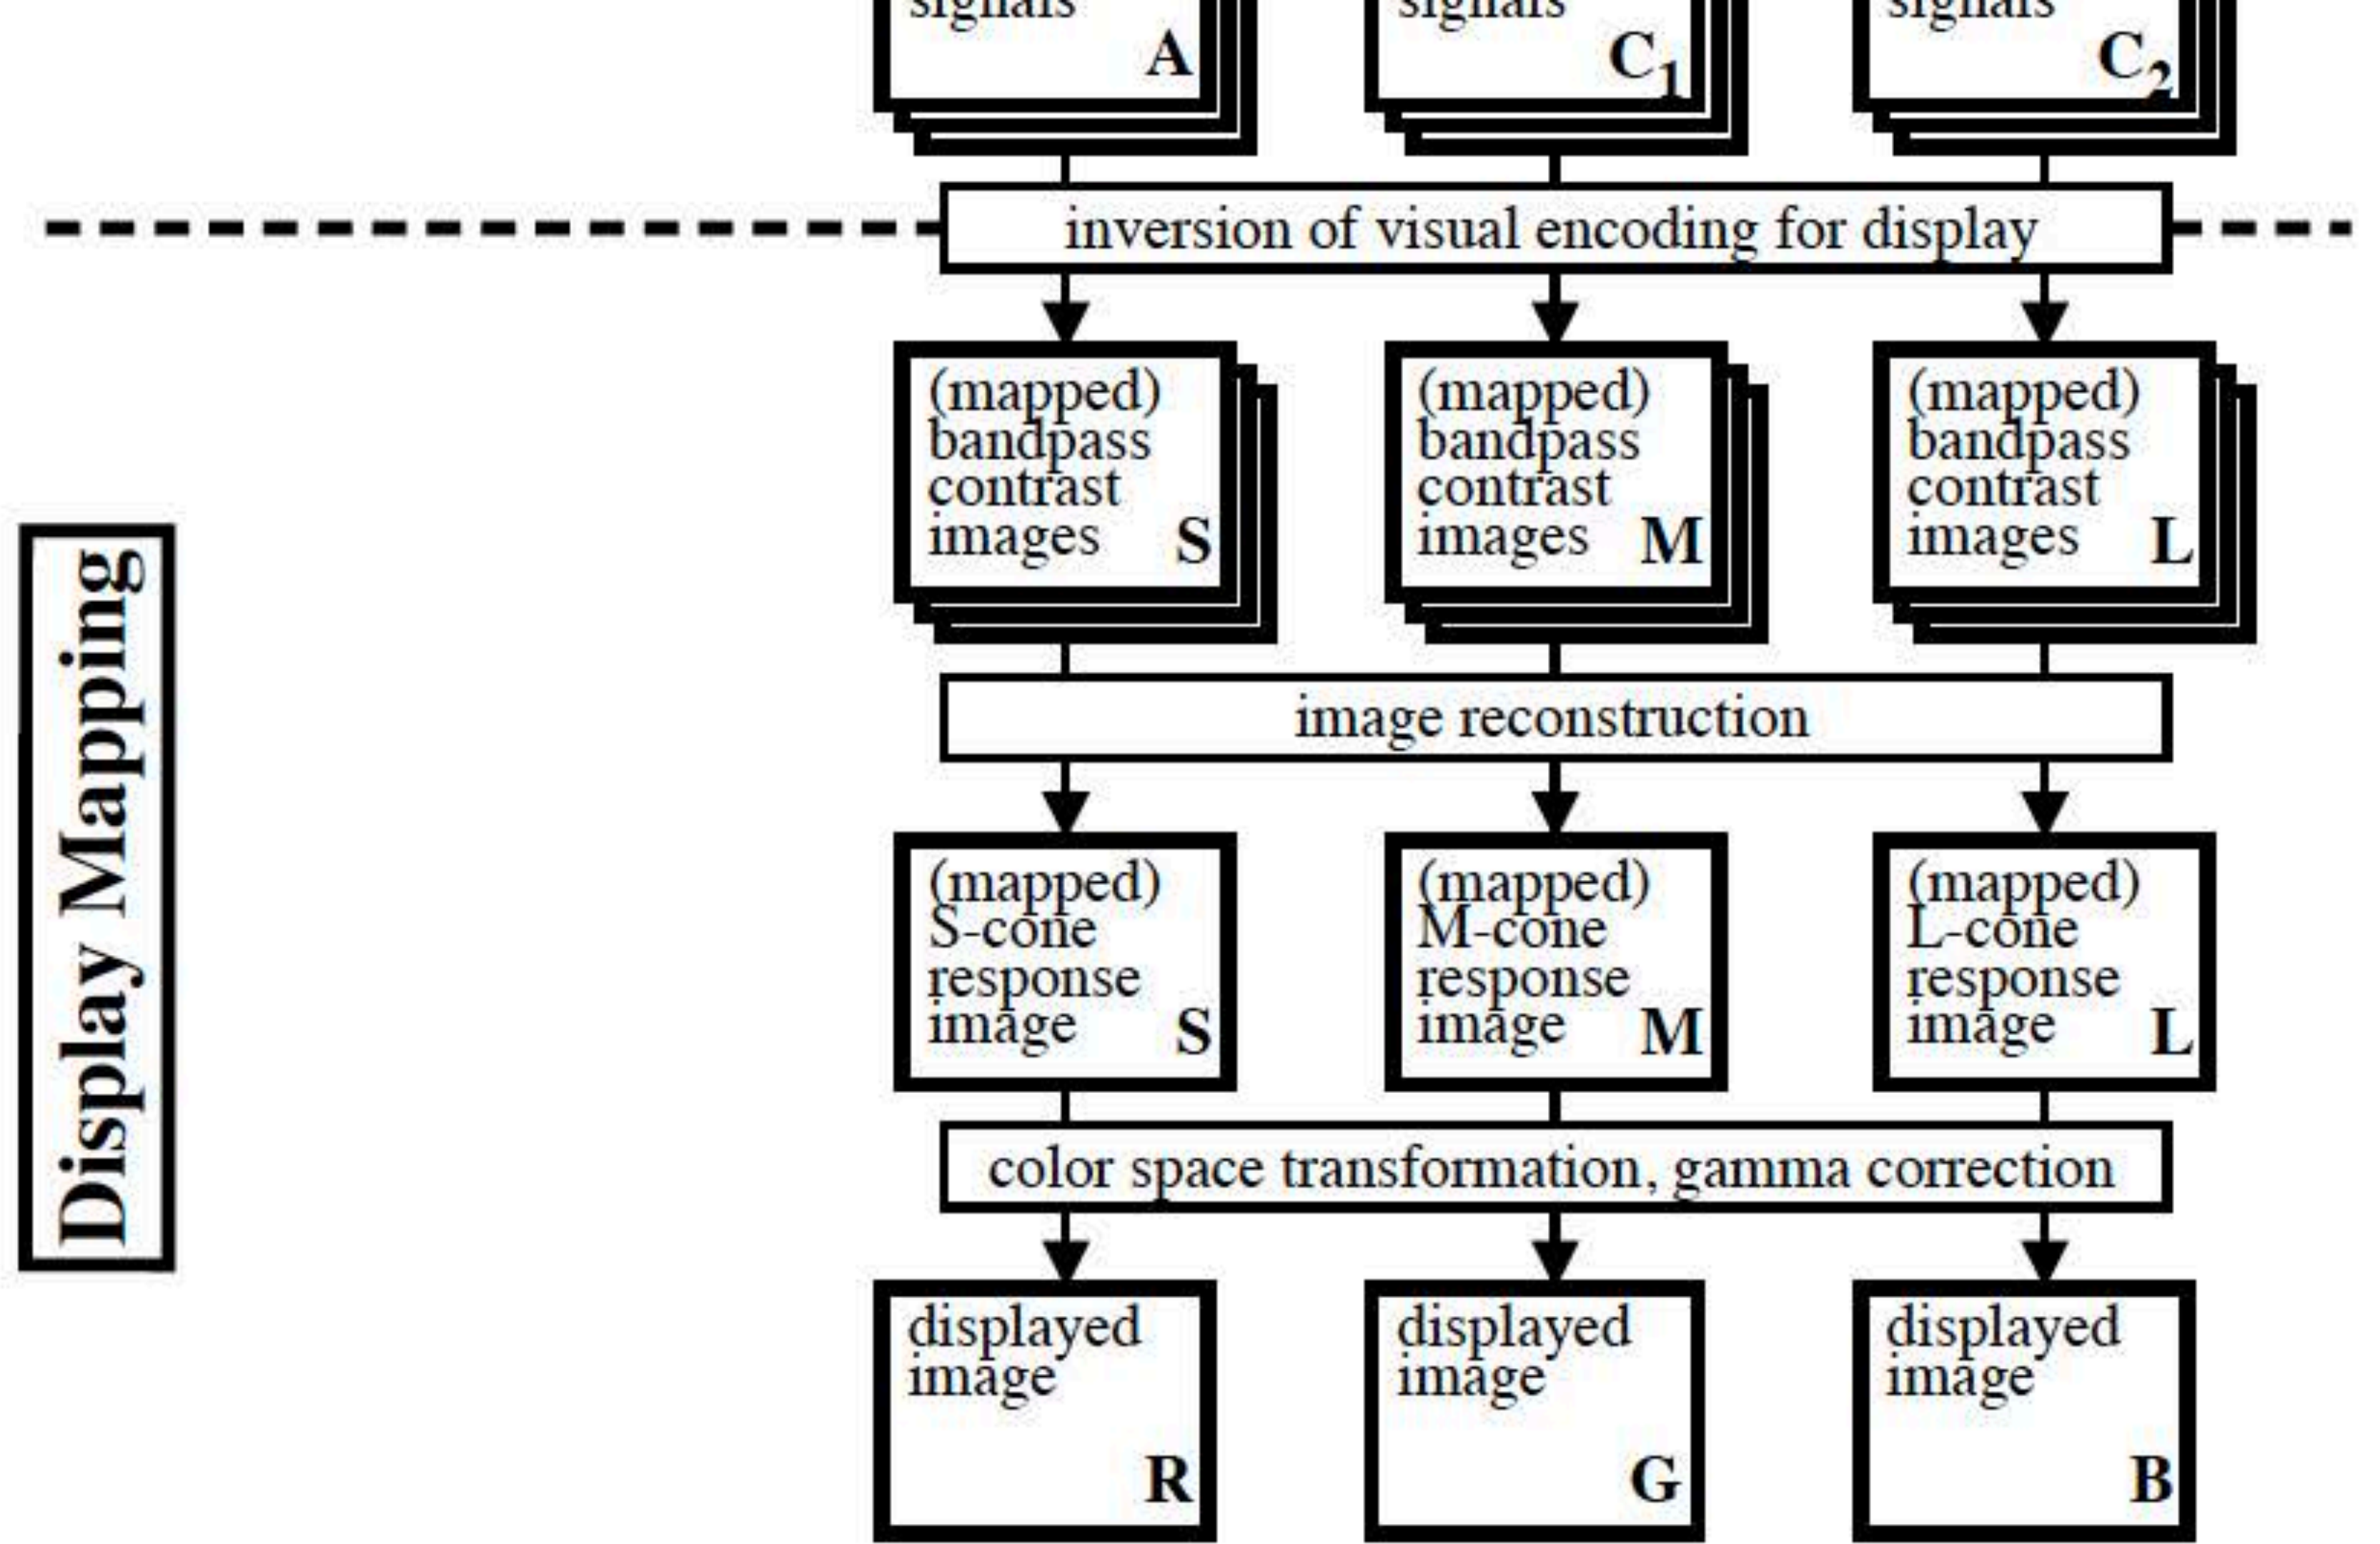


Figure 7: Flow chart of the computational model of adaptation and spatial vision for realistic tone reproduction.

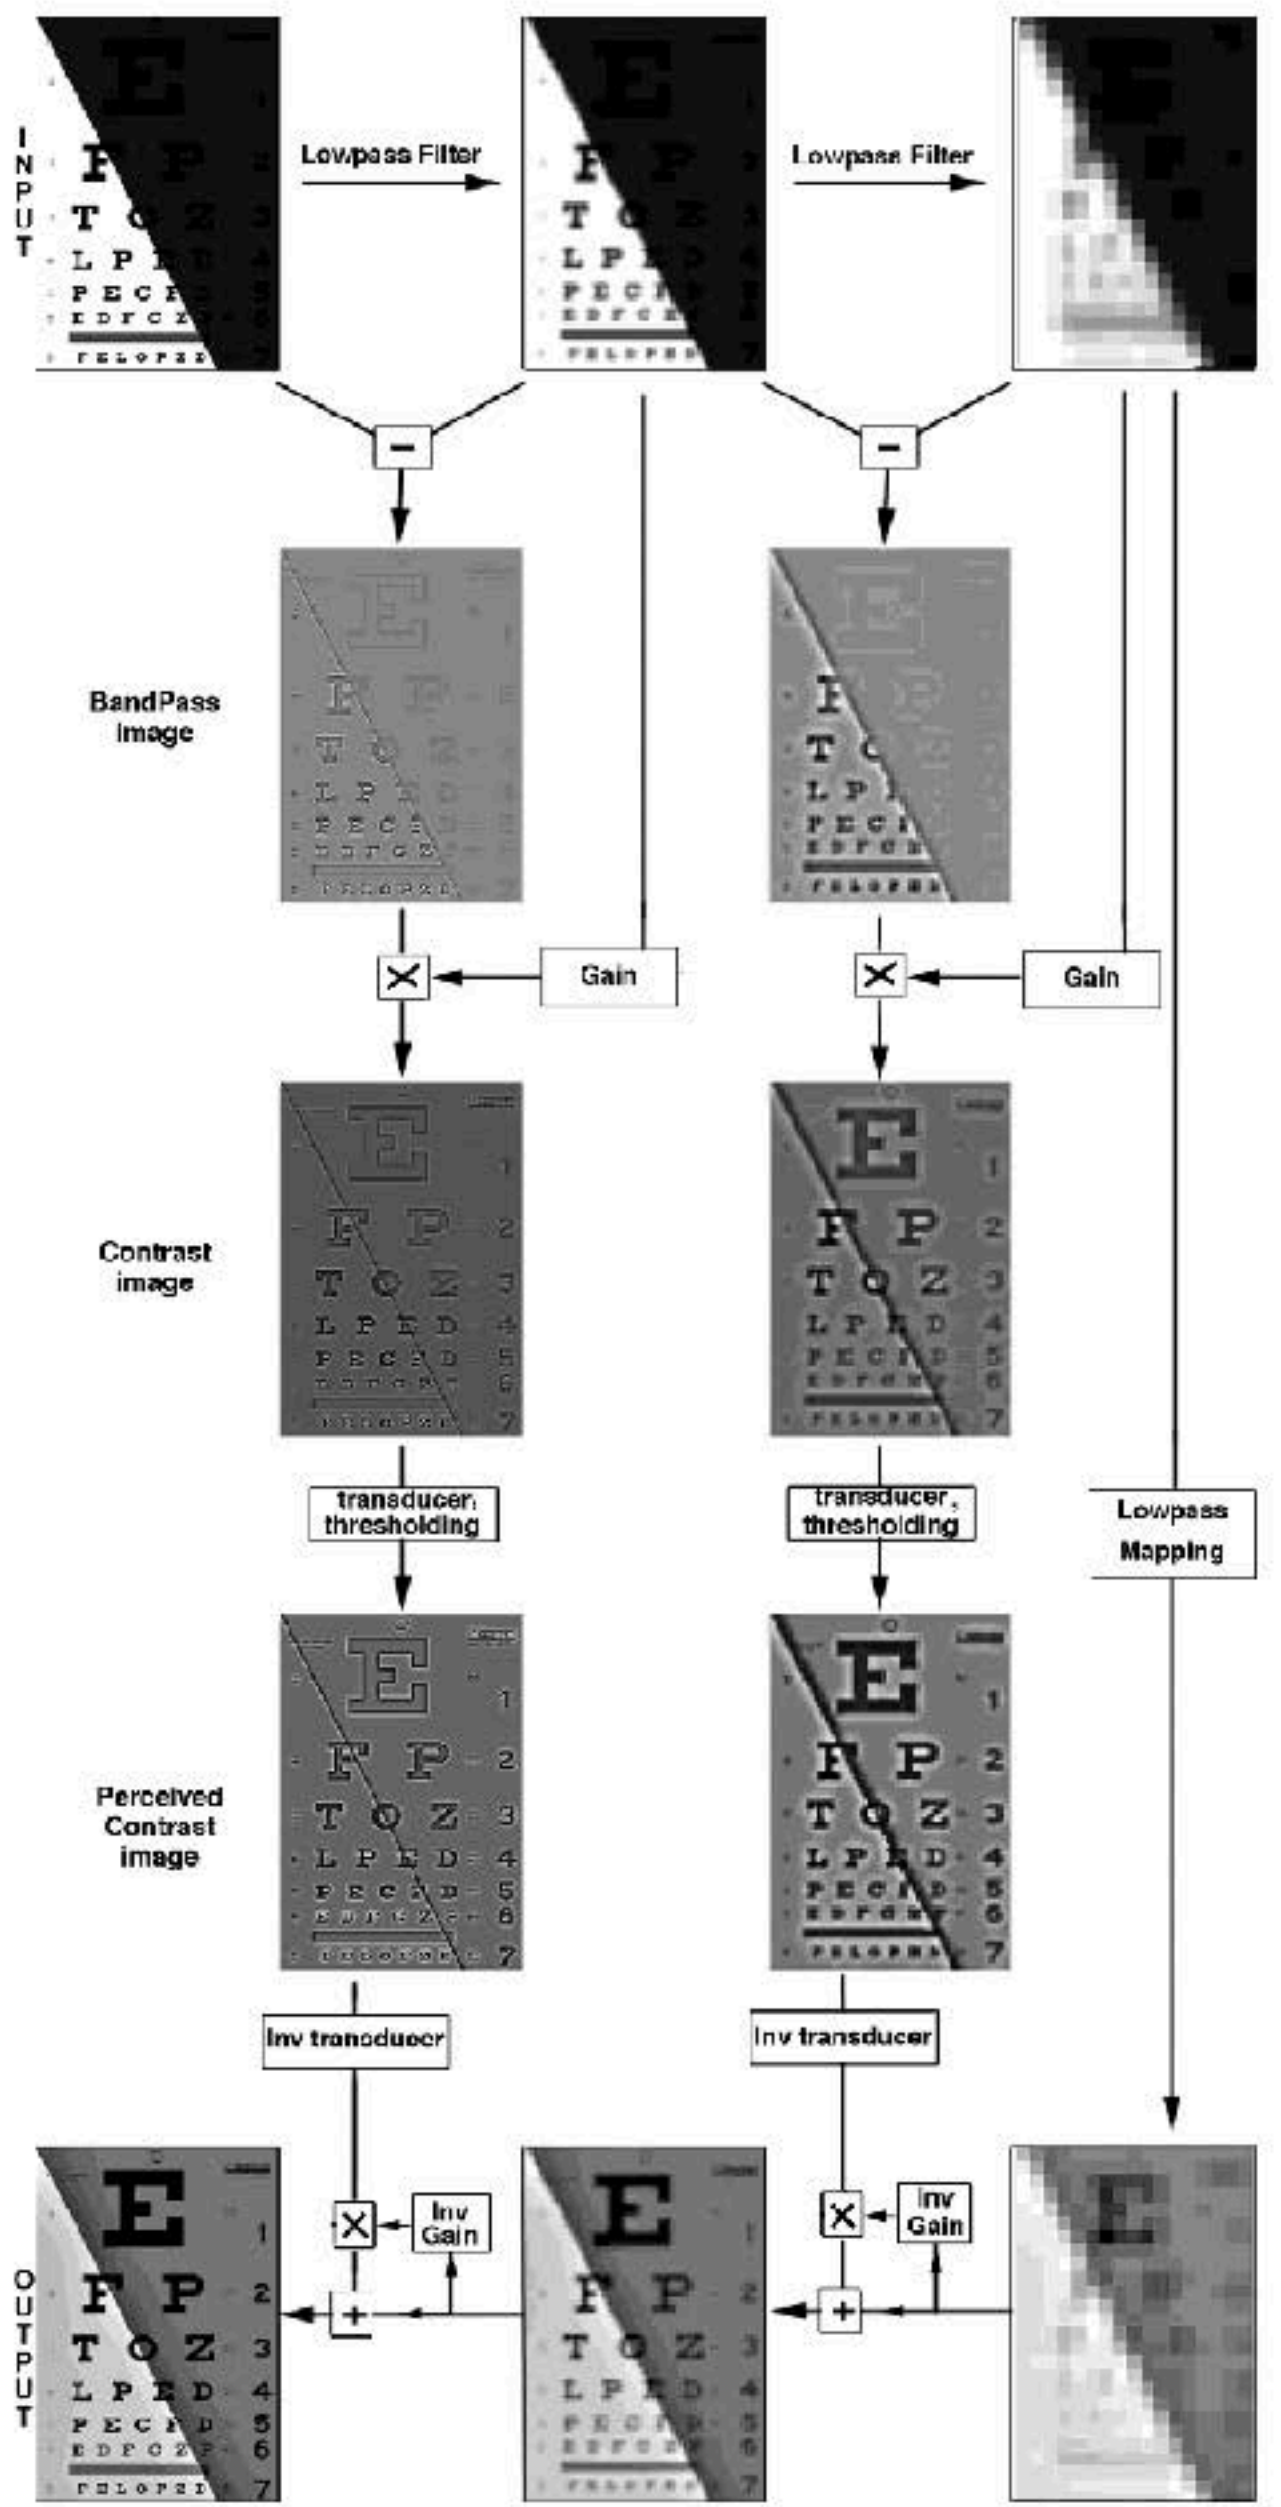



Figure 8: Pictorial representation of the computational model. Note that only 2 out of the 6 spatial mechanisms of one of the channels have been shown for the purpose of illustration. Original image is a Snellen chart with a 30:1 shadow boundary.



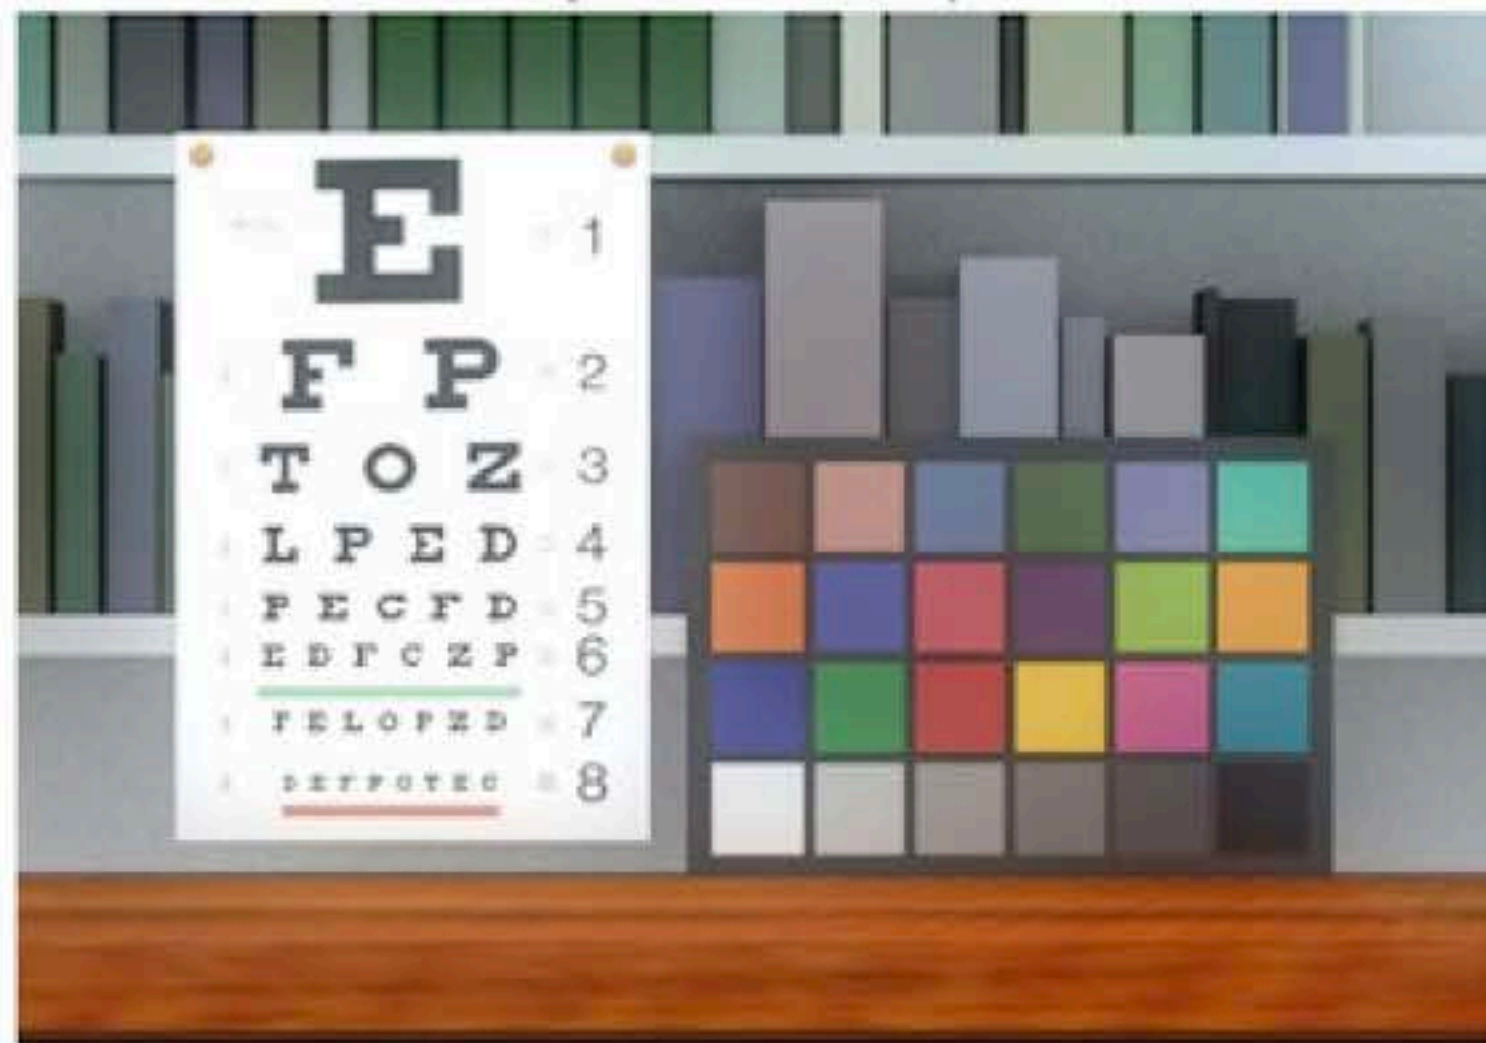
(0.1cd/m²)



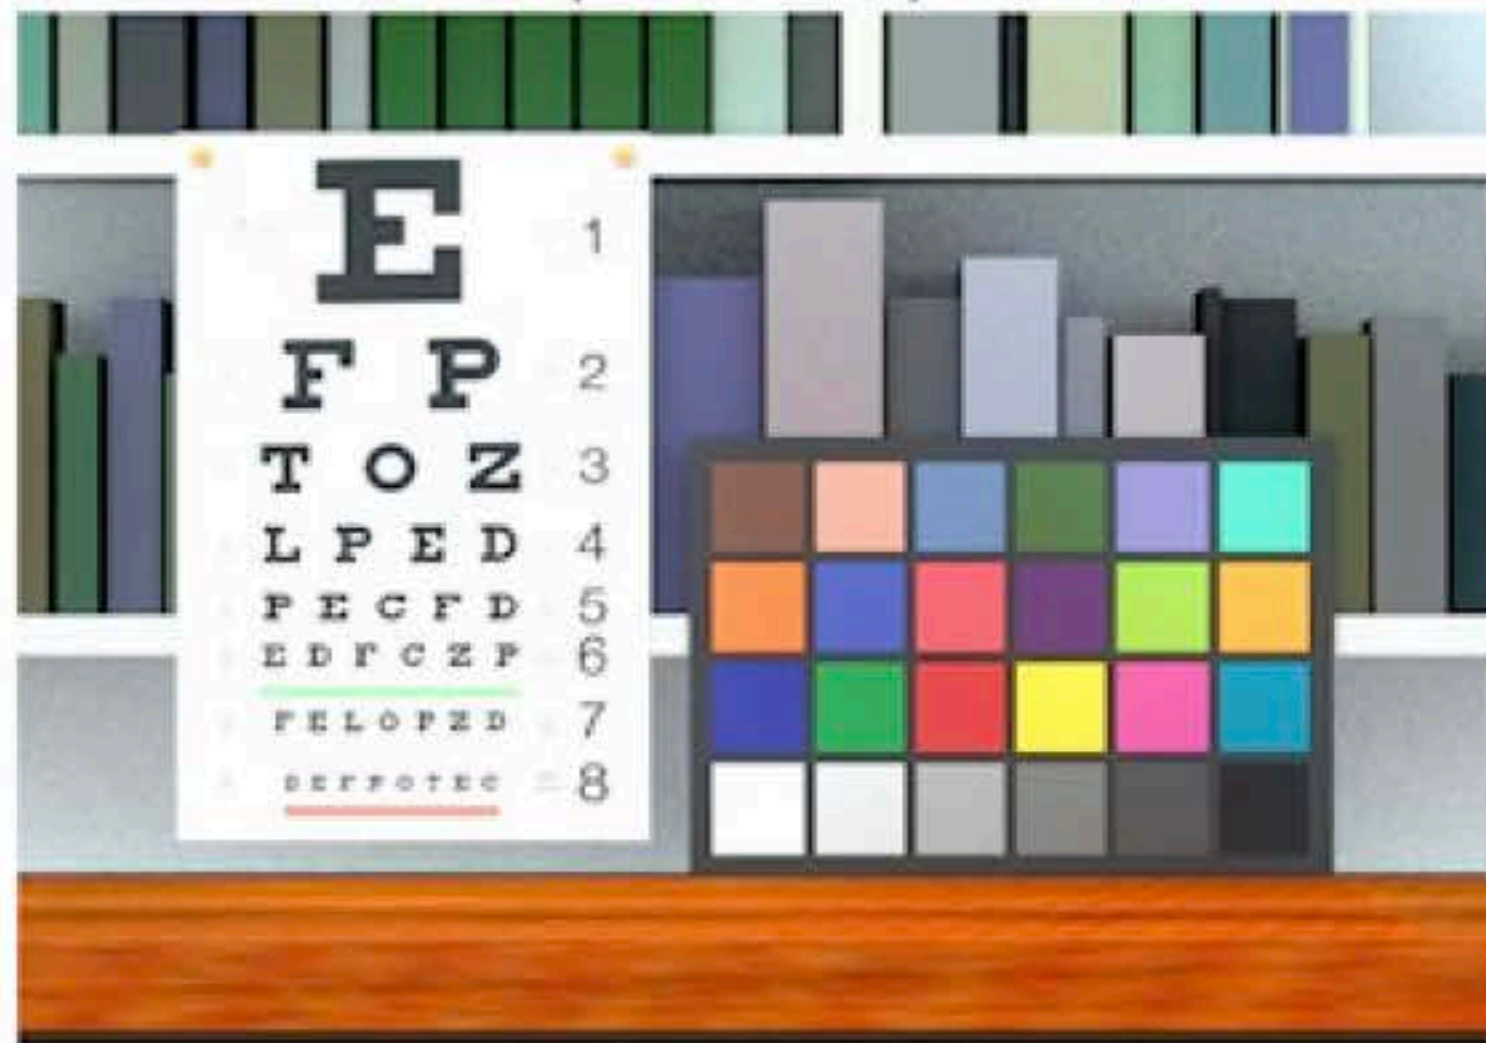
(1 cd/m²)




(10 cd/m²)



(100 cd/m²)



(1000 cd/m²)



(10,000 cd/m²)

Figure 9: Application of the model to a wide range of illumination levels.

reddish illuminant



tungsten illuminant



bluish illuminant



before adaptation

after adaptation



Figure 10: Illustration of chromatic adaptation.








Figure 11: Tone mapping of high-dynamic range images. The images on the top are linear mappings of the original high-dynamic range images. The images on the bottom are the mapping obtained by application of the visual model.

iCAM

iCAM framework for image appearance, differences, and quality

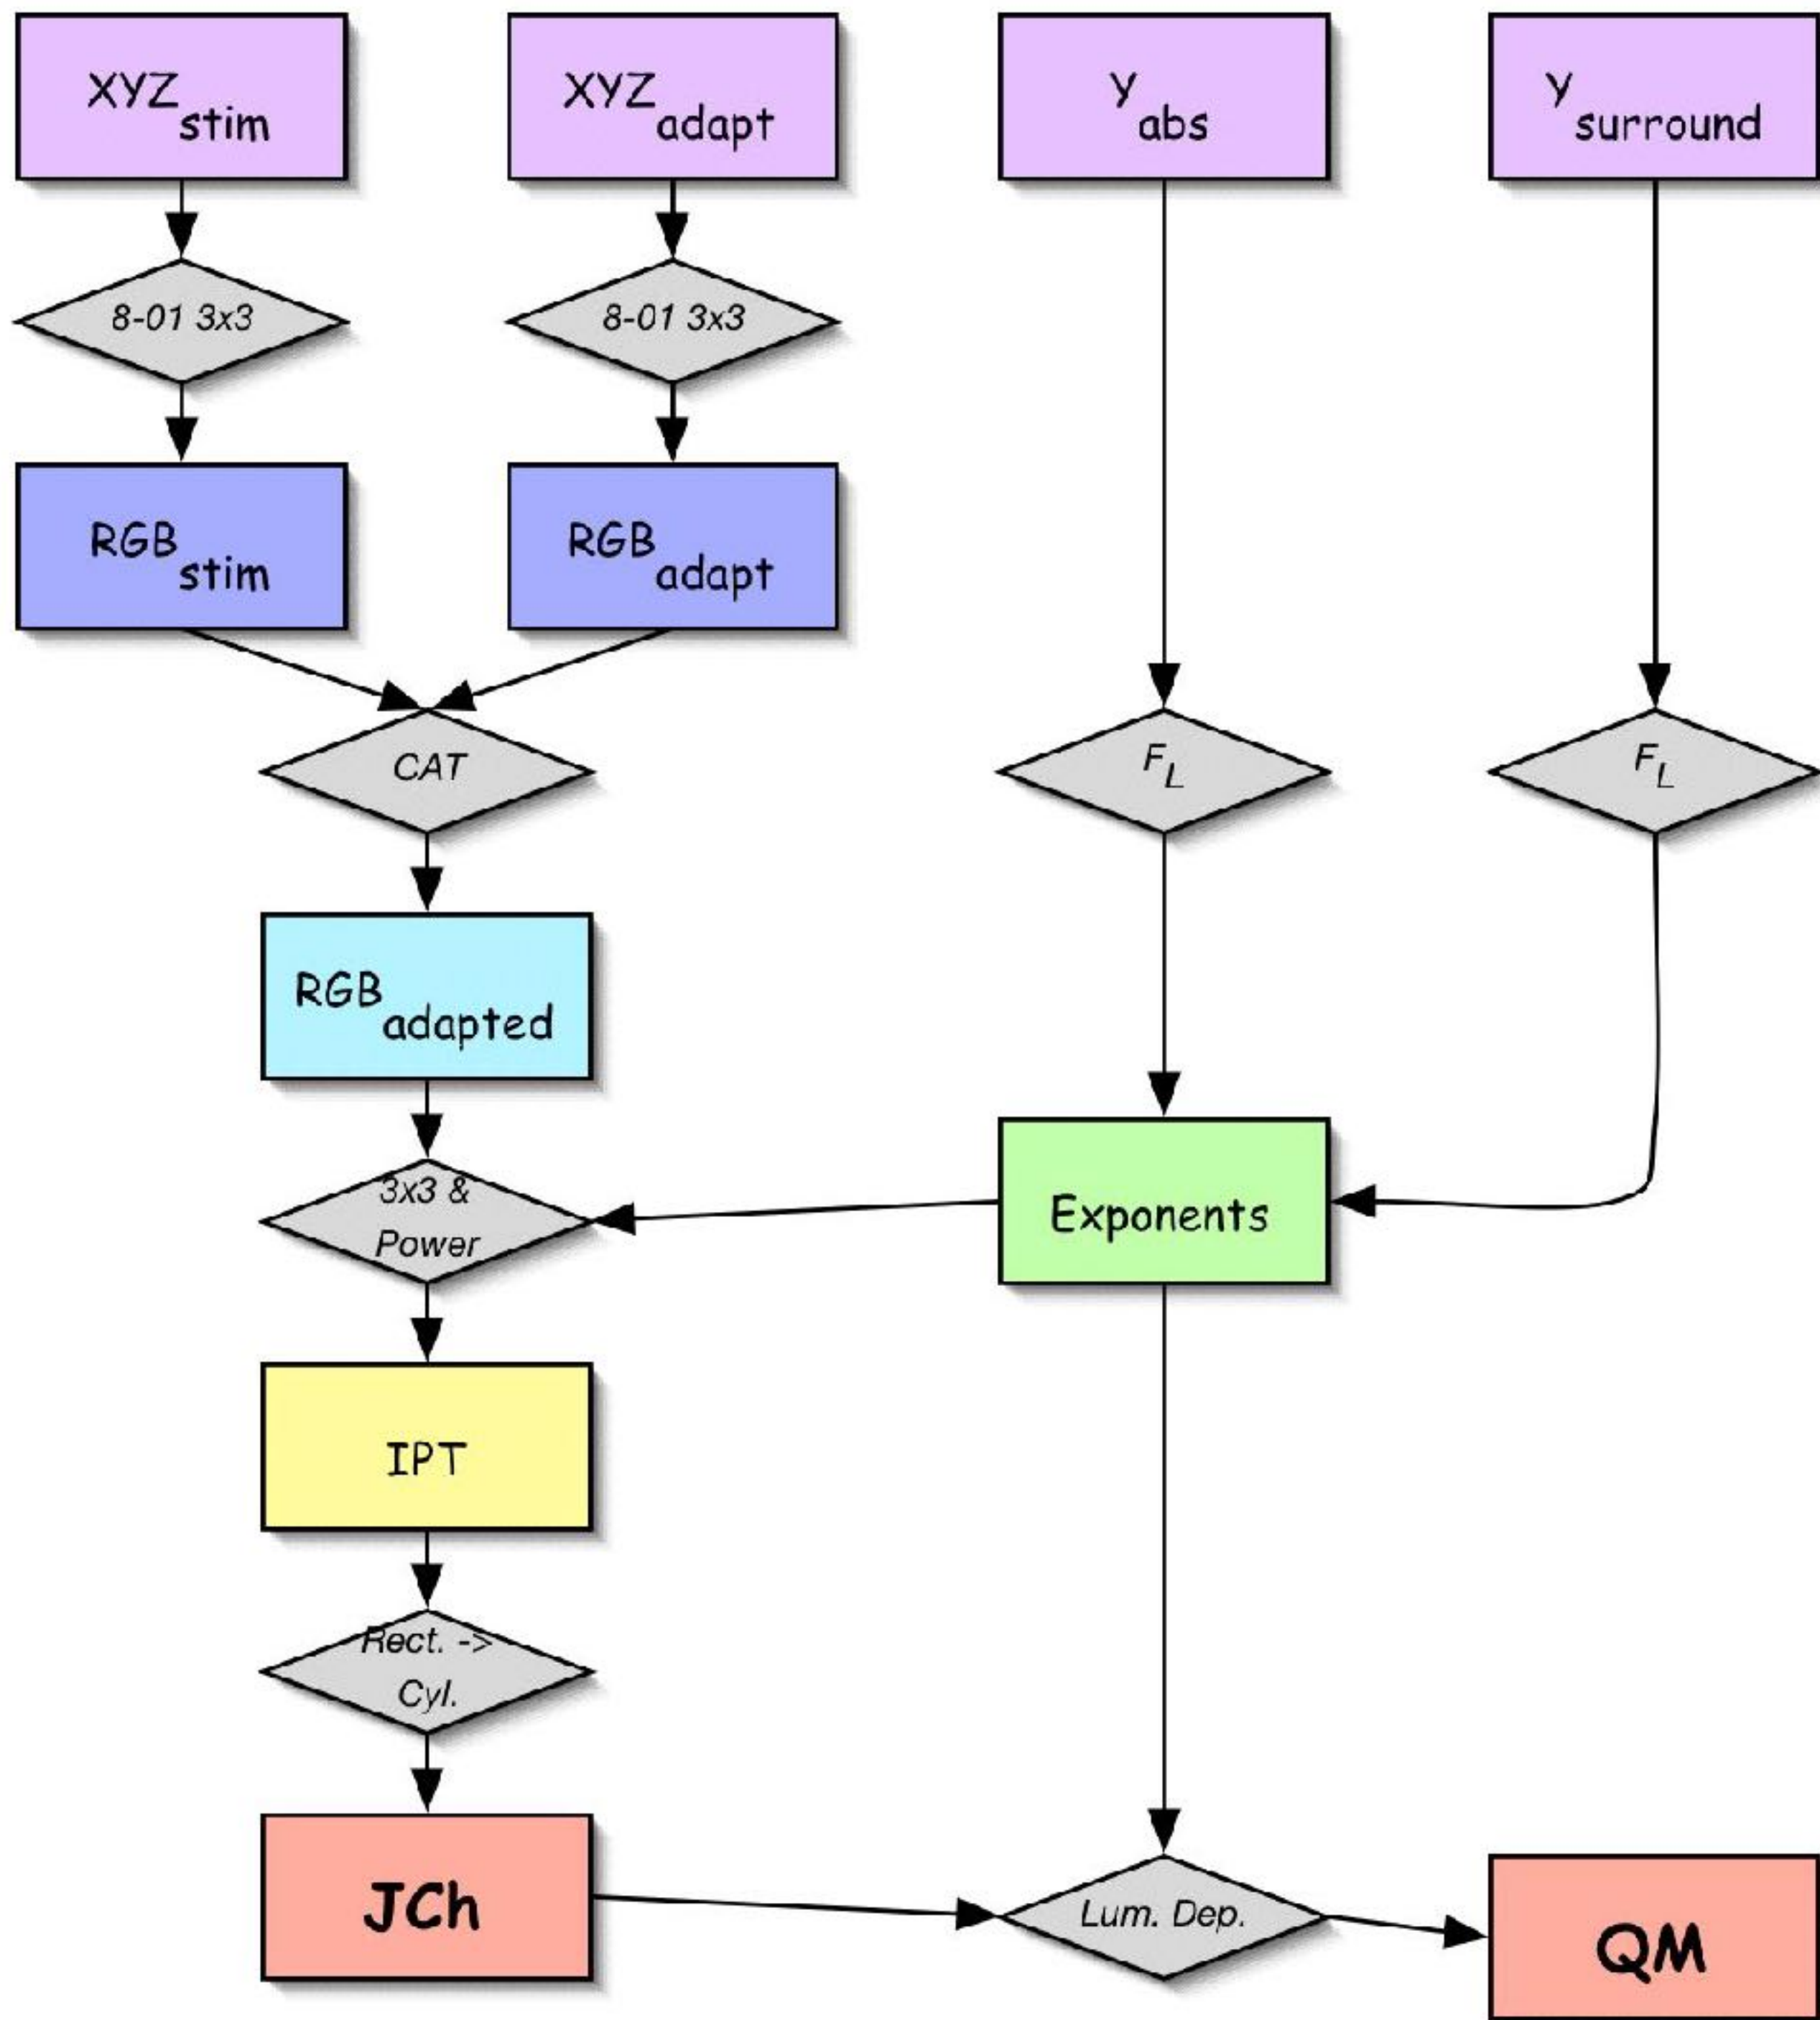
Mark D. Fairchild

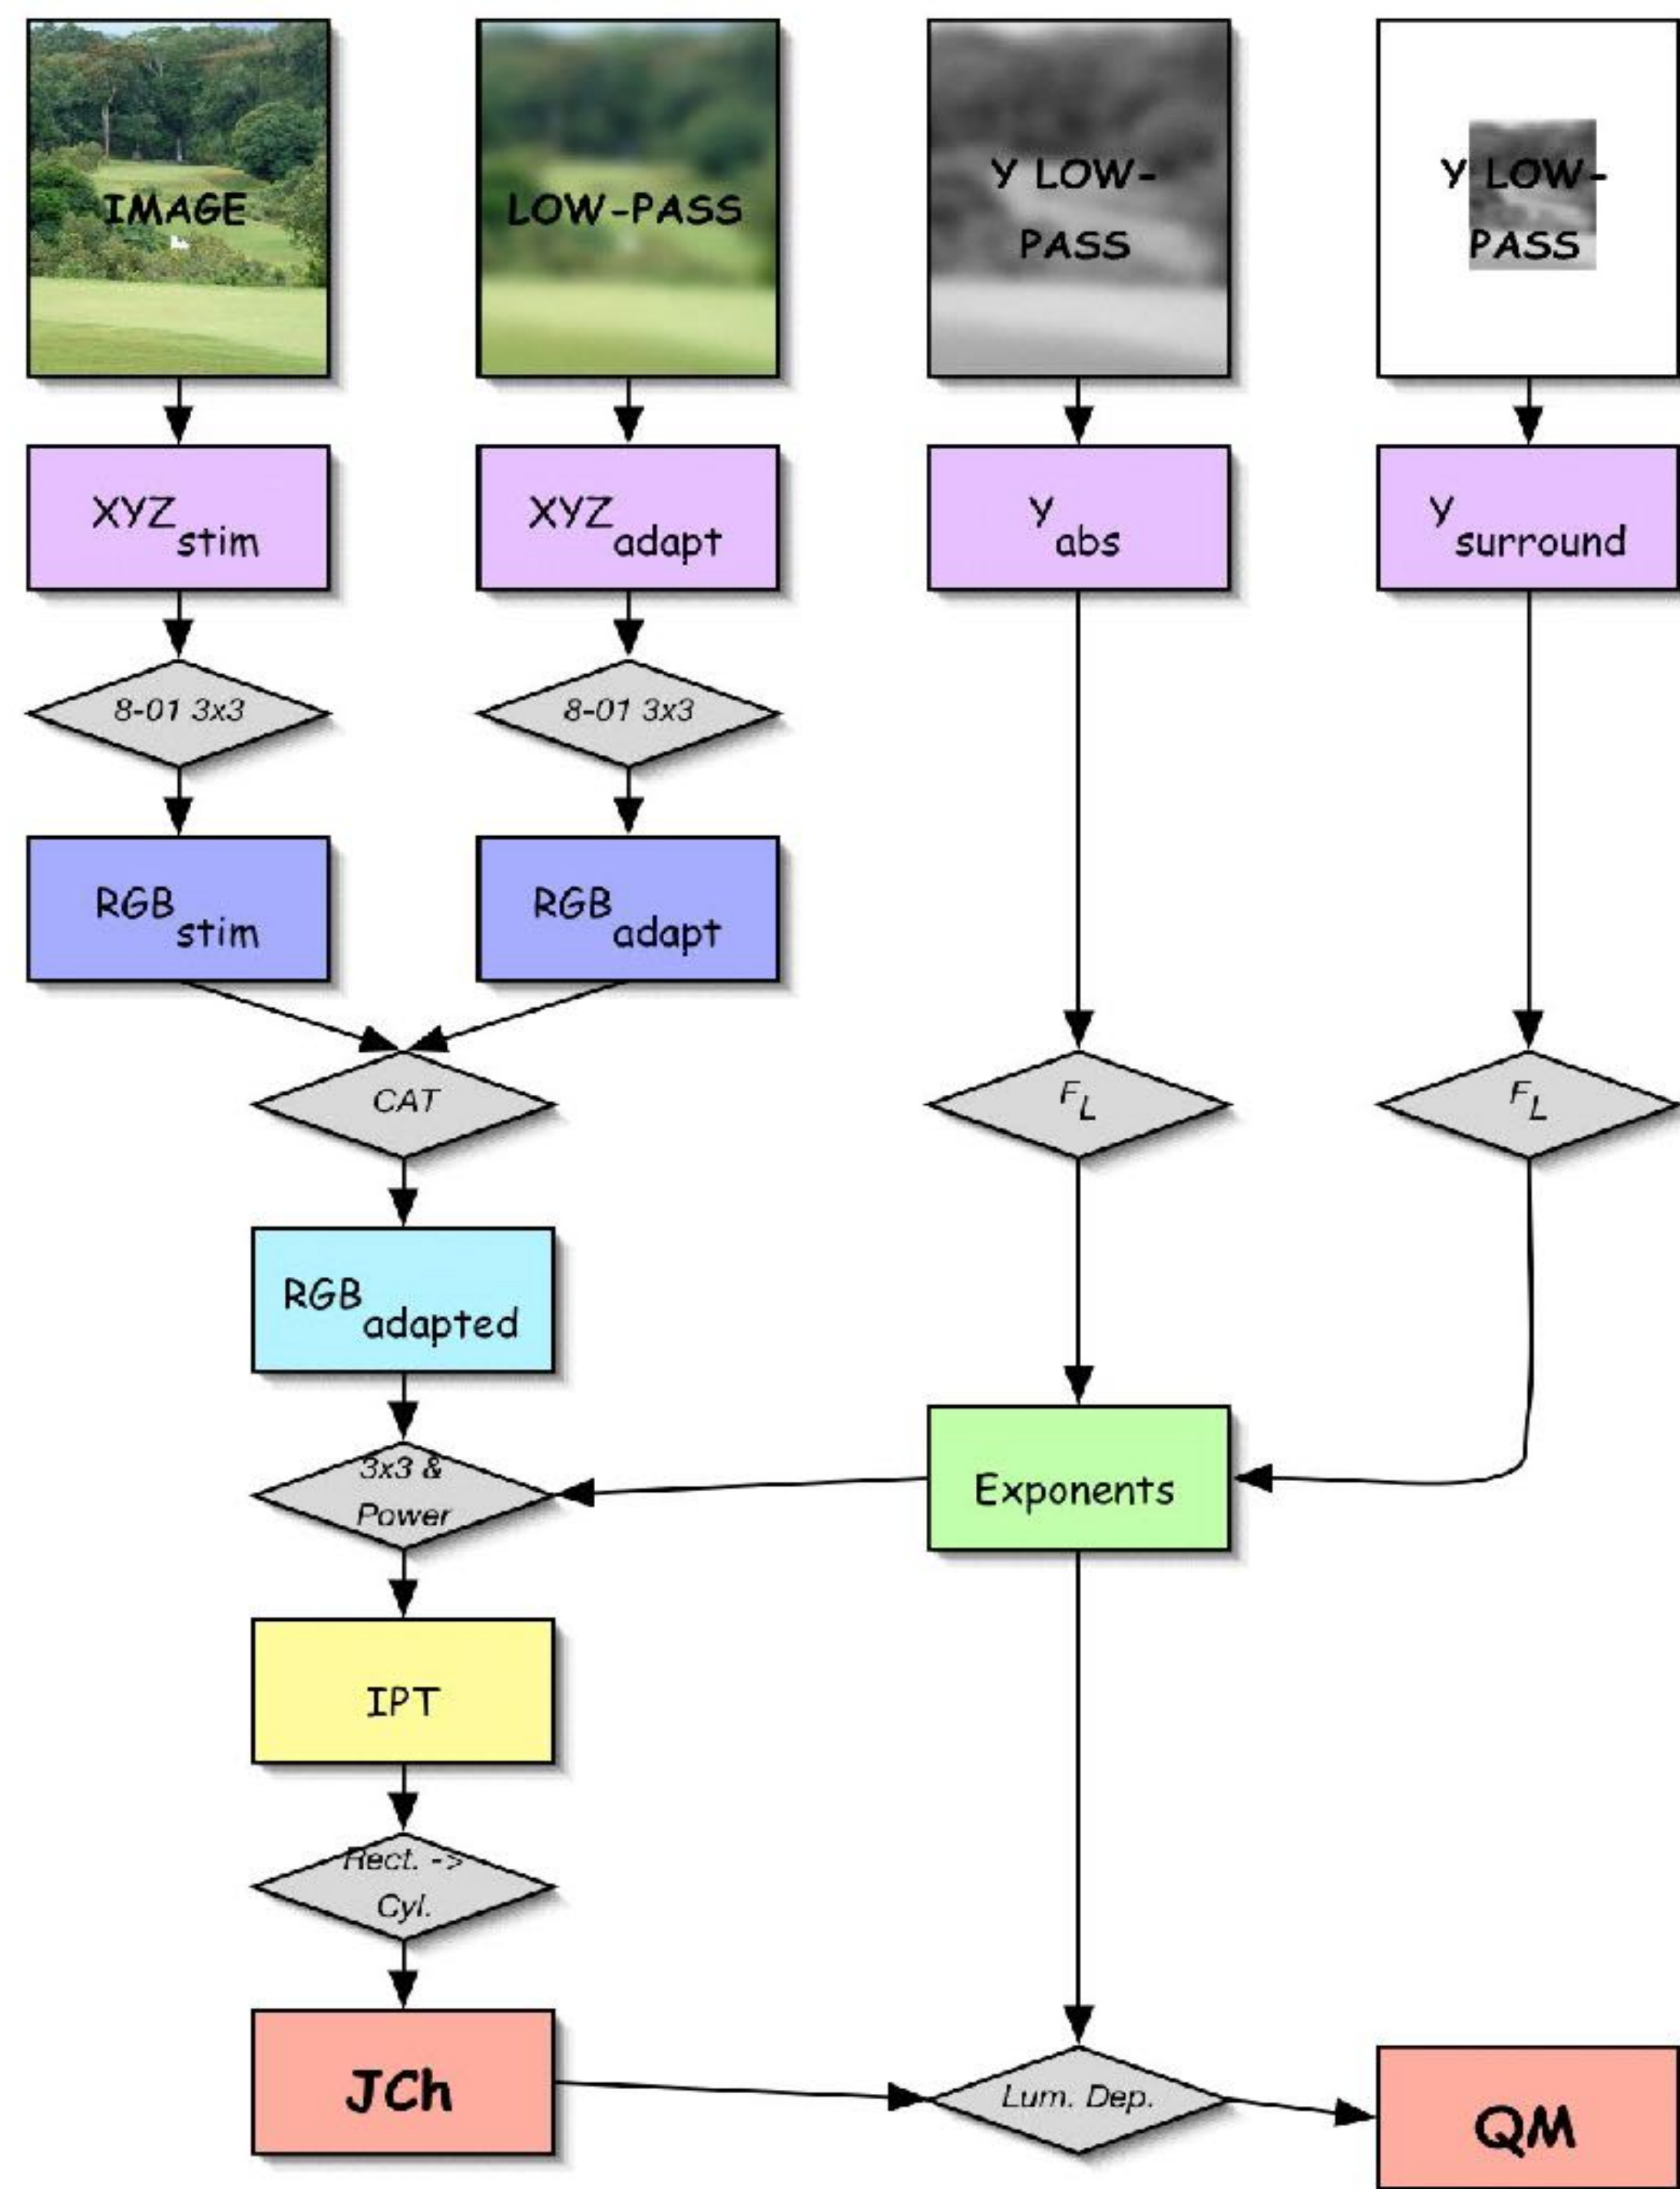
Garrett M. Johnson


Rochester Institute of Technology
Chester F. Carlson Center for Imaging Science
Munsell Color Science Laboratory
54 Lomb Memorial Drive
Rochester, New York 14623-5604
E-mail: mdf@cis.rit.edu


Abstract. *Traditional color appearance modeling has recently matured to the point that available, internationally recommended models such as CIECAM02 are capable of making a wide range of predictions, to within the observer variability in color matching and color scaling of stimuli, in somewhat simplified viewing conditions. It is proposed that the next significant advances in the field of color appearance modeling and image quality metrics will not come from evolutionary revisions of colorimetric color appearance models alone. Instead, a more revolutionary approach will be required to make appearance and difference predictions for more complex stimuli in a wider array of viewing conditions. Such an approach can*

the inks as well as measures of the dot area coverage for halftone systems. In electronic systems like television, system measurements such as signal voltages were used to colorimetrically quantify the imaging system.¹ It should be noted that vision-based measurements of imaging systems for image quality do have a long history, as illustrated by the example of Schade's pioneering work.² As imaging systems evolved in complexity and openness, the need for device-independent image measures became clear.

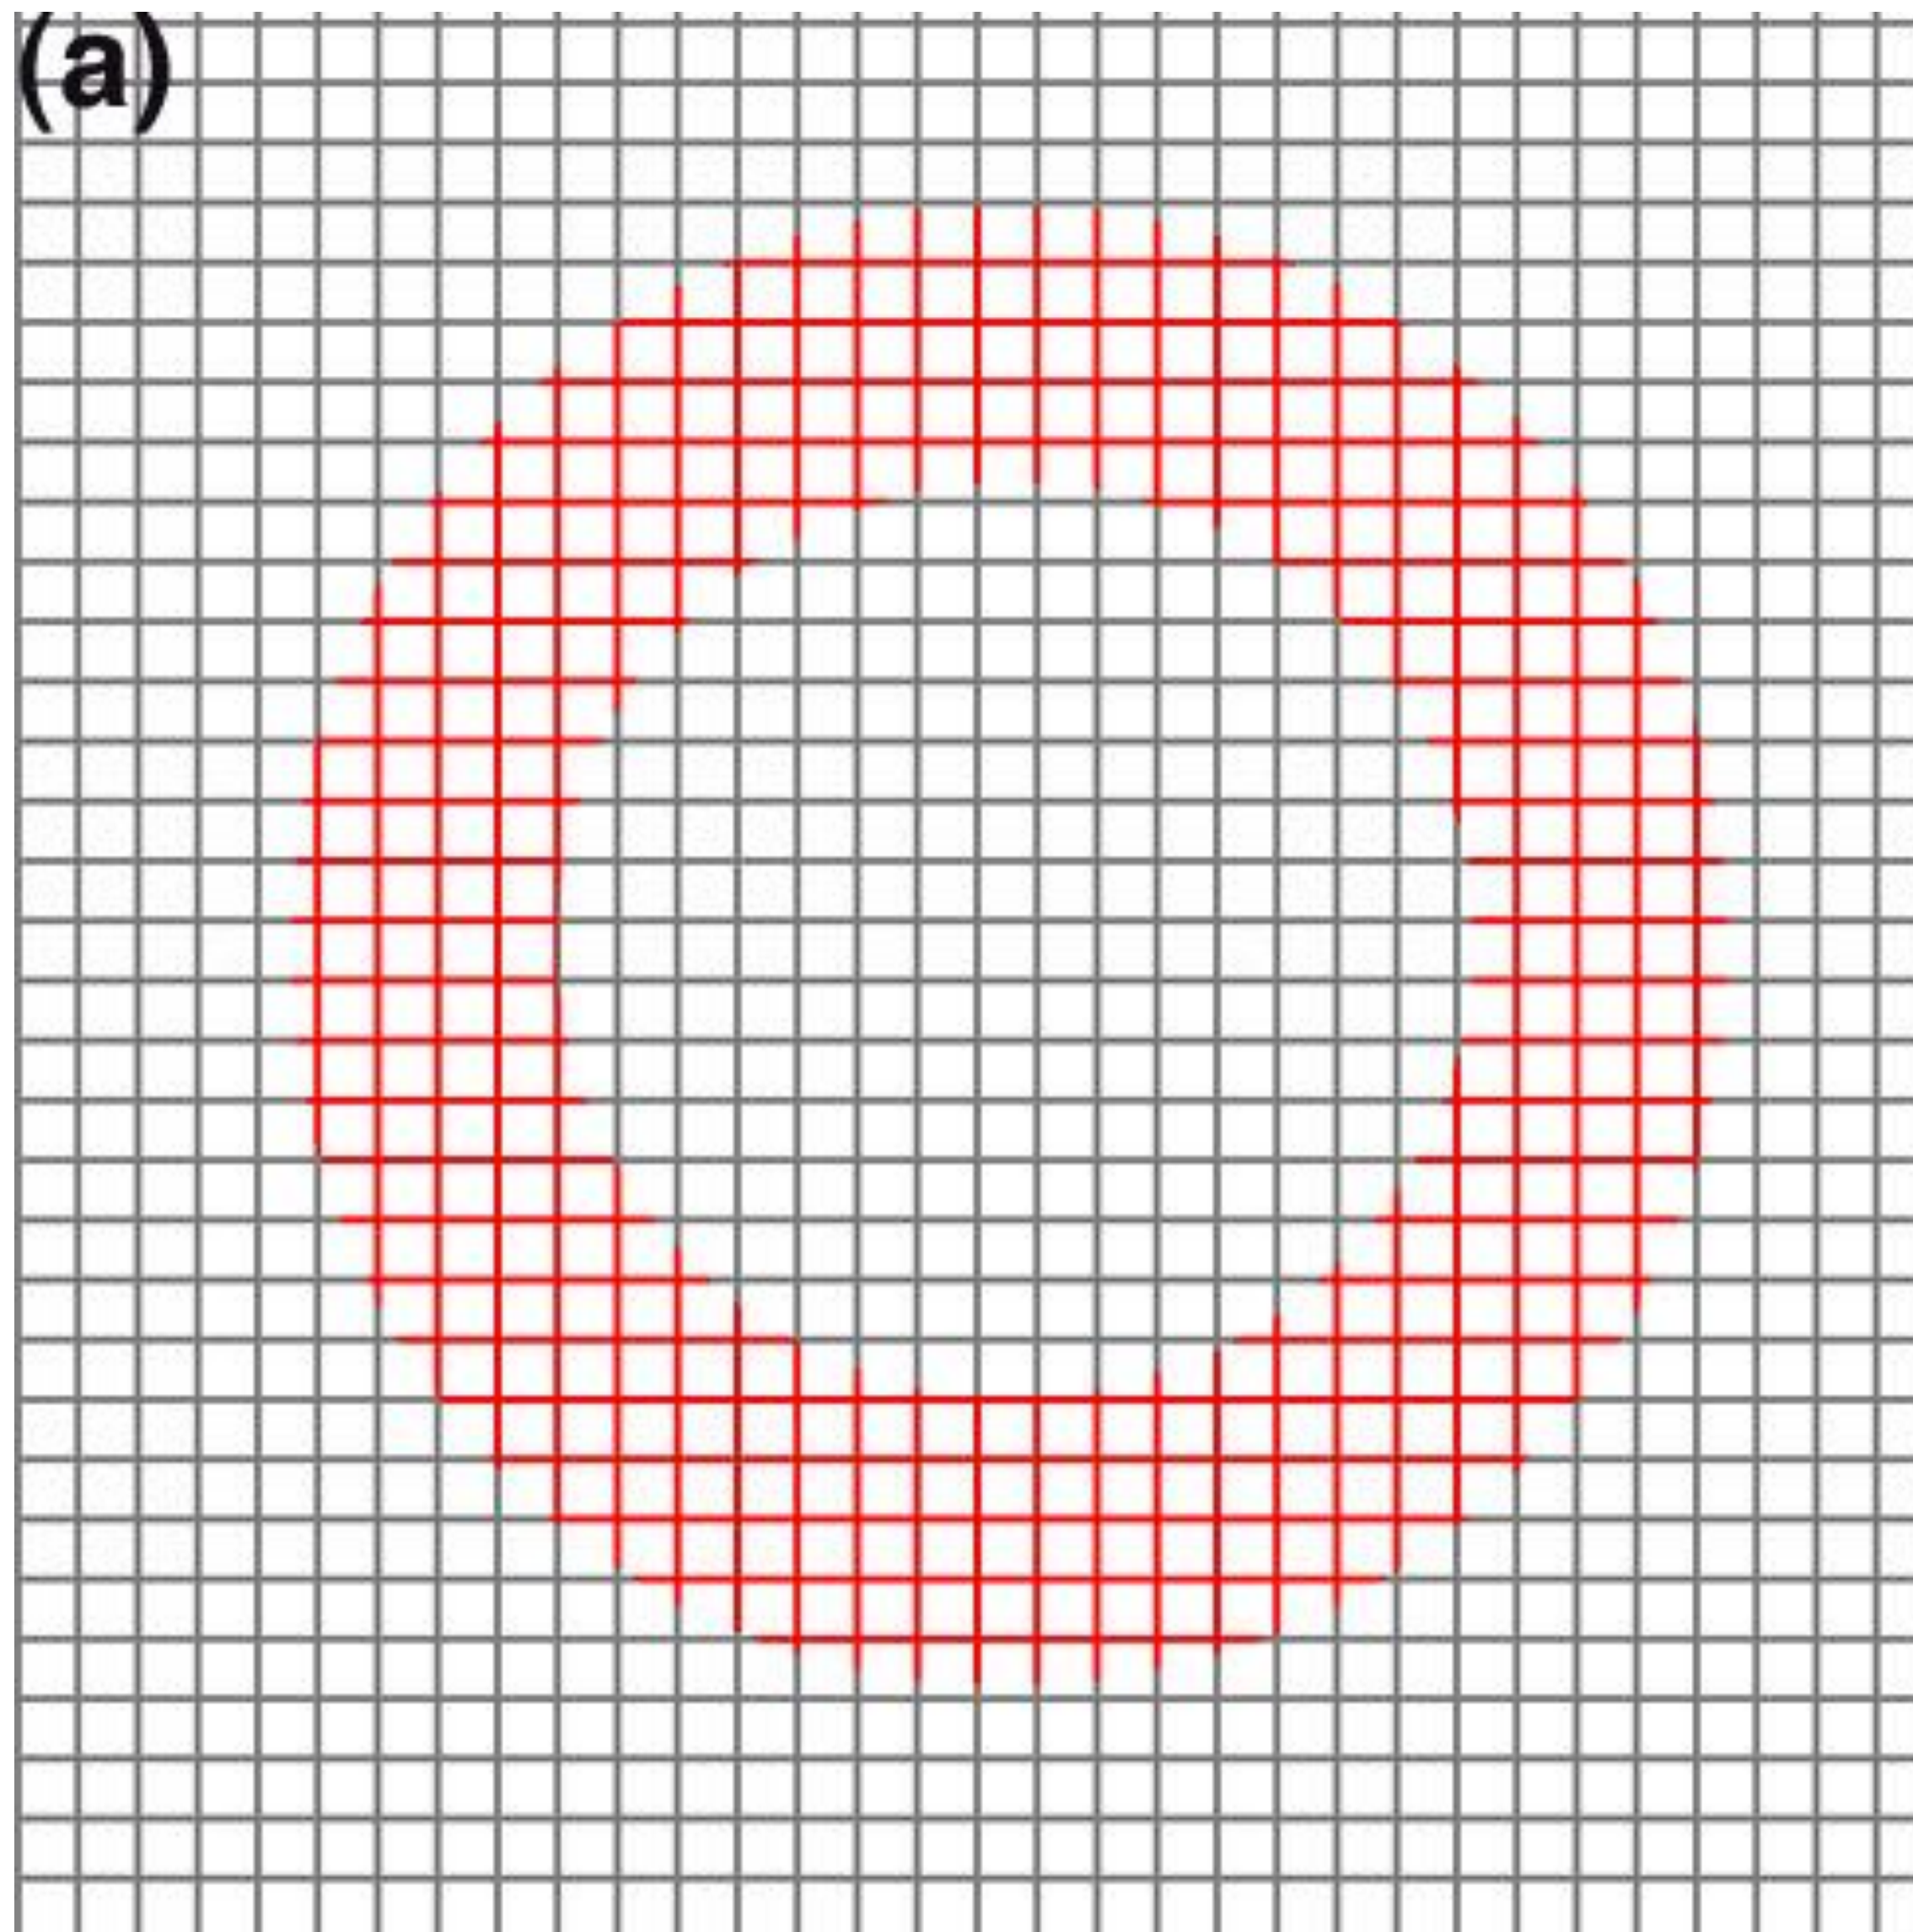






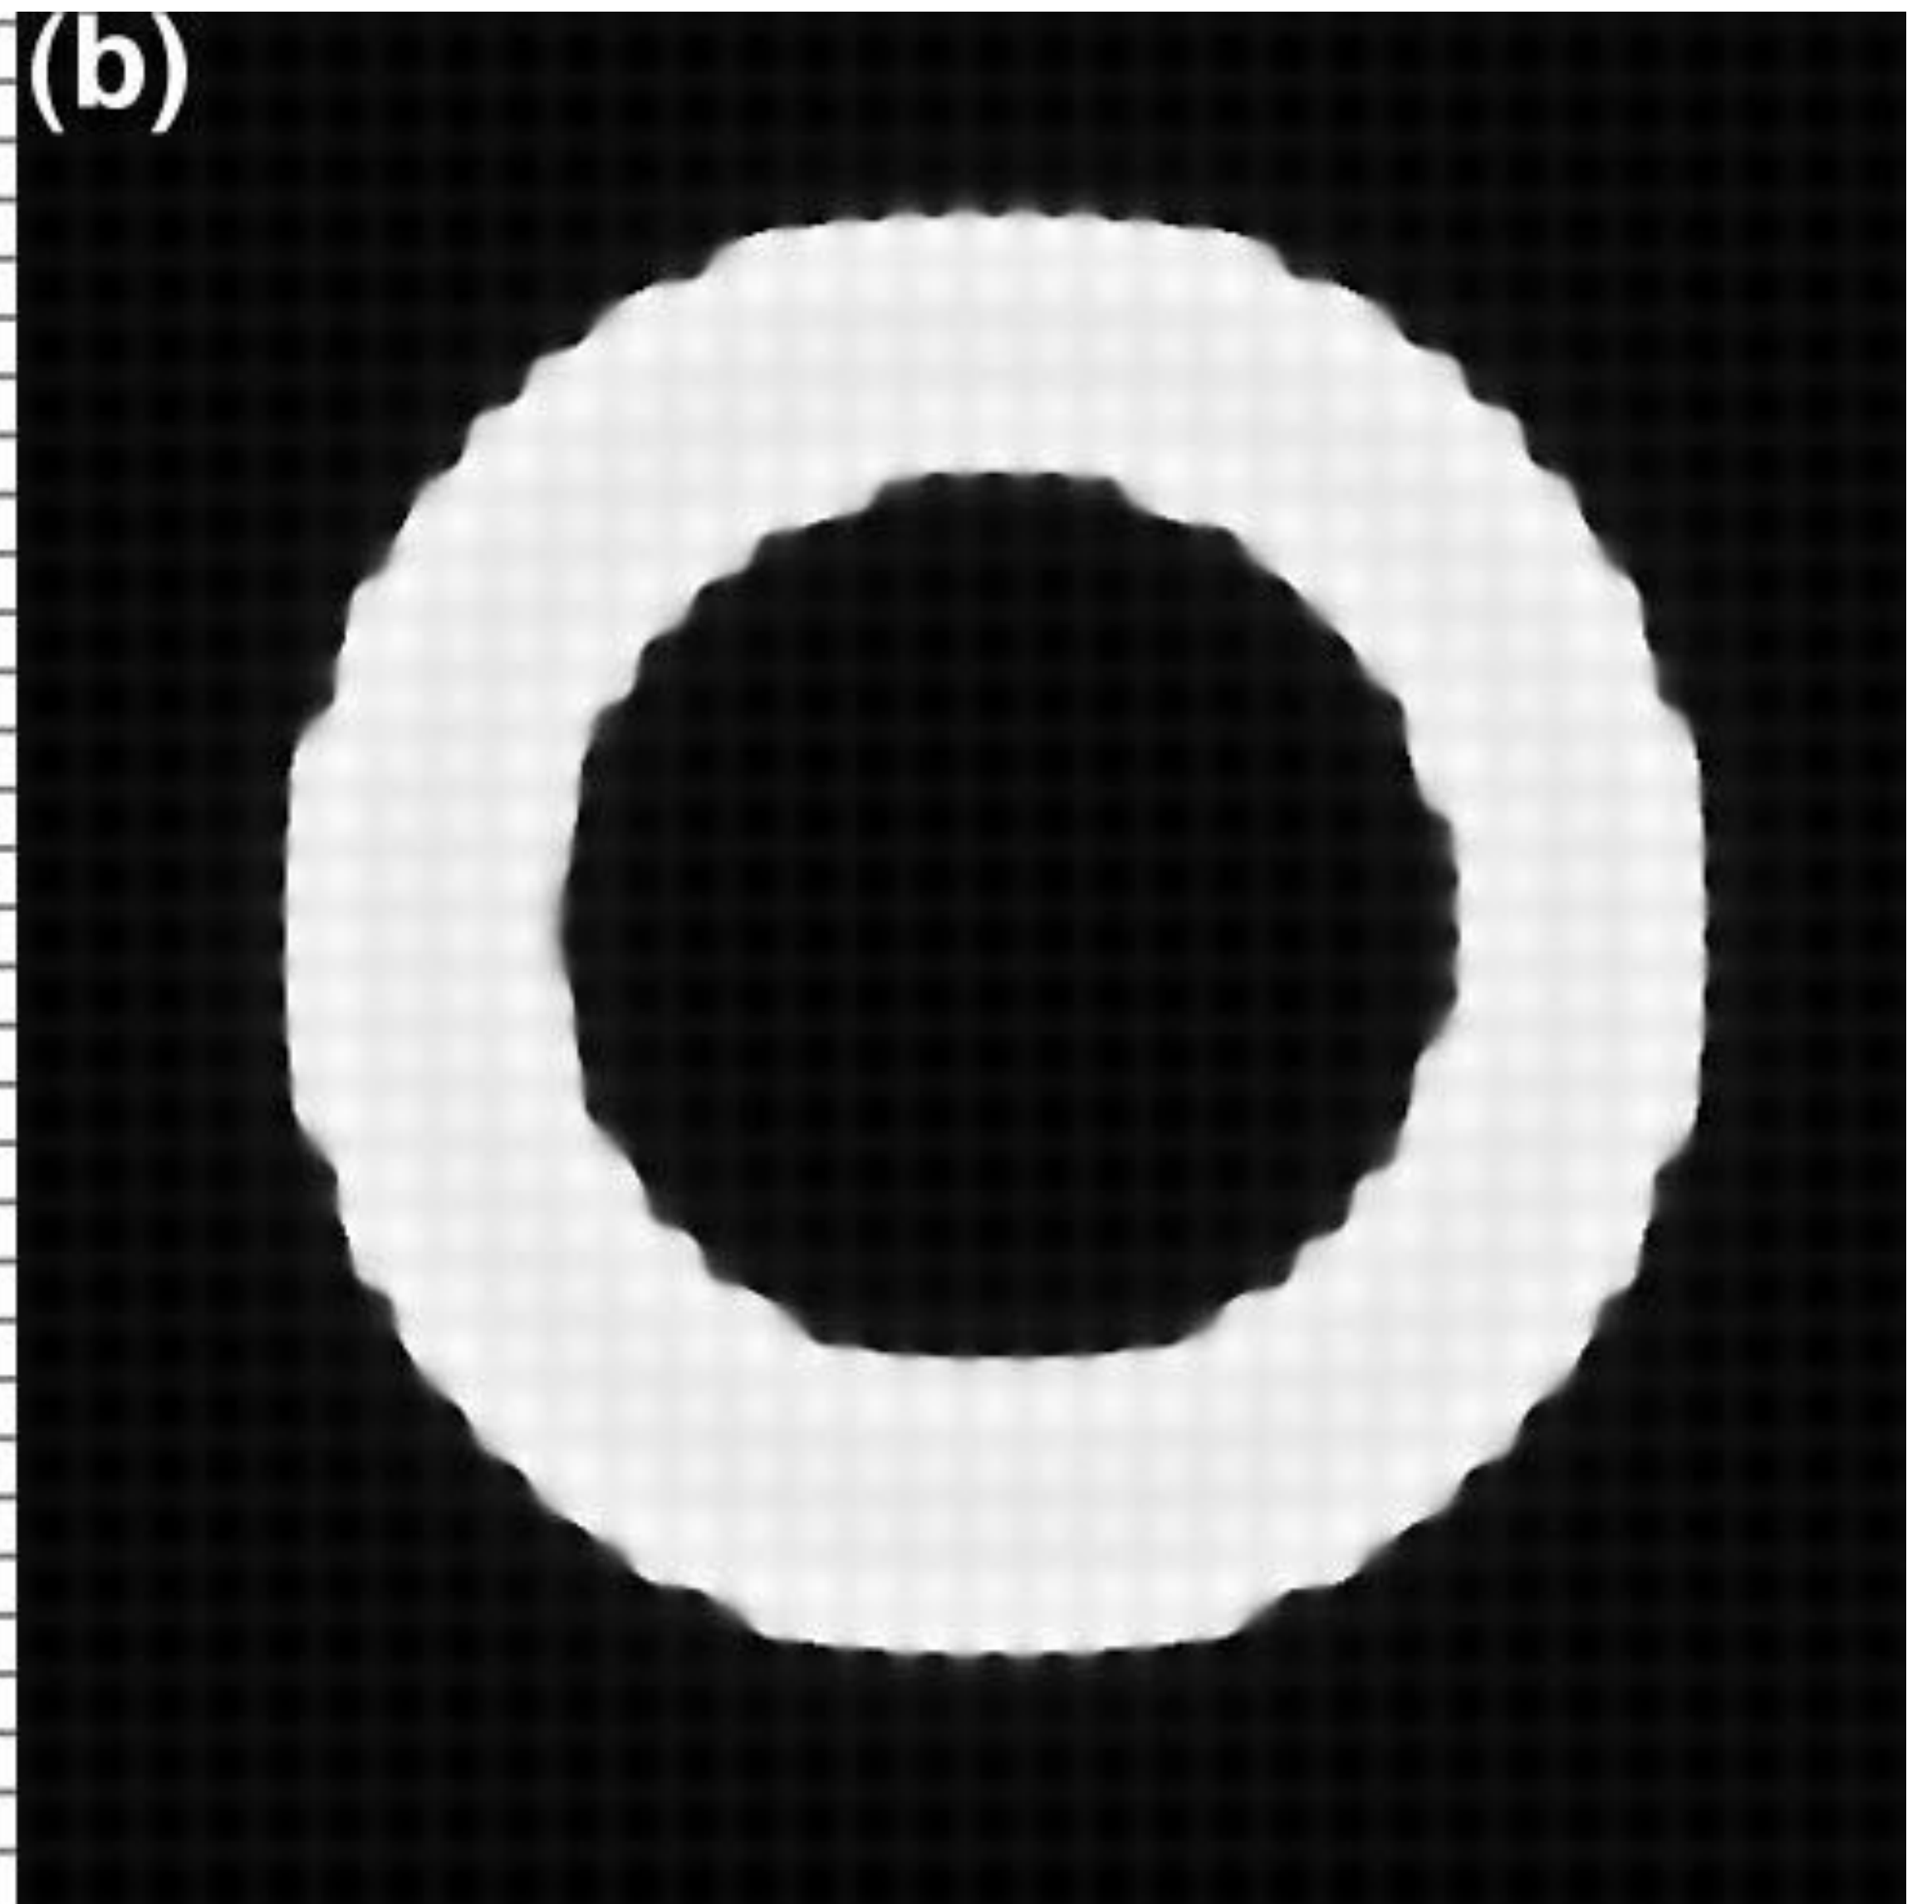


(a)



Stimulus


(b)



Hue Map



Two Input Images are Given, an Original and
Reproduction



The Input Images are Transformed into an Opponent
Color Space



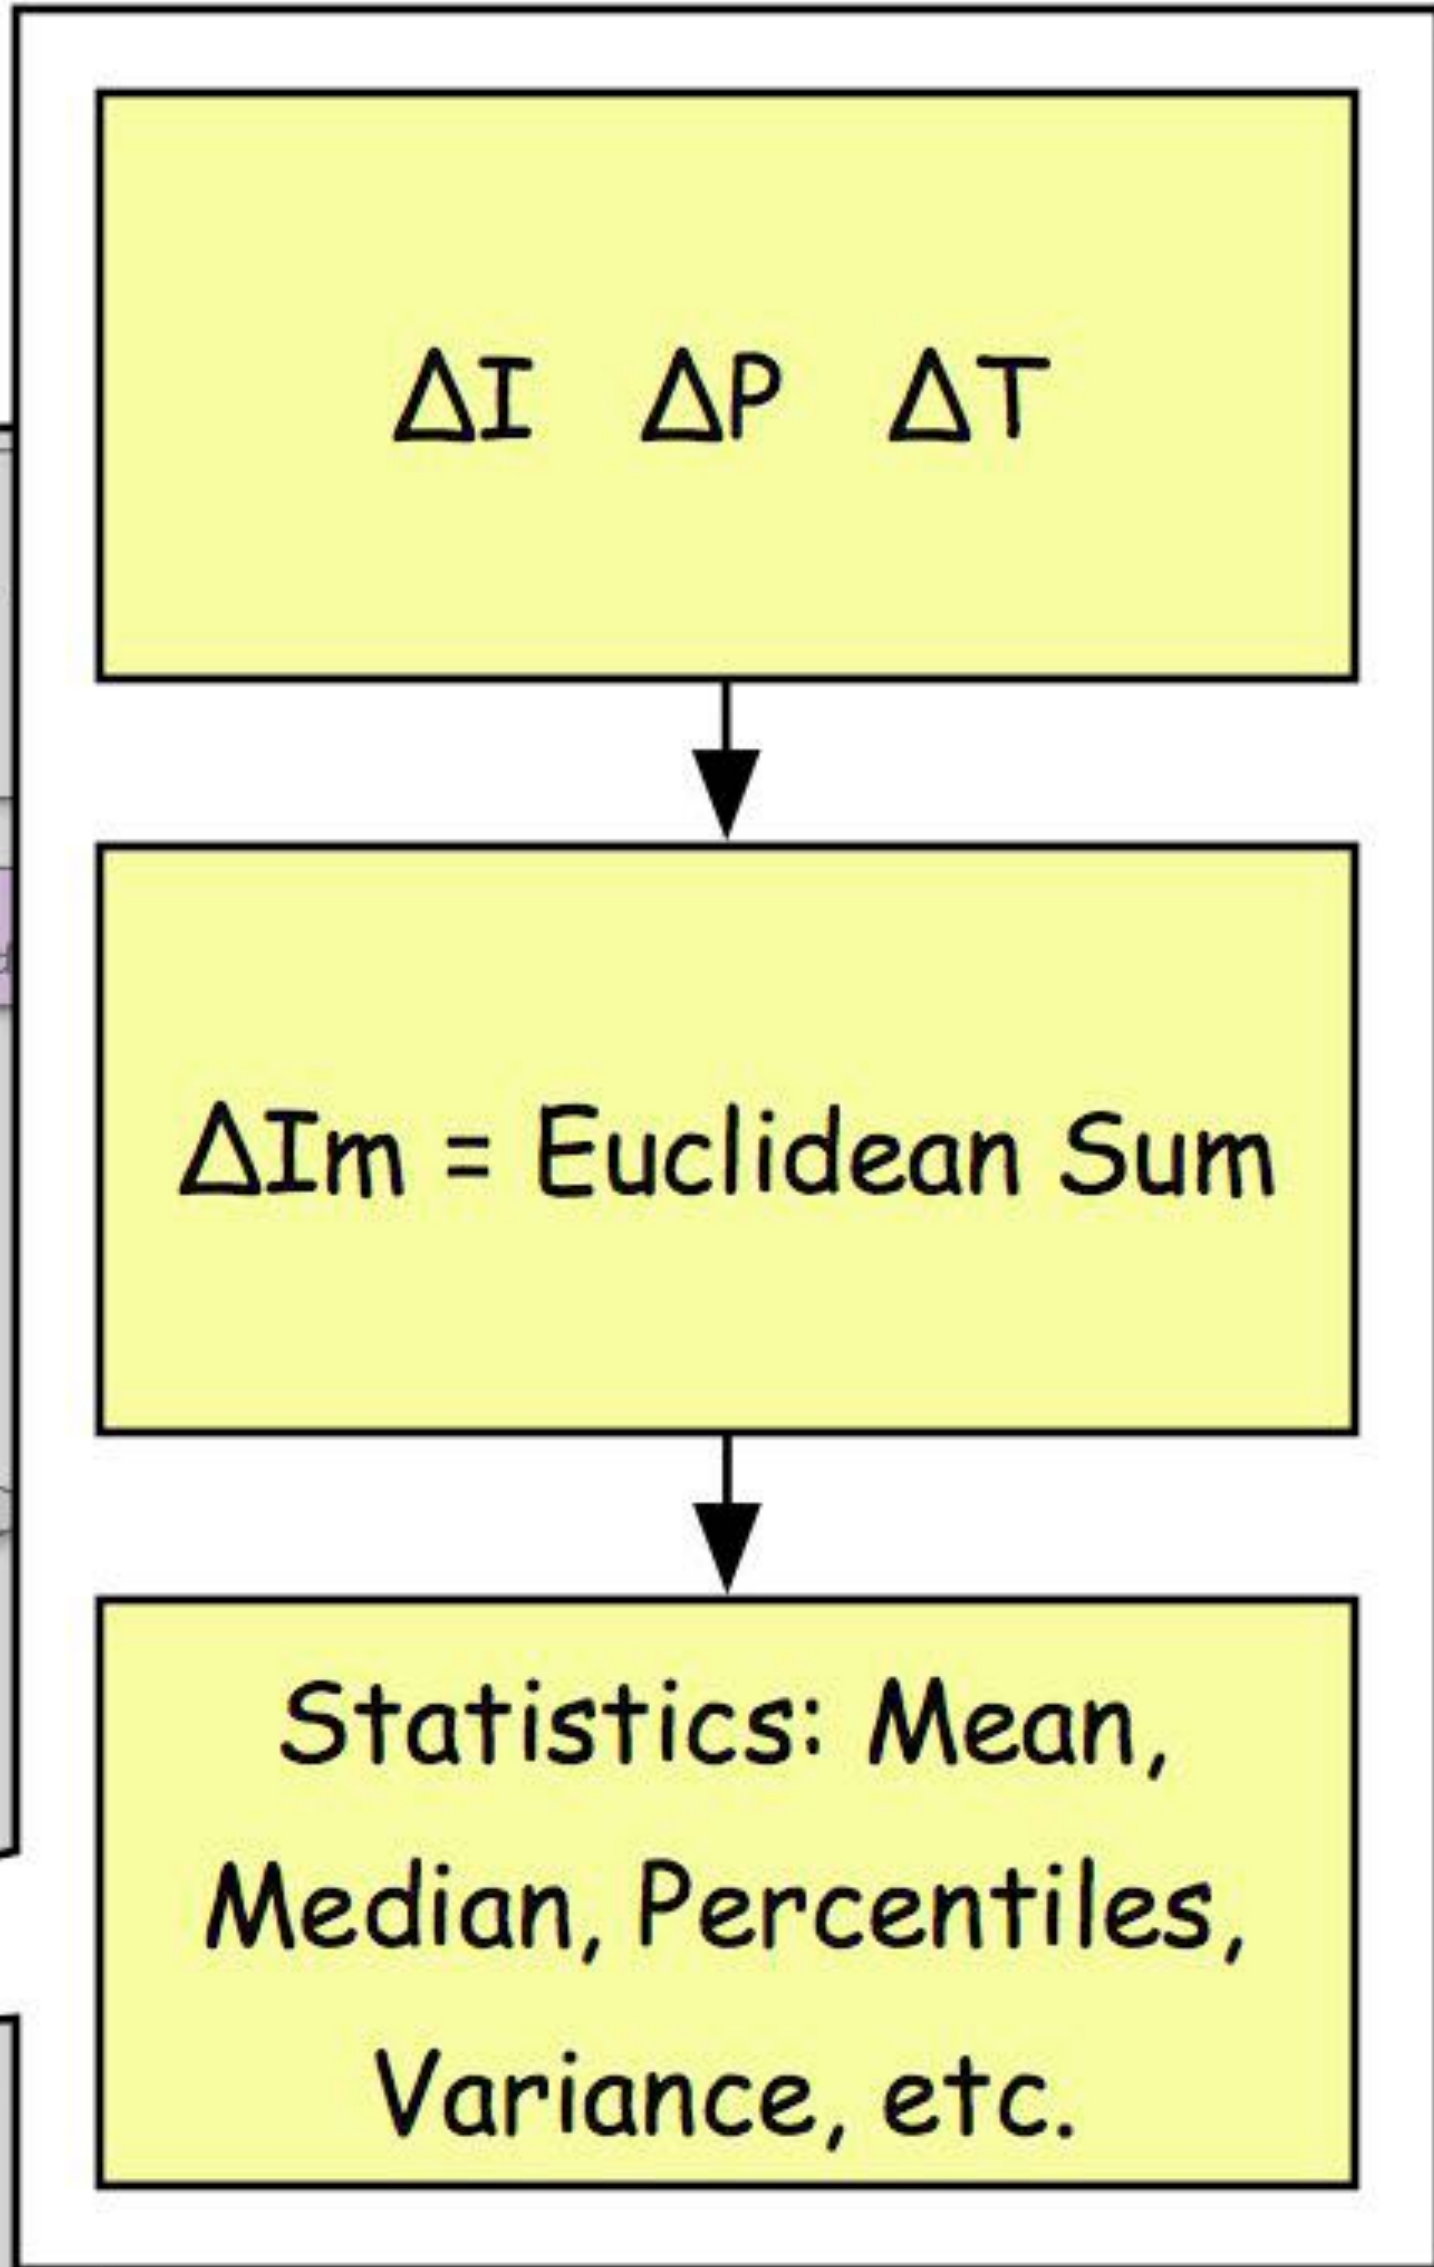
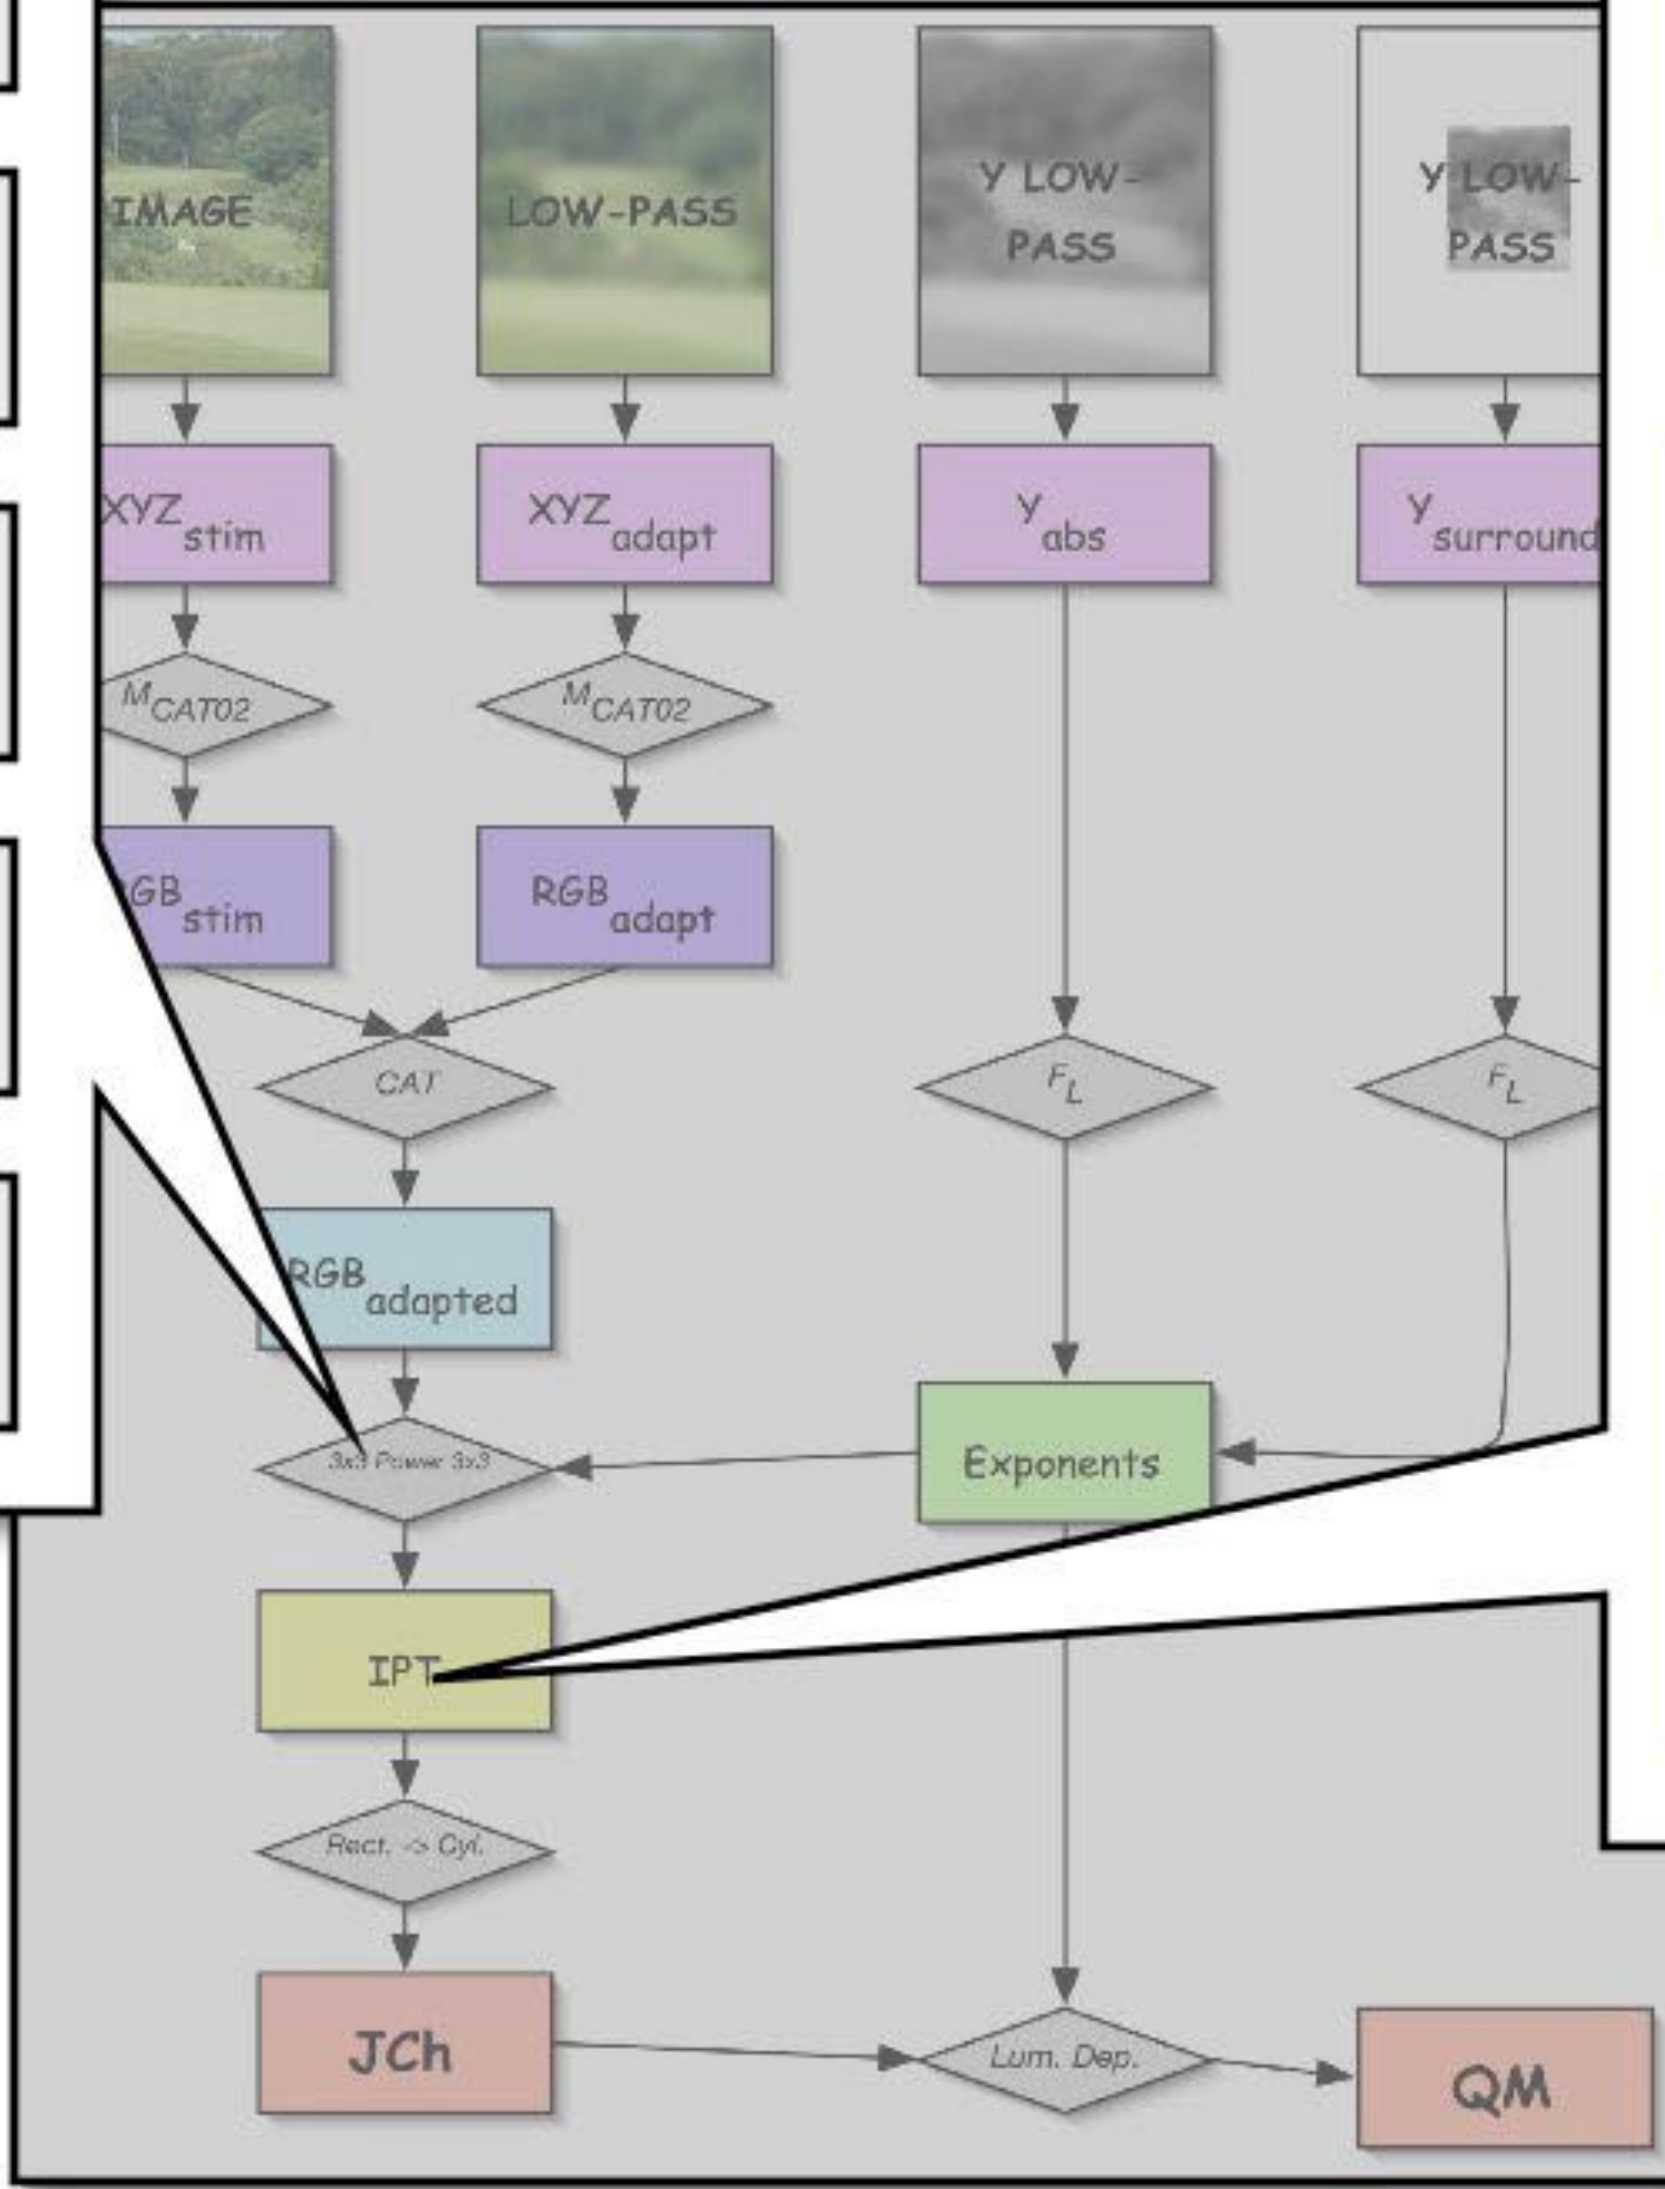
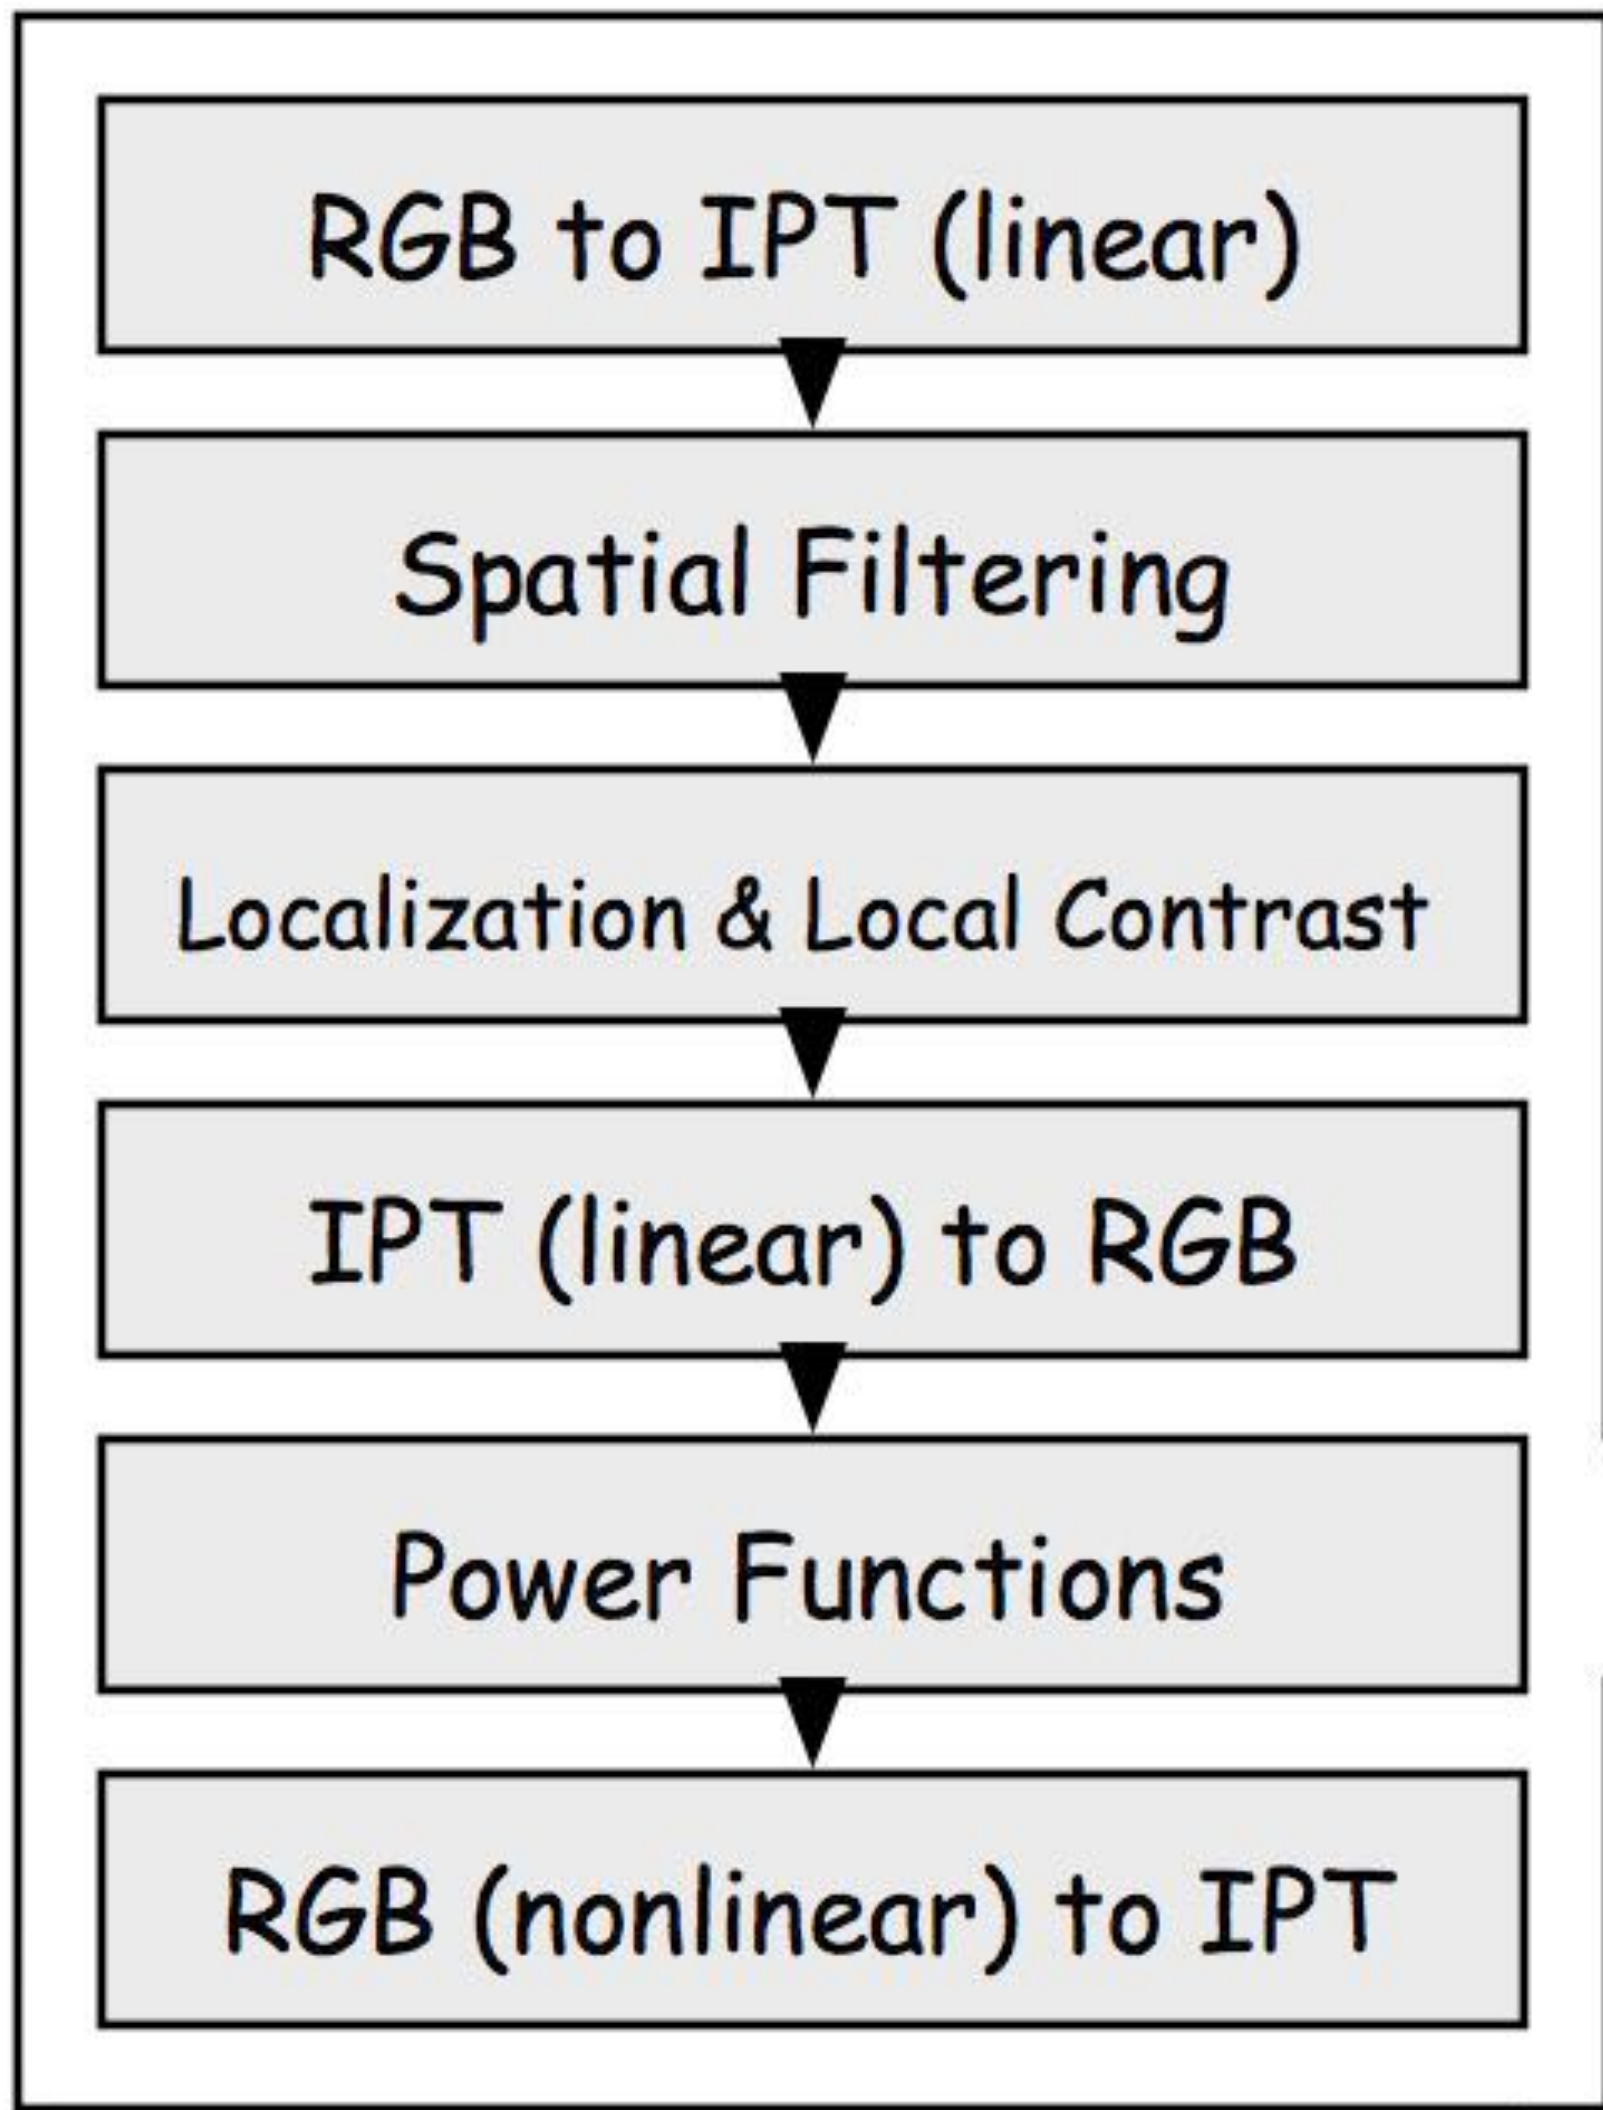
The opponent channels are filtered using Contrast
Sensitivity Functions, which are adapted based on the
spatial information in the image. The filtering
decreases information that is not visible, and
increases information where it is most visible.



Models of local attention and local contrast are
applied to the filtered images

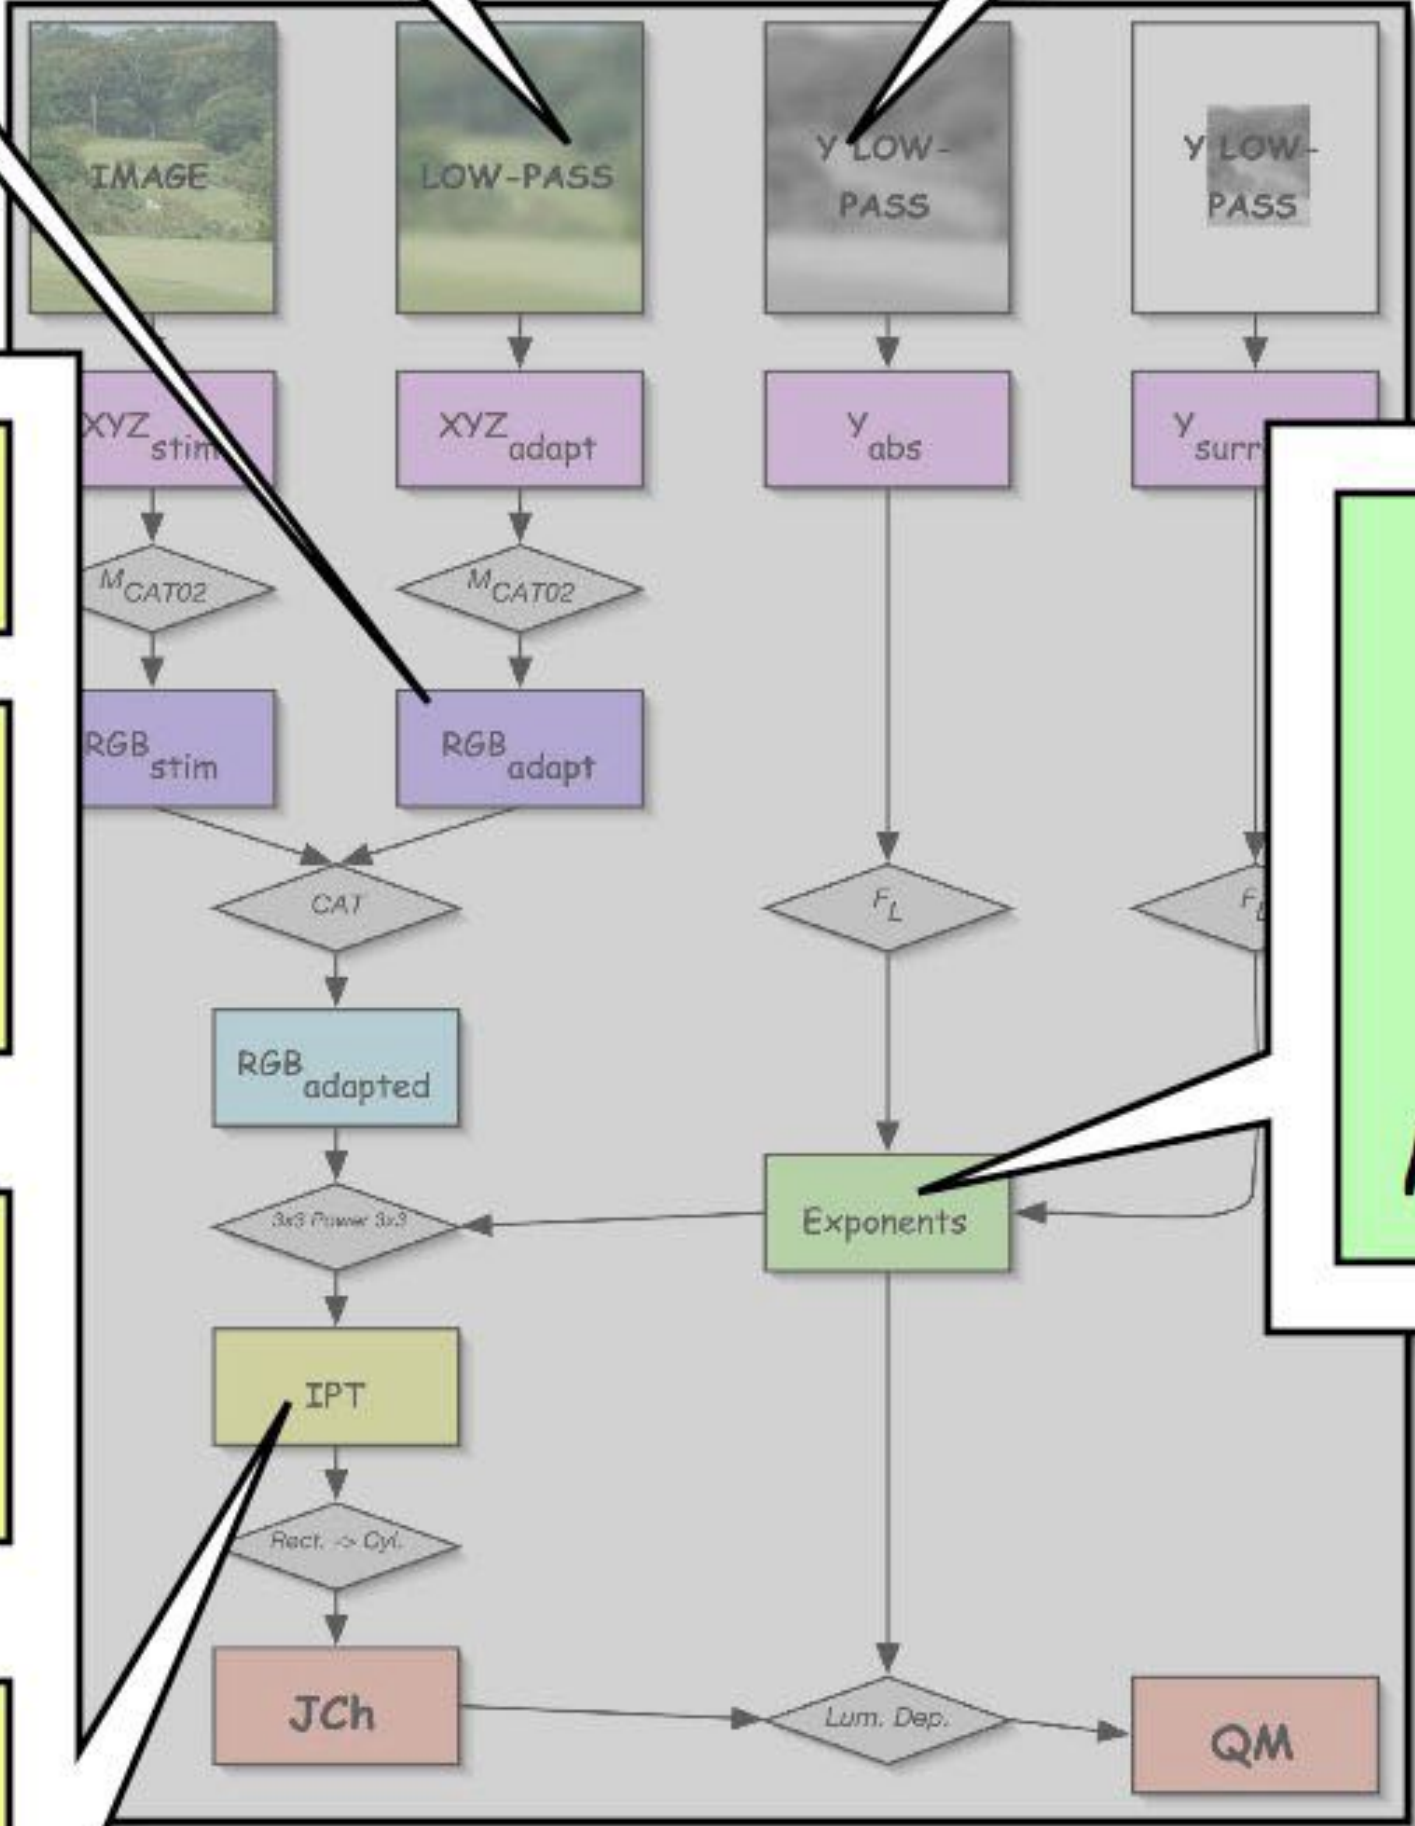
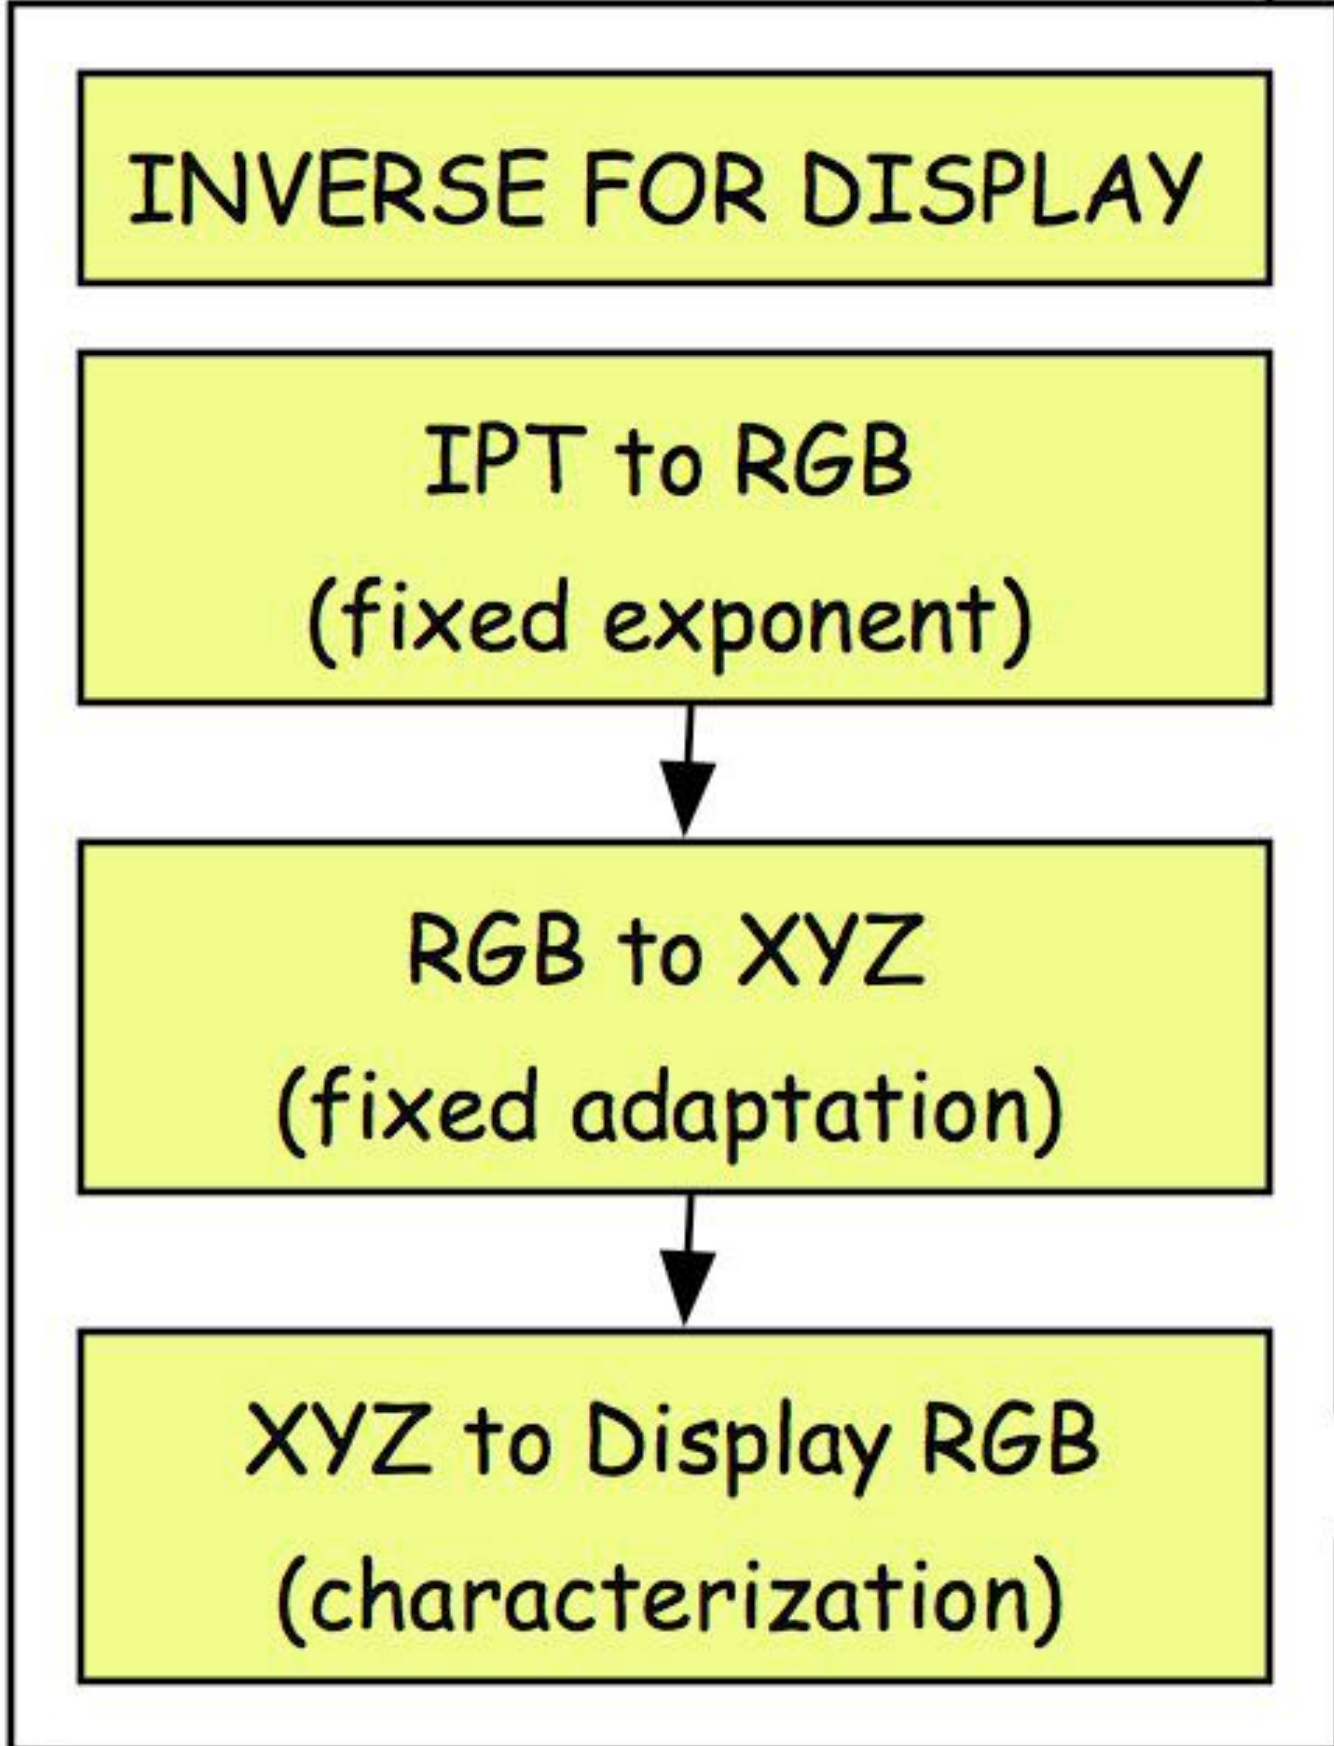


The filtered images are then converted into CIELAB
coordinates, and a Pixel-by-Pixel color difference is
calculated.




Luminance
Adaptation
Only

Gaussian Low-Pass
2-sigma = 1/4 Image Width
~4 Degrees



CIECAM02 F_L Function
Norm. to 1.0 at 1000 cd/m²
Clipped to Min. of 0.3
Multiplied by IPT Exp. (0.43)

(a)



(b)



Temporal Integrator
Based on Previous 10 Sec.
Psychophysically Derived

Gaussian Low-Pass
2-sigma = 1/4 Image Width
~4 Degrees

Luminance
Adaptation
Only

INVERSE FOR DISPLAY

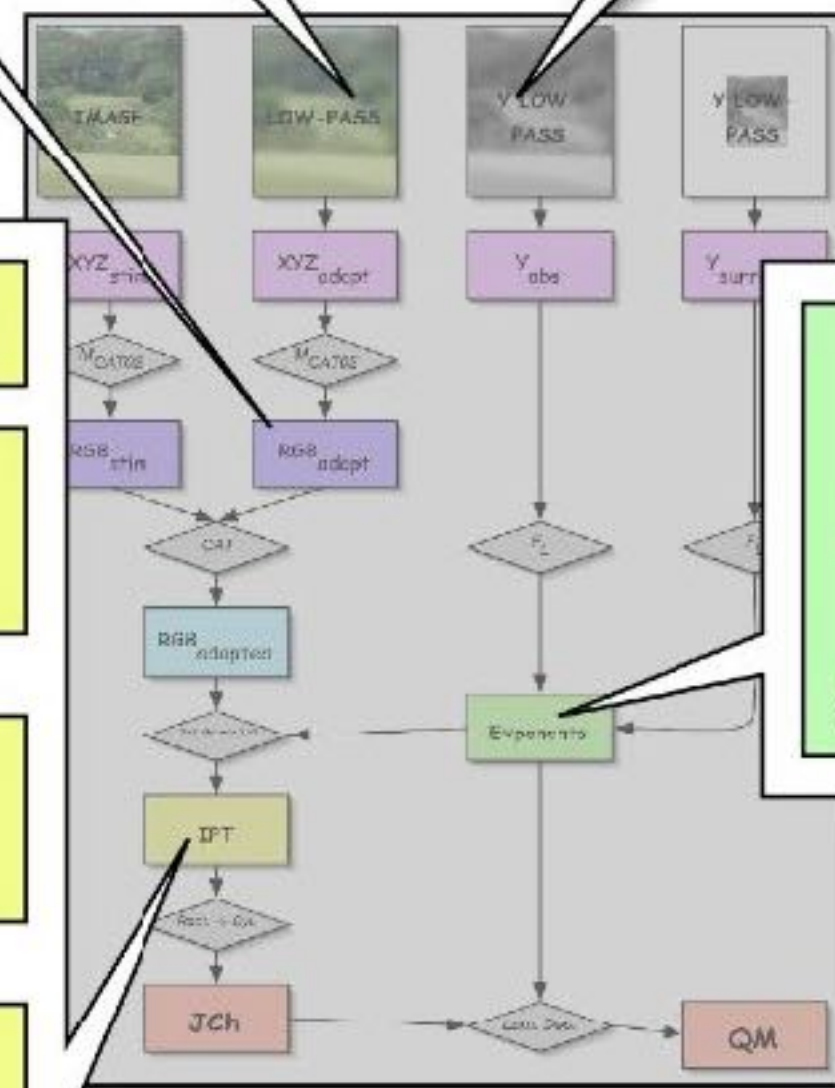
IPT to RGB
(fixed exponent)

RGB to XYZ
(fixed adaptation)

XYZ to Display RGB
(characterization)

CIECAM02 F_L Function
Norm. to 1.0 at 1000 cd/m^2
Clipped to Min. of 0.3
Multiplied by IPT Exp. (0.43)

Rendered
frame-by-
frame with
temporally
integrated
adaptation
stimulus.



HDR Photographic Survey

The HDR Photographic Survey

Mark D. Fairchild, Rochester Institute of Technology, Munsell Color Science Laboratory, Rochester, NY/USA

Abstract

High-dynamic-range (HDR) imaging is one of the remaining frontiers for significant advancement in consumer and professional color imaging systems. Systems for the capture, processing, and display of HDR images are gradually moving from the research labs and specialized applications to more mainstream usage and impact on consumer photography. The HDR Photographic Survey is a research project and public-domain database of images and scene data aimed at improving these systems by providing images and data that can be freely used by researchers around the world. This allows for more efficient testing and improvement of HDR algorithms and displays through enhanced inter-comparisons of results from various researchers and the availability of images with colorimetric and color appearance reference data from the original scenes. This paper describes the process of collecting the images and scene data, the system characterization, and the creation and use of the database.

Introduction

In the 1860's and 1870's outstanding and dedicated photographers such as Timothy O'Sullivan, William Henry Jackson, John K. Hillers, and others took part in expeditions to the American west

as the Zone System, involved pre-visualization of the desired print while at the original scene, careful exposure of the photographic negative as a relatively-HDR record of the scene, and then careful printing with dodging, burning, and retouching to create low-dynamic-range prints that matched the actual or imagined appearance of the original scene. So it could be said that he recorded the appearance of the scene with local adaptation and then rendered that appearance in the darkroom. This is exactly what current researchers try to automate in some way with visual models of spatial adaptation aimed at developing HDR rendering algorithms.[5] In fact, Adams' techniques have rather directly inspired one very successful HDR rendering algorithm.[6]





Figure 1. Eighteen individual exposures, each separated by one stop, used to construct the Luxo Double Checker HDR characterization image.

Inspiration

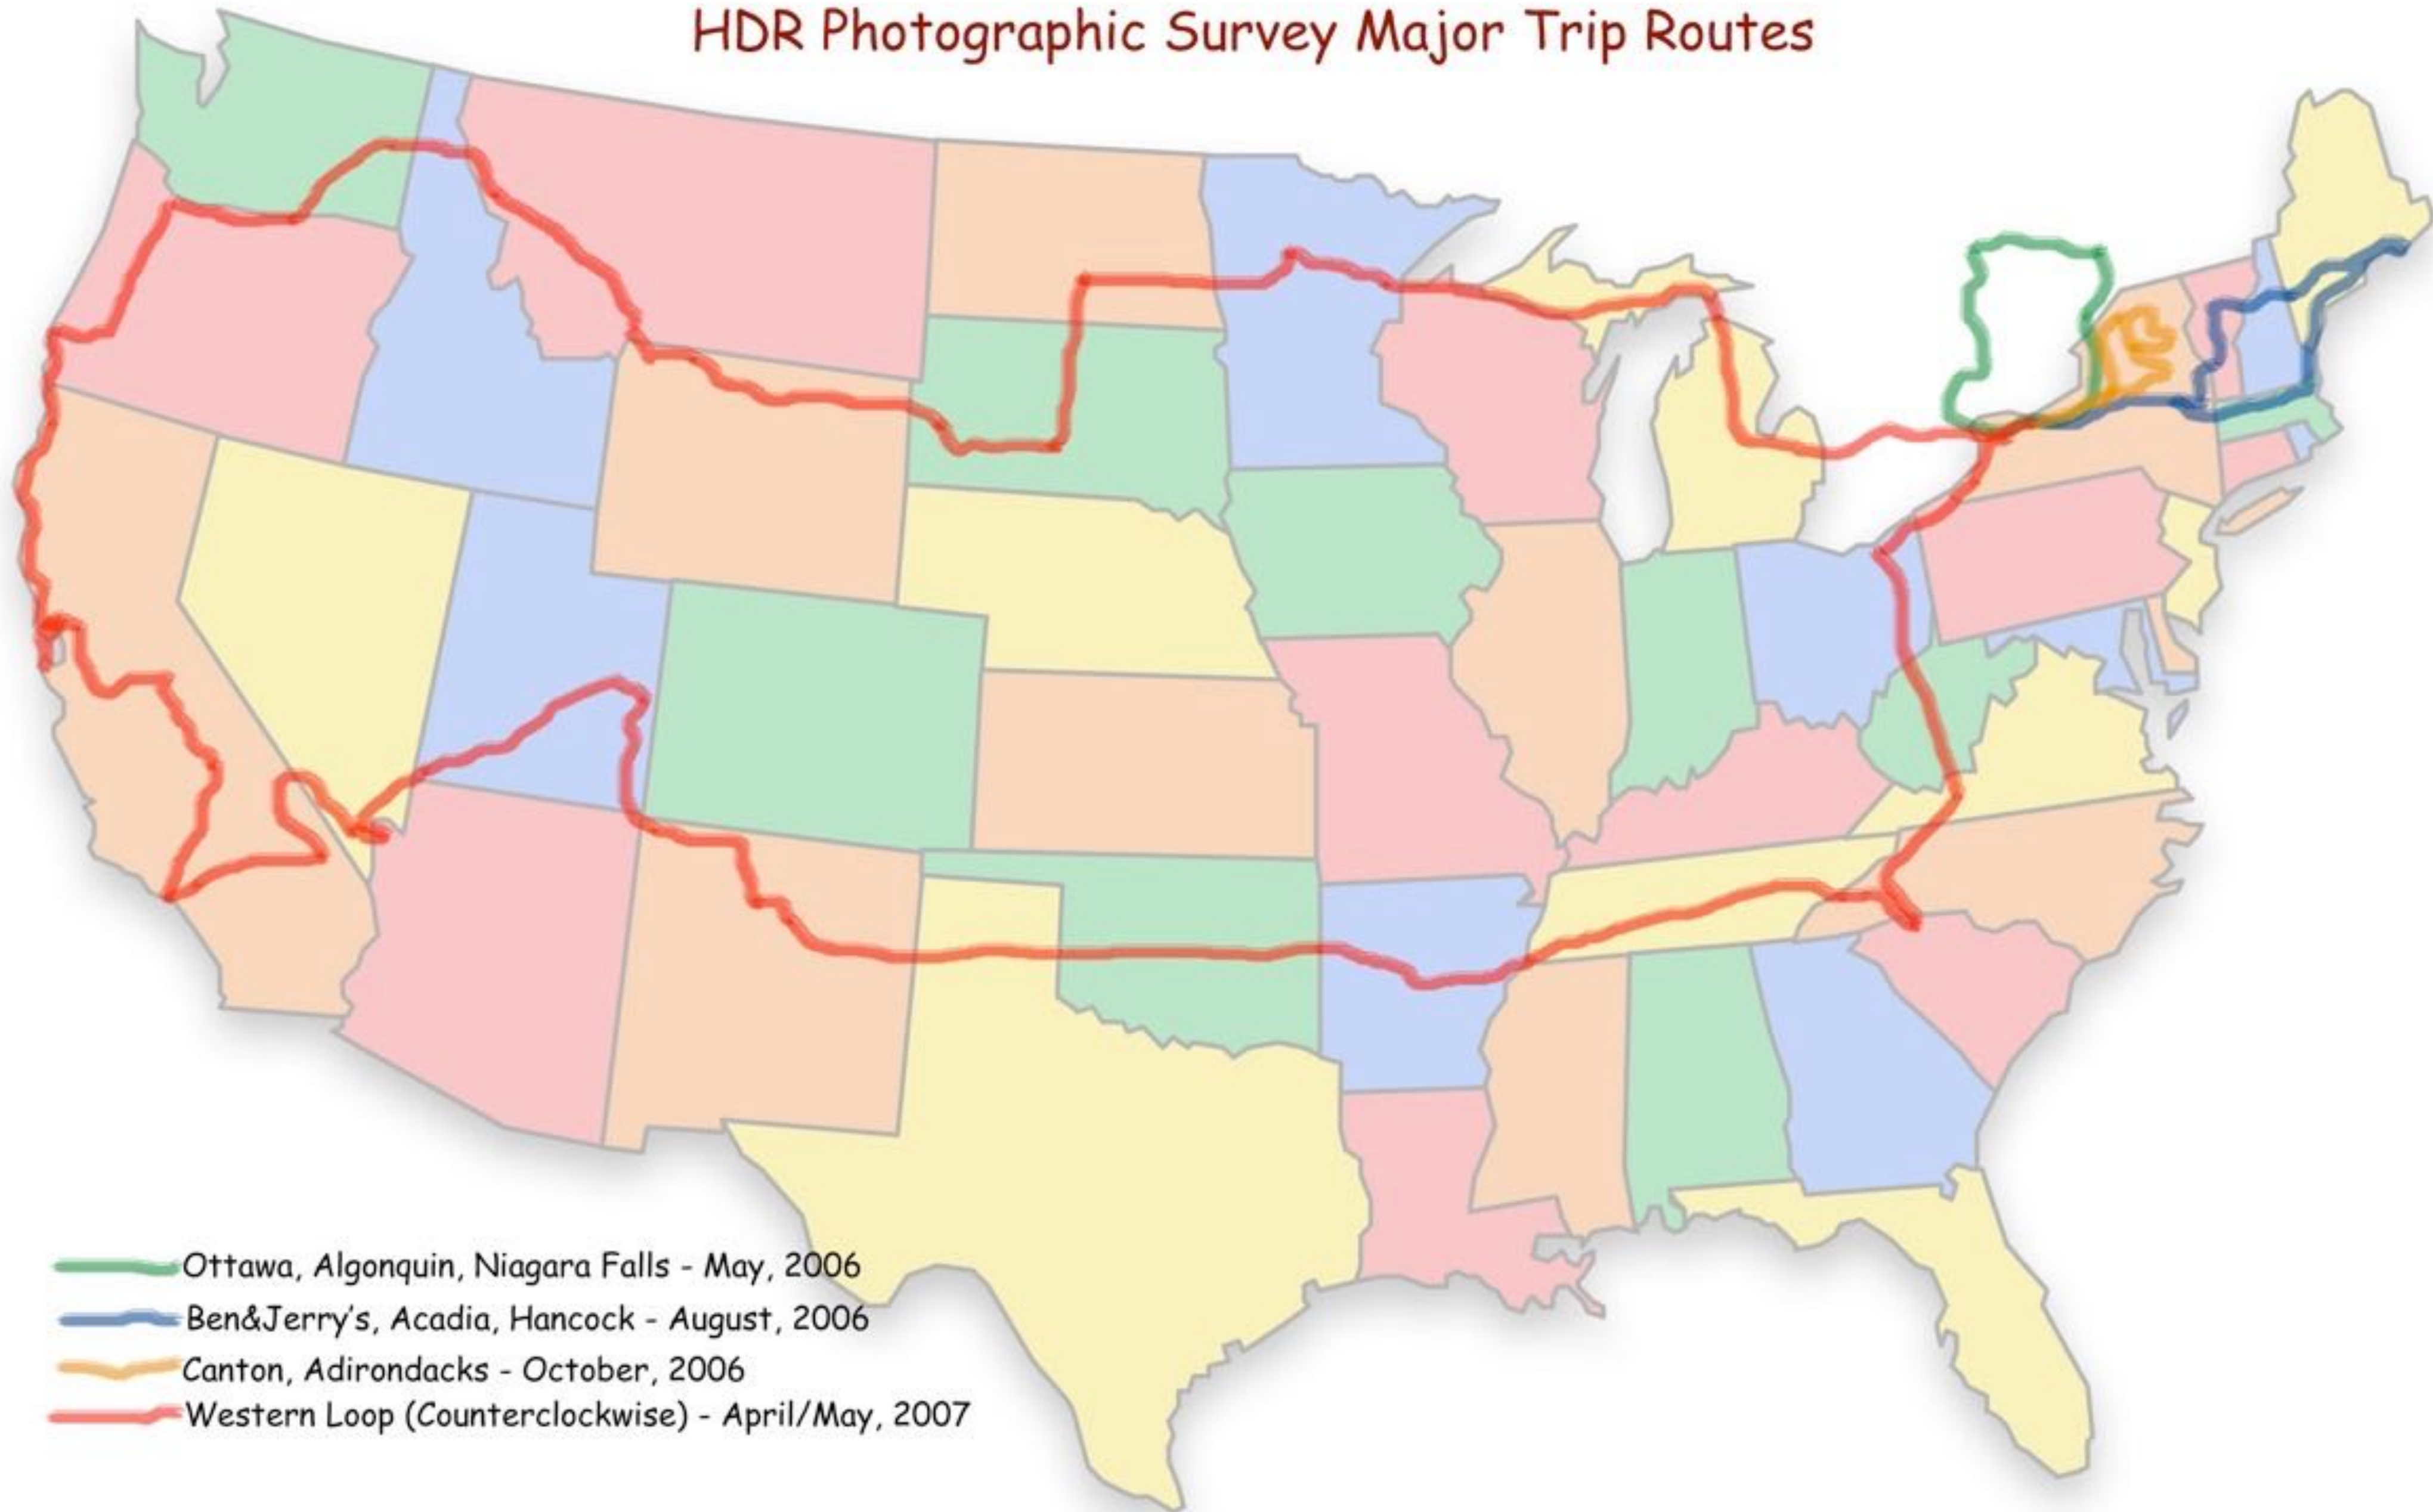
19th Century Photographic Surveys of the American West



William Henry Jackson, Mt. of the Holy Cross



HDR Photographic Survey Major Trip Routes



Procedures

- > Calibrated/Characterized Nikon D2x
- > 9 Exposures, 1-Stop Increments (typical)
- > Assembly into Linear HDR Image (CS2)
- > OpenEXR, 32-Bit Floating Point
- > Scene Colorimetry, Appearance, Other Data
- > Mirror Ball Image

- > All Details Online

Measurement Map



Camera Characterization

- > Nikon D2x
- > Linear HDR Image Data
- > 12-Bit RAW --> 16-Bit TIFF --> 32-Bit OpenEXR
- > 3x3 Matrix from Linear RGB to CIE XYZ
- > Verified in a Second Research Project

- > Mean $\Delta E_{ab} = 2.5$ (Max = 5.5)

Example Data Sheet

Paul Bunyan Data.xls

| | A | B | C | D | E | F | G | H | I | J |
|----|---|---|--------------------|--------------|------------|---------------|----------------------|--------------------|---------|---------------|
| 2 | | Scene Name: | Paul Bunyan | | | Date: | 4/19/07 | Time: | 5:11PM | |
| 3 | | Location: | Bemidji, MN | Latitude: | 47°28.239N | Longitude: | 94°52.745W | Direction: | 149° | |
| 4 | | Weather: | 65°+ Sunny, 28.59" | | | | | Elevation: | 1352' | |
| 5 | | | | | | | | | | |
| 6 | | Camera: | D2x | F/#: | 16 | Focal Length: | 20mm | Num. of Exposures: | 9 +/- 4 | |
| 7 | | Colorimeter: | CS100 | Meas. Angle: | 1° | | | | | |
| 8 | | | | | | | | | | |
| 9 | | Scene Elements: | | | | | | | | |
| 10 | | NAME | Y | x | y | J | C | h | Q | Texture (Y/N) |
| 11 | | Babe Under Eye | 4780 | 0.221 | 0.244 | 8 | 8 | B | 200 | Y |
| 12 | | Babe Leg in Shadow | 522 | 0.237 | 0.249 | 5 | 6 | B | 100 | Y |
| 13 | | Paul Right Thigh | 875 | 0.208 | 0.228 | 4 | 7 | B10G | 80 | Y |
| 14 | | Paul Red Center | 1760 | 0.533 | 0.324 | 7.5 | 8 | Y90R | 180 | Y |
| 15 | | Sky Middle | 2560 | 0.243 | 0.252 | 8 | 6 | B | 100 | N |
| 16 | | Trees Middle | 936 | 0.379 | 0.408 | 5 | 5 | G10Y | 120 | Y |
| 17 | | | | | | | | | | |
| 18 | | | | | | | | | | |
| 19 | | | | | | | | | | |
| 20 | | | | | | | | | | |
| 21 | | | | | | | | | | |
| 22 | | | | | | | | | | |
| 23 | | | | | | | | | | |
| 24 | | | | | | | | | | |
| 25 | | Mirror Ball Shot: | Y | N | | | | | | |
| 26 | | | | | | | | | | |
| 27 | | Summary Stats (Post): | Min. | Max. | Dyn. Range | Mean | Luminance Multiplier | | | |
| 28 | | | 0.01 | 2,690 | 1.2K:1 | 333 | 300 | | | |
| 29 | | | | | | | | | | |
| 30 | | Added Notes: | | | | | | | | |
| 31 | | Got my BSU shirt first!!! | | | | | | | | |
| 32 | | Sun Fading Fast ... | | | | | | | | |
| 33 | | Waited to finish measurements in the sun. | | | | | | | | |
| 34 | | | | | | | | | | |
| 35 | | | | | | | | | | |

Ready
Sum=0
2m=0

Sheet1 Sheet2 Sheet3

SCRL CAPS
2CBT CV62

9-Stop Mosaic



Linear OpenEXR



Locally Rendered



9-Stop Mosaic



Linear OpenEXR



Locally Rendered




Comparison



Linear



CS2 Local



iCAM06

Comparison 2



Linear



CS2 Local



iCAM06



Website

`<markfairchild.org/HDR.html>`

iCAM06



ELSEVIER

Available online at www.sciencedirect.com



J. Vis. Commun. Image R. 18 (2007) 406–414

JOURNAL OF
VISUAL
Communication &
IMAGE
Representation

www.elsevier.com/locate/jvci

iCAM06: A refined image appearance model for HDR image rendering

Jiangtao Kuang ^{*}, Garrett M. Johnson, Mark D. Fairchild

Munsell Color Science Laboratory, Rochester Institute of Technology, 54 Lomb Memorial Dr., Rochester, NY 14623, USA

Received 15 November 2006; accepted 12 June 2007

Available online 27 June 2007

Abstract

A new image appearance model, designated iCAM06, was developed for High-Dynamic-Range (HDR) image rendering. The model, based on the iCAM framework, incorporates the spatial processing models in the human visual system for contrast enhancement, photoreceptor light adaptation functions that enhance local details in highlights and shadows, and functions that predict a wide range of color appearance phenomena. Evaluation of the model proved iCAM06 to have consistently good HDR rendering performance in both preference and accuracy making iCAM06 a good candidate for a general-purpose tone-mapping operator with further potential applications to a wide-range of image appearance research and practice.

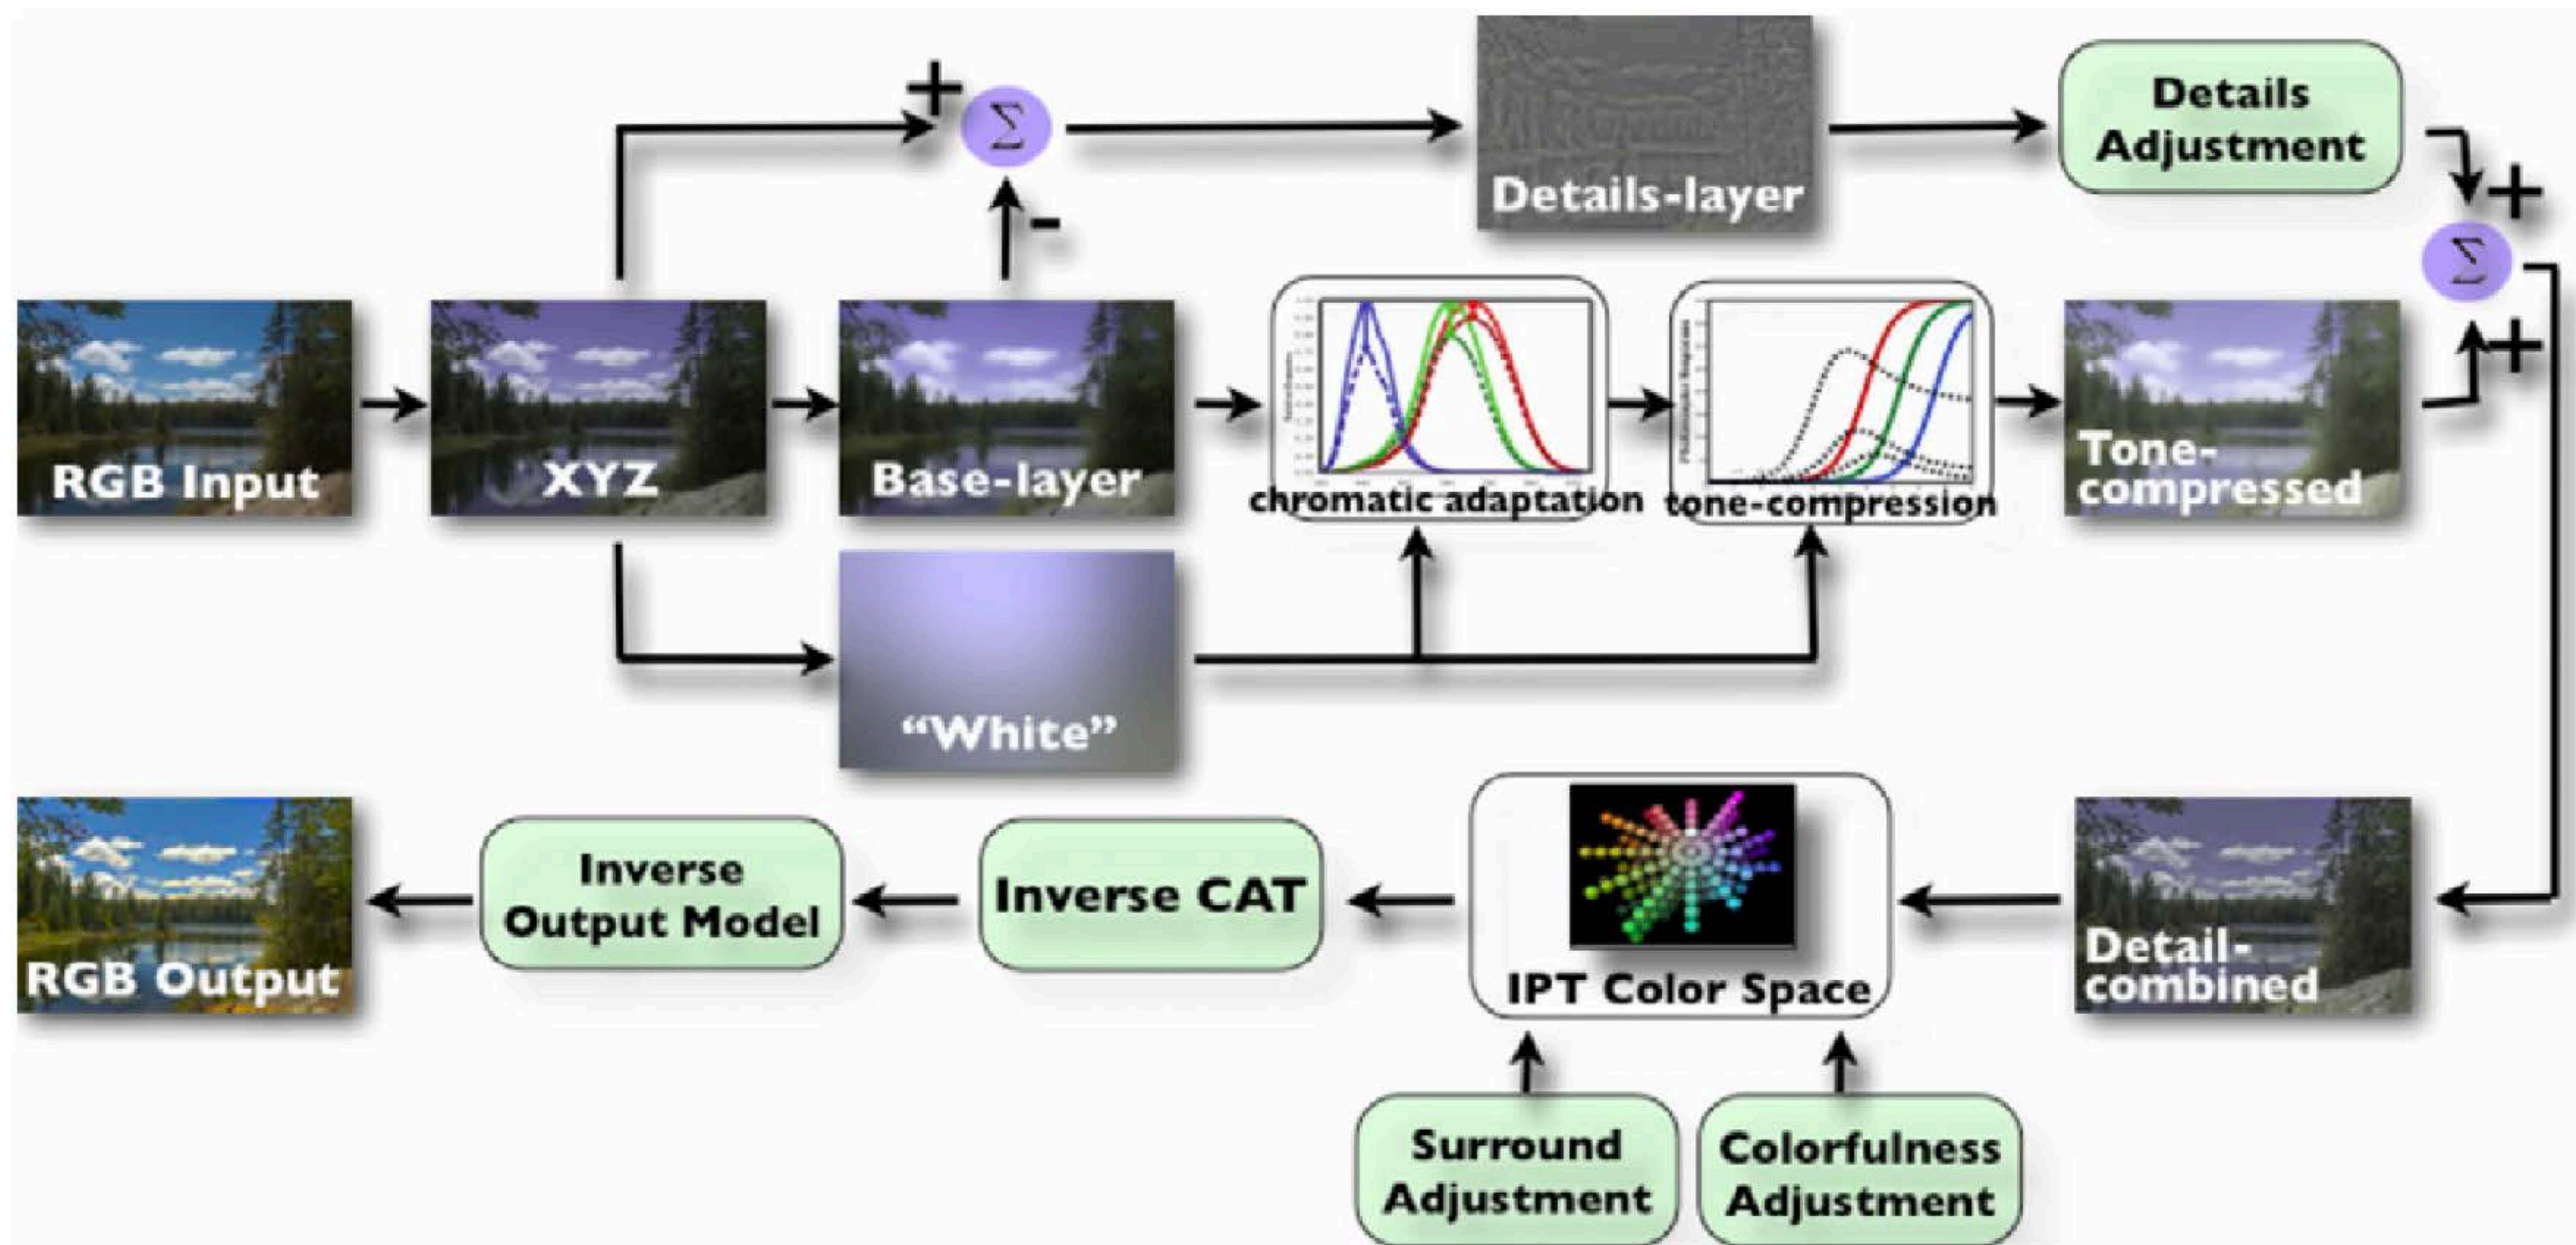


Fig. 1. Flowchart of iCAM06.




Fig. 2. Cone response after adaptation plotted against log intensity ($\log \text{cd/m}^2$) for three adaptation levels in iCAM06. Open circles: reference whites; filled circles: adapting luminances.




Fig. 3. Response functions for the rods in iCAM06. The rod signal, A_s , is plotted against log luminances for three adaptation levels -1 , 0 and 1 , respectively.




Fig. 4. Exponent function for details adjustment accounting for the Stevens effect.





Fig. 5. Multiplication functions for colorfulness adjustment accounting for Hunt effect.




Fig. 7. Overall preference scores of tone-mapping operators over 12 HDR images (The operators are labeled as Histogram Adjustment (H), iCAM, Photographic Reproduction (P), Photoshop Exposure and Gamma (PG), Photoshop Local Adaptation (PL), Bilateral Filter (B) and iCAM06. The same labels are used in this paper.).




Fig. 6. Thumbnails of experimental images (scenes).




Fig. 10. Overall accuracy scores of tone-mapping operators for 4 real-world scenes.




Fig. 8. Preference scores for 12 test images by image.

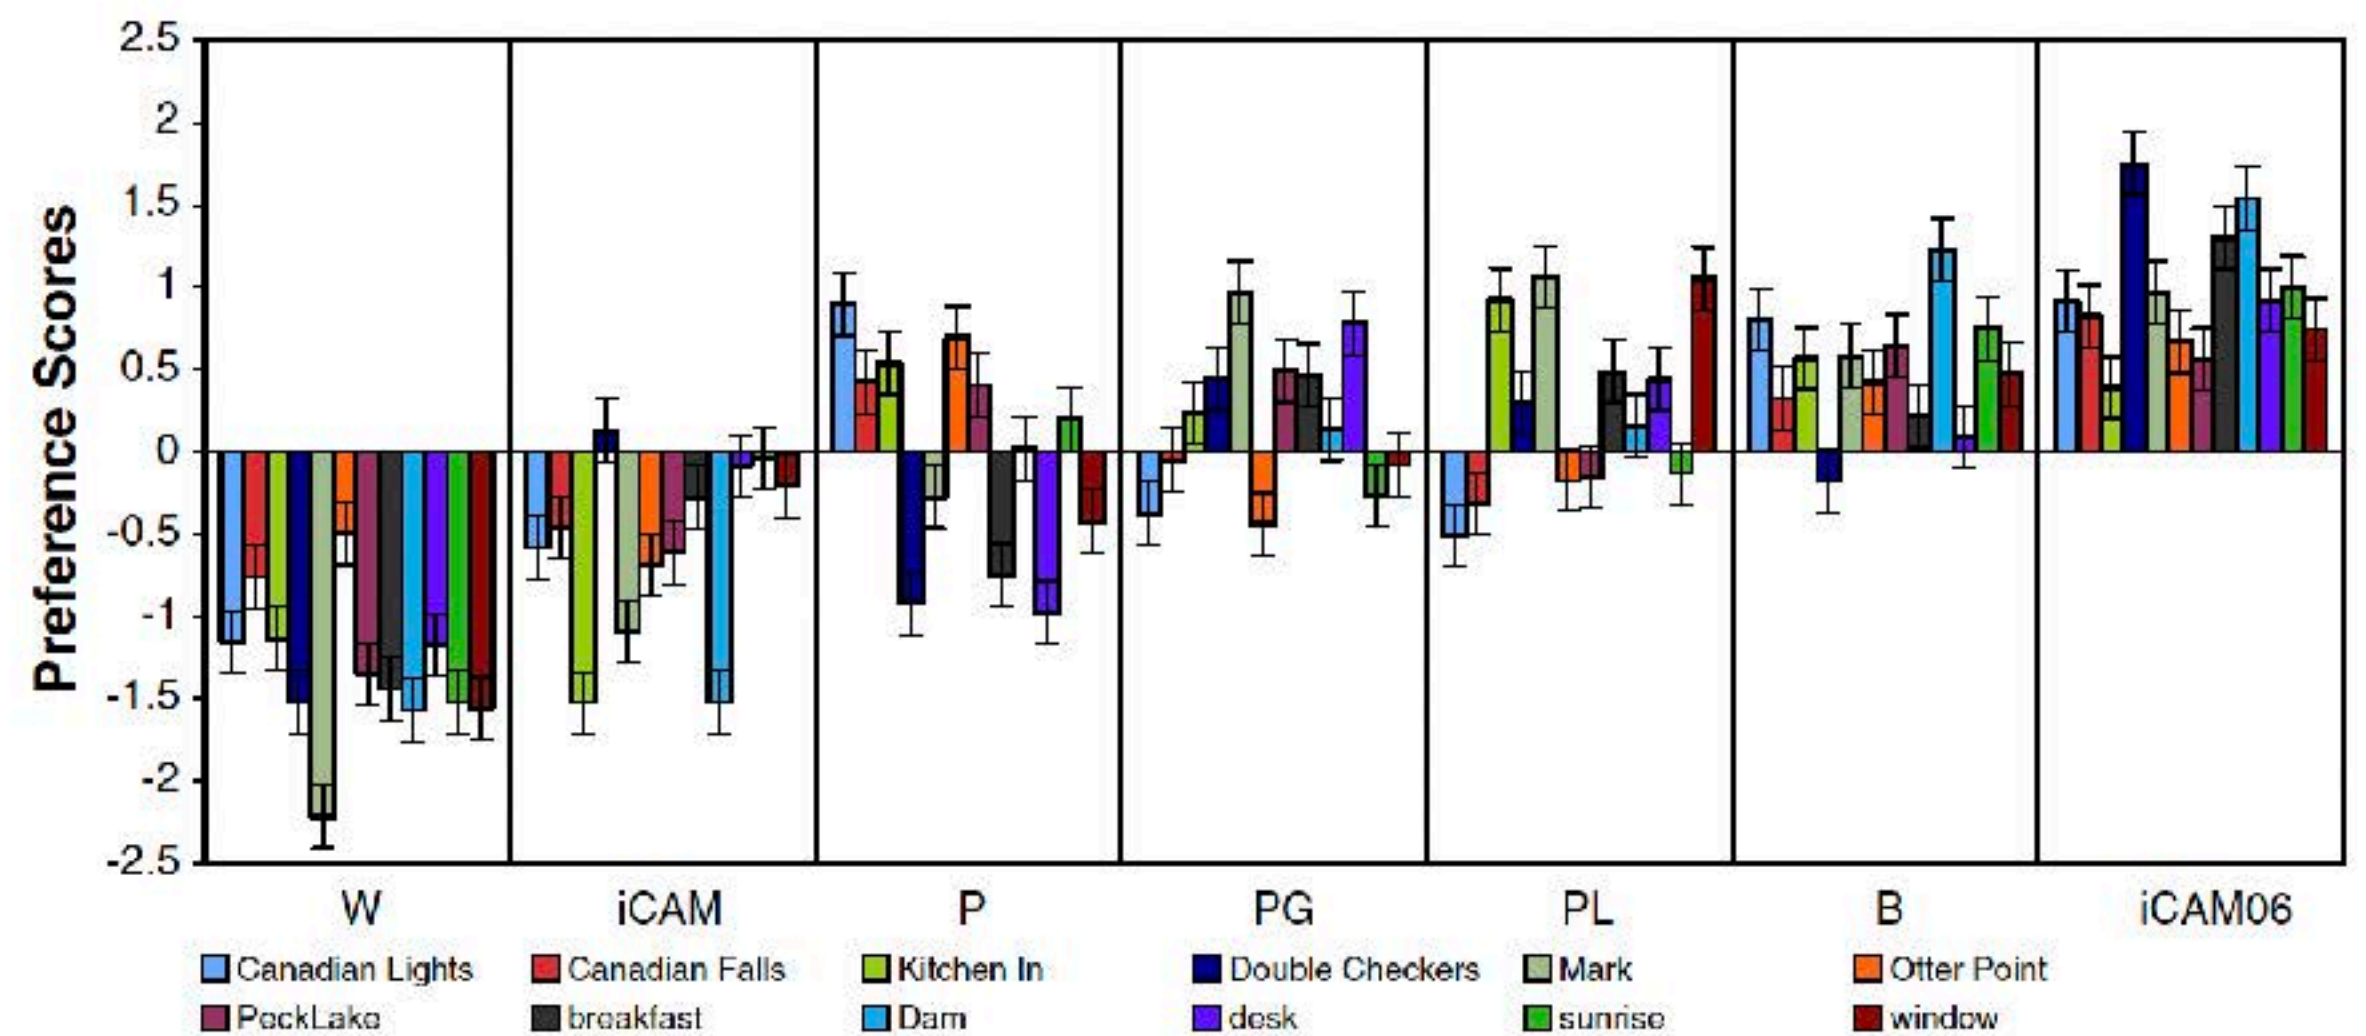


Fig. 9. Preference scores for 12 test images by algorithm.

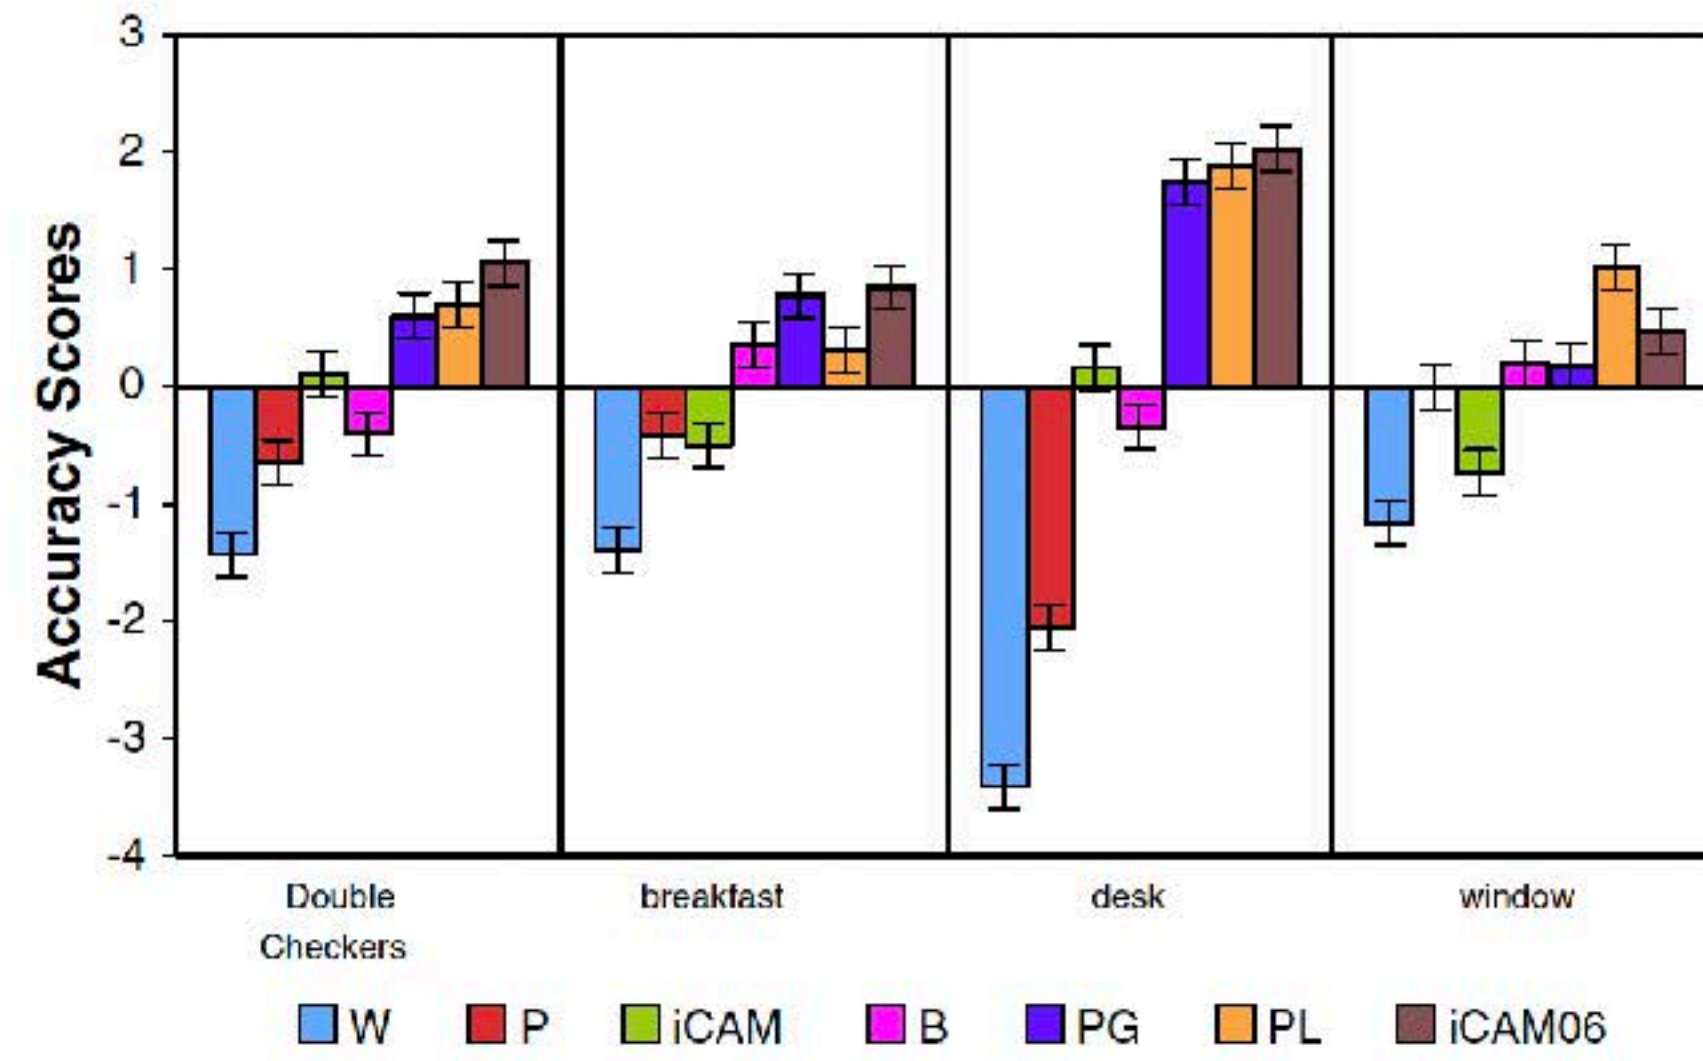


Fig. 11. Accuracy scores for 4 test HDR scenes by scene.




Fig. 12. Accuracy scores for 4 test HDR scenes by algorithm.

Large Color Differences

RESEARCH ARTICLE

WILEY

Distance metrics for very large color differences

Saeedeh Abasi^{1,2}  | Mohammad Amani Tehran¹ | Mark D. Fairchild²

¹Textile Engineering Department,
Amirkabir University of Tehran,
Tehran, Iran

²Program of Color Science/Munsell Color
Science Laboratory, Rochester Institute of
Technology, Rochester, New York

Correspondence

Mohammad Amani Tehran, Textile
Engineering Department, Amirkabir
University of Tehran, Tehran, Iran.
Email: amani@aut.ac.ir

Abstract

Small, supra-threshold color differences are typically described with Euclidean distance metrics, or dimension-weighted Euclidean metrics, in color appearance spaces such as CIELAB. This research examines the perception and modeling of very large color differences in the order of 10 CIELAB units or larger, with an aim of describing the salience of color differences between distinct objects in real-world scenes and images. A psychophysical experiment was completed to compare directly large color-difference pairs designed to probe various Euclidean and non-Euclidean distance metrics. The results indicate that very large color differences are best described by HyAB, a combination of a Euclidean metric in hue and chroma with a city-block metric to incorporate lightness differences.

KEYWORDS

TABLE 1 One stimulus with two color pairs in first category

| | Pair number | L* | a* | b* | CIELAB | CIEDE2000 | Color presentation |
|------------|--------------------|-----------|-----------|-----------|---------------|------------------|---|
| Stimulus A | Pair 1 | 60 | -15 | 6.5 | 39.28 | 28.35 |  |
| | | 60 | 11.5 | -22.5 | | |  |
| | Pair 2 | 60 | -15 | 6.5 | 40.80 | 29.74 |  |
| | | 71 | 11.5 | -22.5 | | |  |

TABLE 3 One stimulus with two color pairs in first category



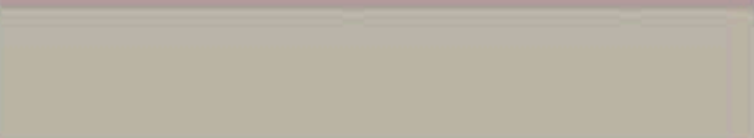

| | Pair number | L* | a* | b* | CIELAB | CIEDE2000 | Color presentation |
|------------|--------------------|-----------|-----------|-----------|---------------|------------------|---|
| Stimulus B | Pair 1 | 60 | -2.25 | 10.5 | 31.70 | 28.97 |  |
| | | 60 | 21.5 | -10.5 | | |  |
| | Pair 2 | 72 | -2.25 | 10.5 | 33.90 | 30.57 |  |
| | | 60 | 21.5 | -10.5 | | |  |

TABLE 4 One stimulus with two color pairs in second category

| | Pair number | L* | a* | b* | CIELAB | CIEDE2000 | Color presentation |
|------------|--------------------|-----------|-----------|-----------|---------------|------------------|---|
| Stimulus A | Pair 1 | 70 | -17.5 | 26.5 | 41.18 | 29.46 |  |
| | | 50 | -17.5 | -9.5 | | |  |
| | Pair 2 | 70 | -17.5 | 26.5 | 42.76 | 31.46 |  |
| | | 50 | -6 | -9.5 | | |  |

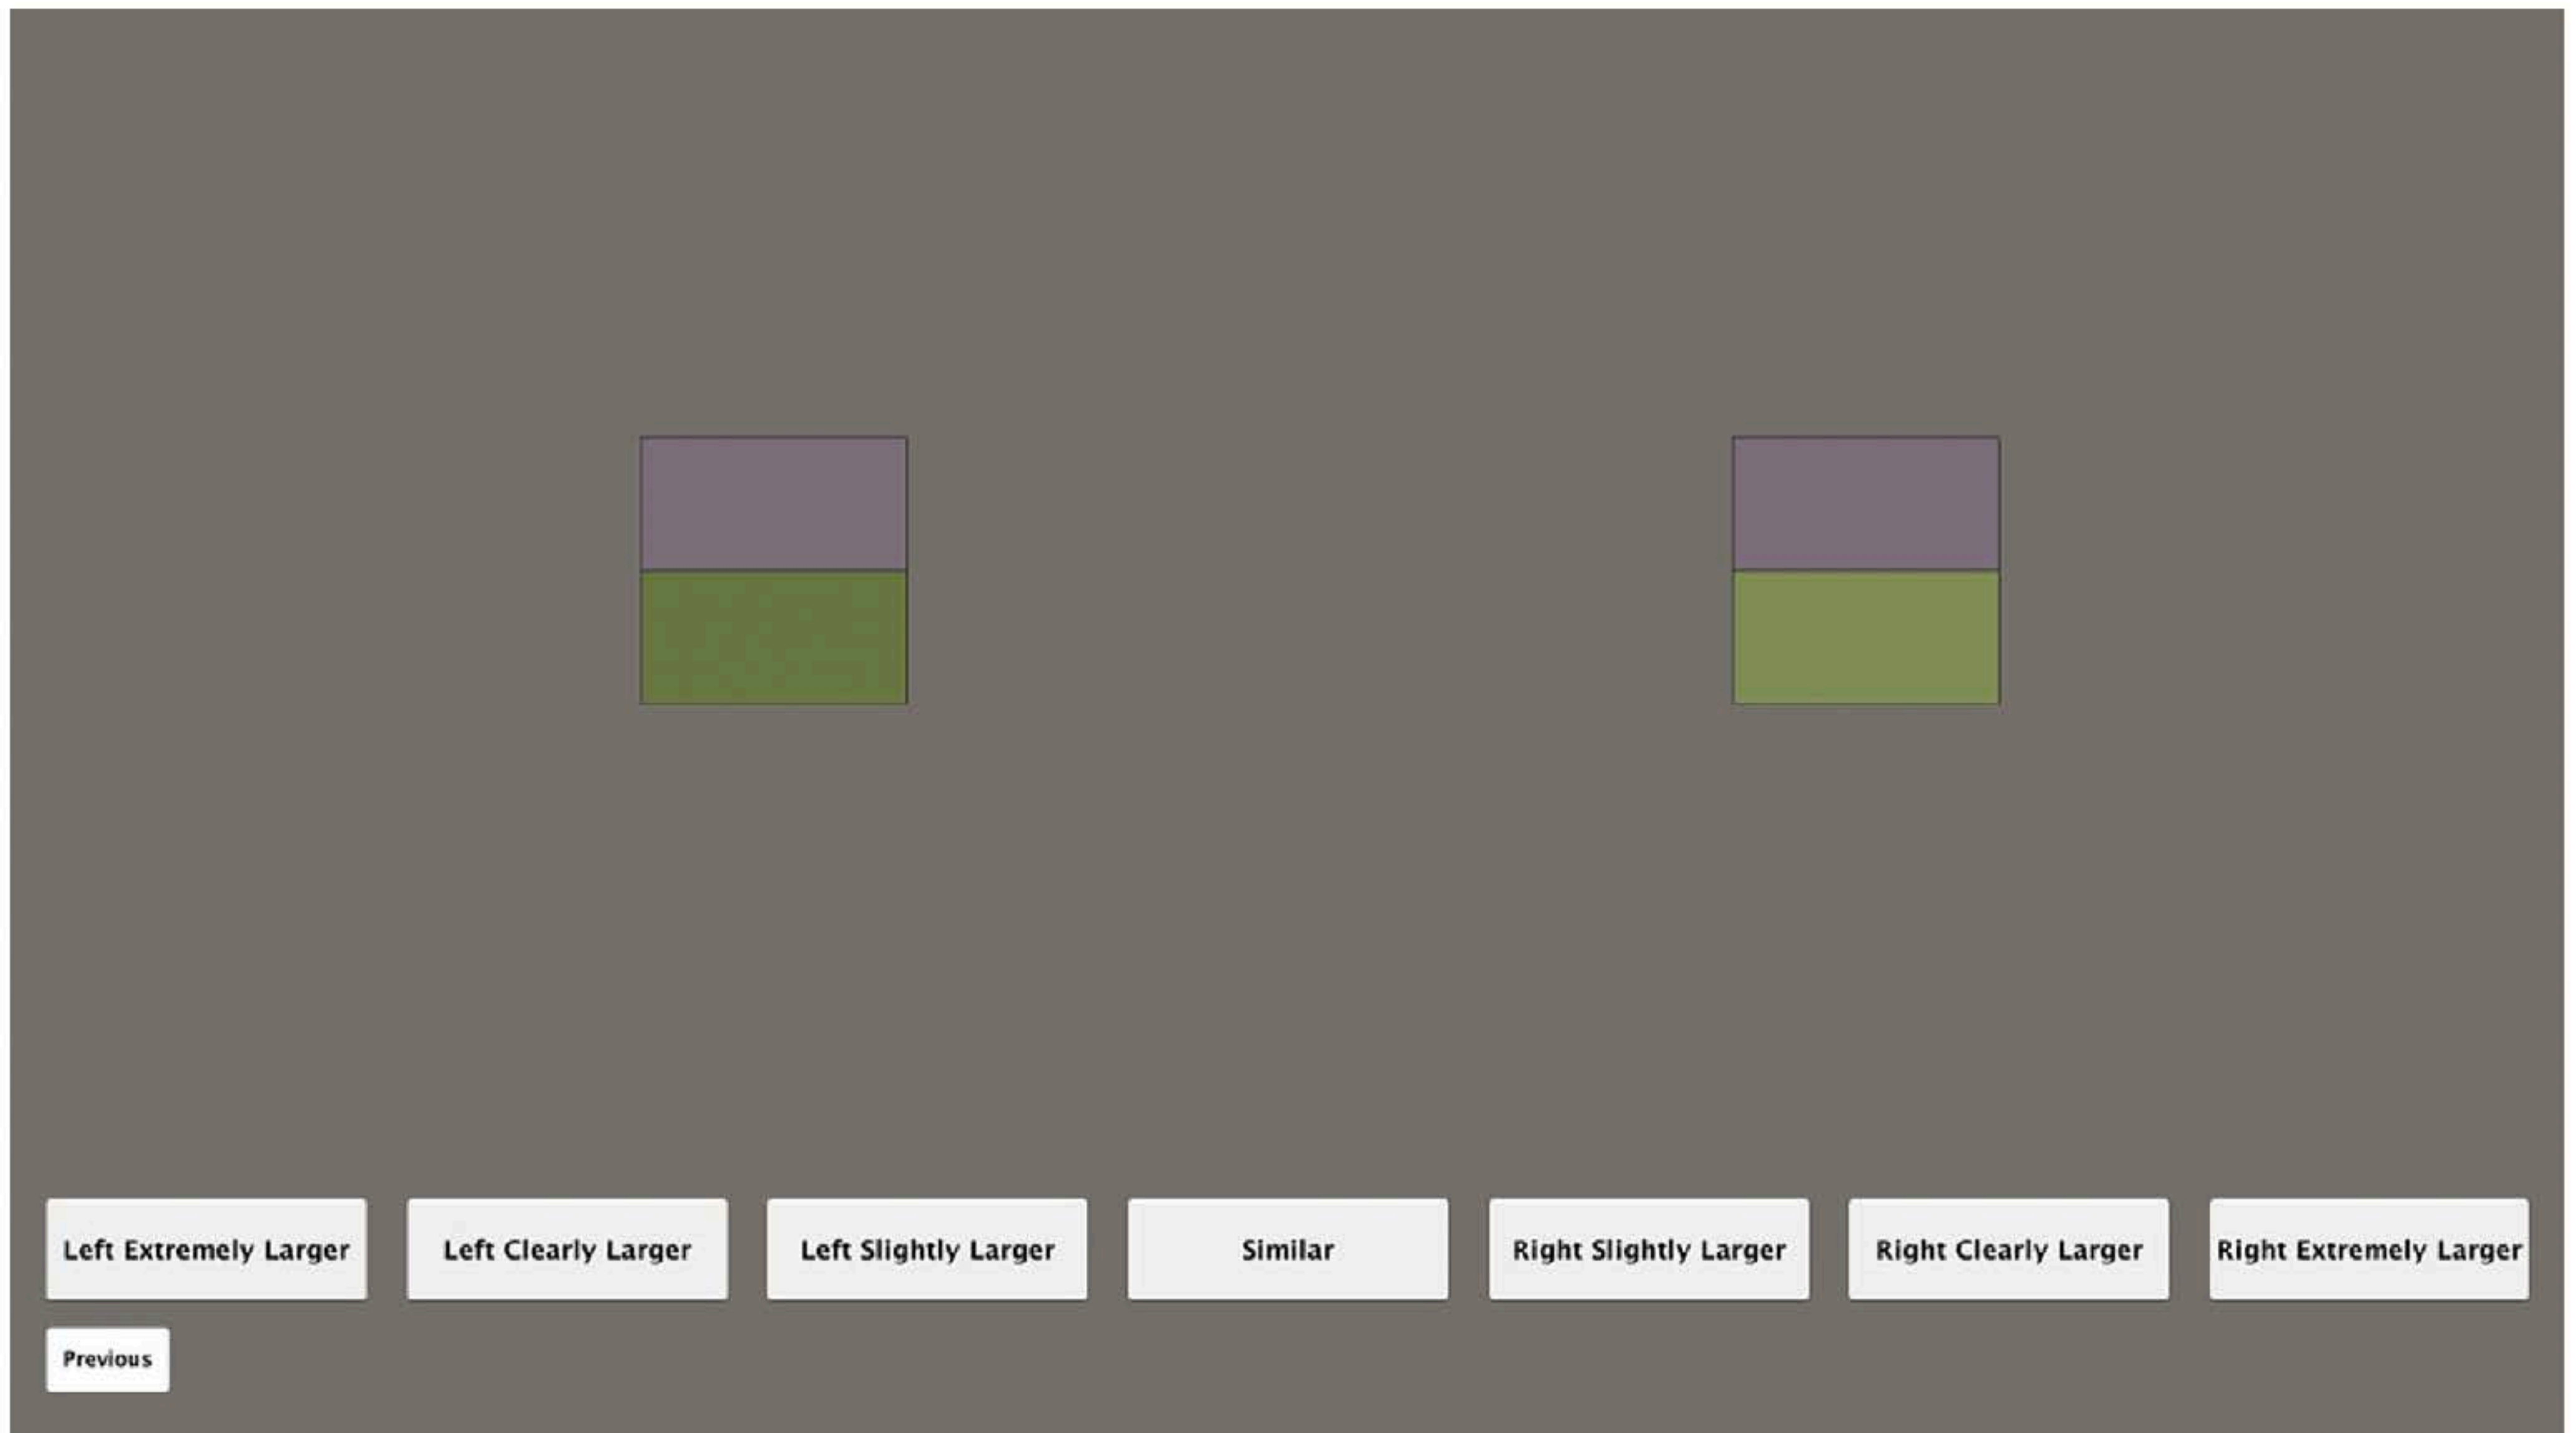


FIGURE 4 Screenshot of observer interface

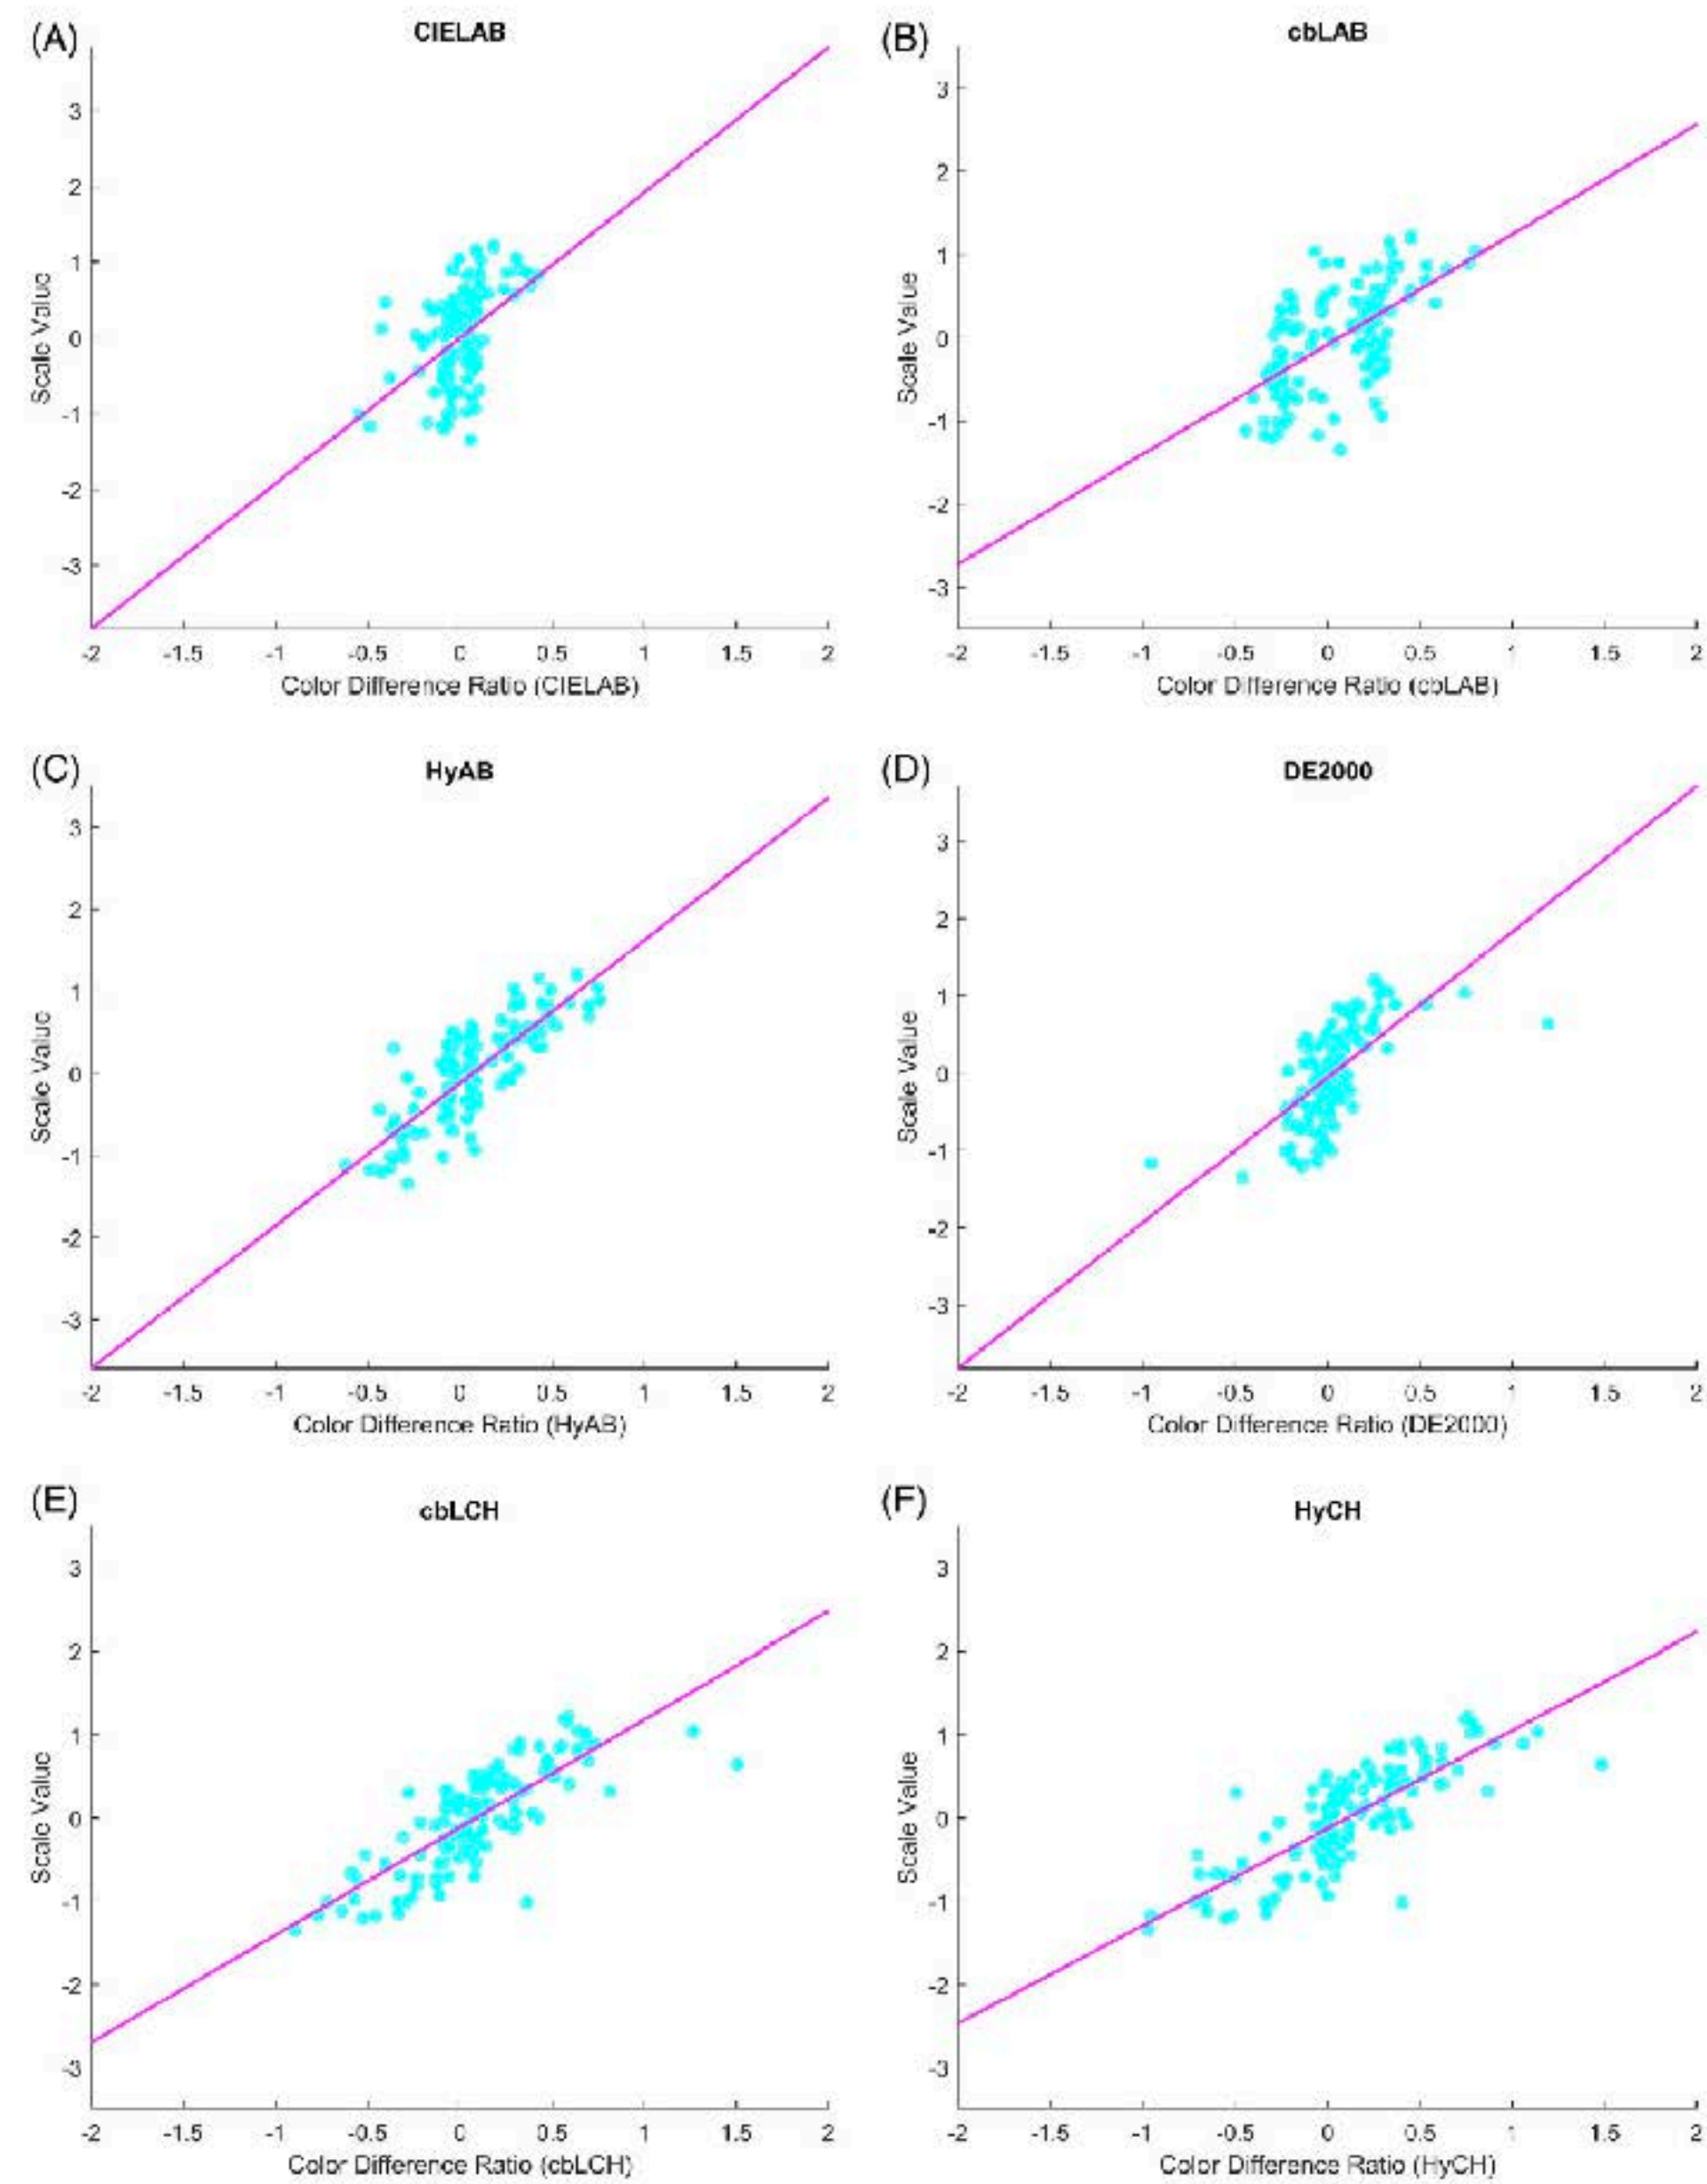


FIGURE 5 The color difference ratio for different formulas vs scale value for all stimuli of first dataset. A, CIELAB; B, cbLAB; HyAB; D, CIEDE2000; E, cbLCH; and F, HyCH

Finally, the overall distance between the two points is the sum of two distances. The color difference (CD) between two points using “ $\Delta L^* + \text{Euclidean}(a^*, b^*)$ ” is calculated using Equation (3):

$$\text{CD1} = |\Delta L^*| + \left[(a^*_1 - a^*_2)^2 + (b^*_1 - b^*_2)^2 \right]^{1/2} \quad (3)$$

The “ $\Delta L^* + \text{Euclidean}(a^*, b^*)$ ” is notated as the “HyAB” color difference formula in the rest of this article—“Hy” to indicate that it is a hybrid model, a combination of city block and Euclidean, and the last letters to show that the Euclidean dimensions are a^* and b^* .

FLIP

https://research.nvidia.com/publication/2020-07_FLIP

FLIP: A Difference Evaluator for Alternating Images

PONTUS ANDERSSON, NVIDIA

JIM NILSSON, NVIDIA

TOMAS AKENINE-MÖLLER, NVIDIA

MAGNUS OSKARSSON, Lund University

KALLE ÅSTRÖM, Lund University

MARK D. FAIRCHILD, Rochester Institute of Technology

Image quality measures are becoming increasingly important in the field of computer graphics. For example, there is currently a major focus on generating photorealistic images in real time by combining path tracing with denoising, for which such quality assessment is integral. We present FLIP, which is a difference evaluator with a particular focus on the differences between rendered images and corresponding ground truths. Our algorithm produces a map that approximates the difference perceived by humans when alternating between two images. FLIP is a combination of modified existing building blocks, and the net result is surprisingly powerful. We have compared our work against a wide range of existing image difference algorithms and we have visually inspected over a thousand image pairs that were either retrieved from image databases or generated in-house. We also present results of a user study which indicate that our method performs substantially better, on average, than the other algorithms. To facilitate the use of FLIP, we provide source code in C++, MATLAB, NumPy/SciPy, and PyTorch.




Fig. 1. The FFLIP pipeline.

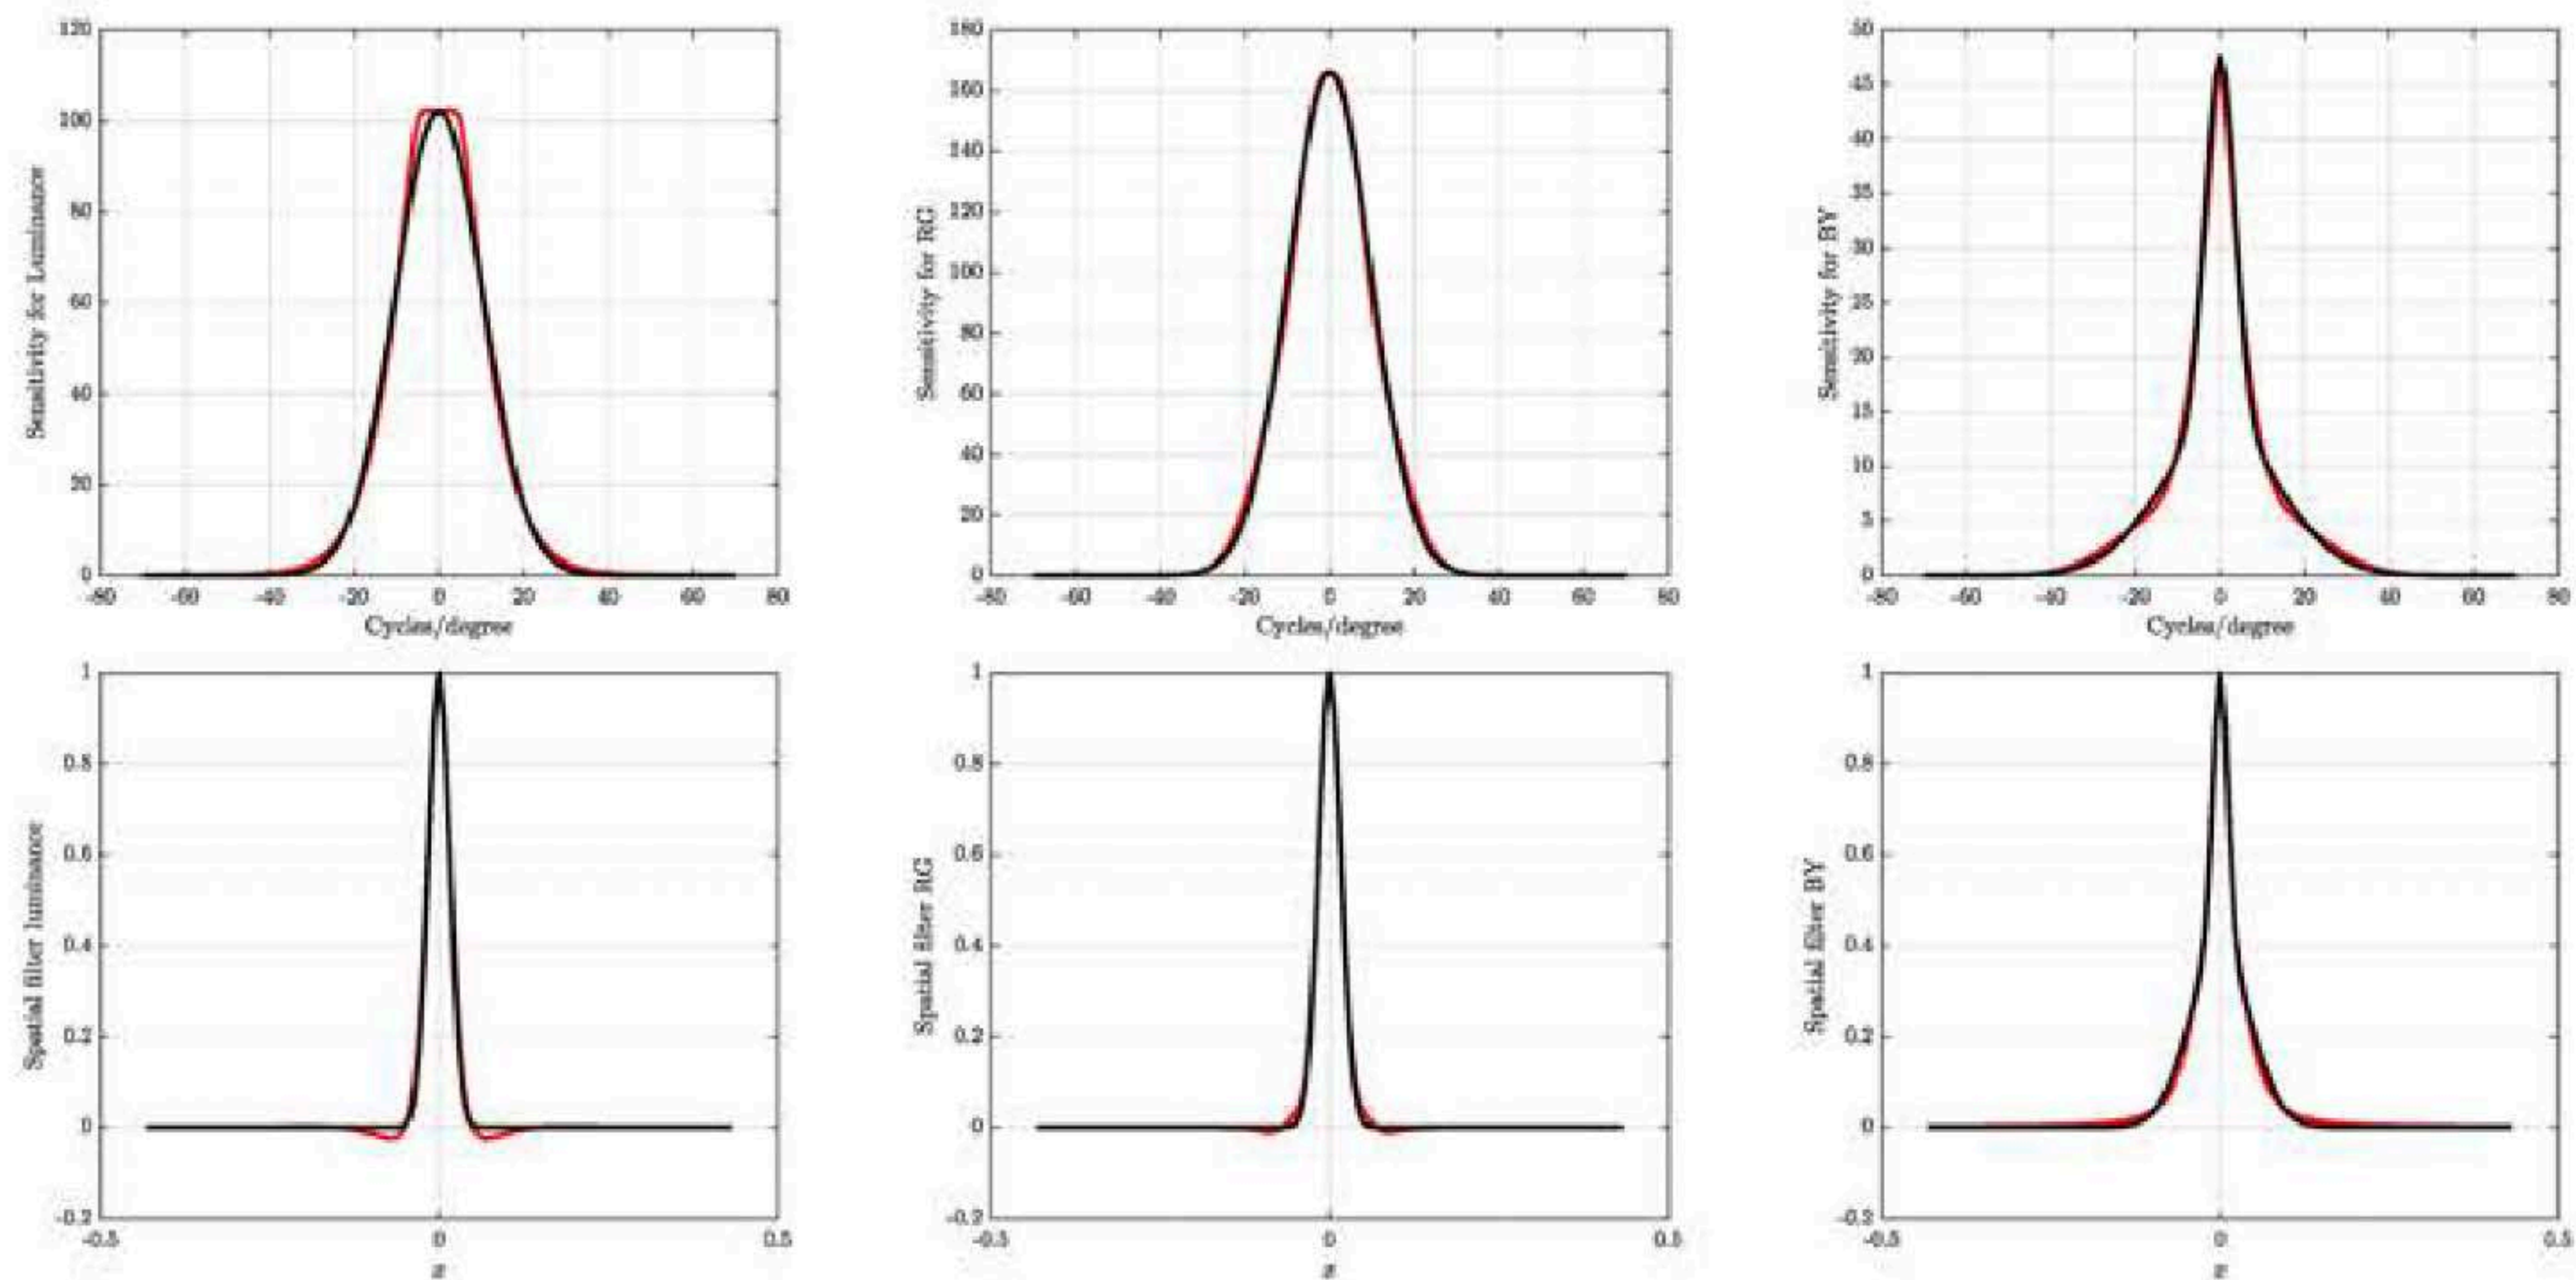


Fig. 2. Top: the original CSFs (Equation 2) are shown in red and our approximation, using a single Gaussian, or two in the case of the blue-yellow channel, in black. Bottom: the red curves are inverse discrete Fourier transformed versions of the CSFs in Equation 2, but they have negative lobes, which we wish to avoid. The black curves are the Gaussians from the top row, analytically inverse Fourier-transformed to the spatial domain. These do not include any negative lobes and are the ones used by FLIP.

Spatial Filtering



$$\Delta E_{\text{HyAB}}(\mathbf{R}, \mathbf{T}) = |L_{\mathbf{R}}^* - L_{\mathbf{T}}^*| + \sqrt{(a_{\mathbf{R}}^* - a_{\mathbf{T}}^*)^2 + (b_{\mathbf{R}}^* - b_{\mathbf{T}}^*)^2}. \quad (8)$$



Plus Hunt and remapping.




Fig. 3. Left: two colors with the same, high luminance, where the second color was generated using a rotation of the chrominance of the first. Right: two colors with the same, low luminance, with the same chrominance components as the colors in the first color pair. The difference between the colors is perceived as larger in the left pair due to the Hunt effect.




Fig. 4. Illustration of our error redistribution, compressing large color distances into a smaller range.




Fig. 5. From left to right: 1) the reference image, 2) the test image, including a small edge generated by a short ramp on either side and some salt and pepper noise, 3) the edge difference (first argument in Equation 9), 4) the point difference (second argument in Equation 9), and 5) the maximum of the edge and point difference in each pixel. Note that all differences are scaled for visibility.

pixels. This is illustrated in Figure 5 where we see how the edge error is large along the edge, but small on the pixels containing salt and pepper noise, while the opposite is true for the point error. Based on this, we choose the feature difference, ΔE_f , between pixels in the reference image, \mathbf{R} , and the test image, \mathbf{T} , to be the maximum of the differences in edge and point feature values, i.e.,

$$\Delta E_f = \left(\frac{1}{\sqrt{2}} \max (|\|\nabla \mathbf{R}\| - \|\nabla \mathbf{T}\||, |\|\nabla^2 \mathbf{R}\| - \|\nabla^2 \mathbf{T}\||) \right)^{q_f}, \quad (9)$$



Fig. 6. From left to right: 1) reference image, 2) test image, including aliasing artifacts, 3) color error, ΔE_c , produced by the pipeline in Section 4.1, and 4) color error enhanced based on feature differences (Equation 10). The heatmap we use is shown to the right, where bright yellow corresponds to maximum error and dark to minimum error. Note that this heatmap was selected since its luminance is a linear ramp from dark to bright. The images are have been enlarged to more clearly show the differences.

$$\Delta E = (\Delta E_c)^{1-\Delta E_f} . \quad (10)$$

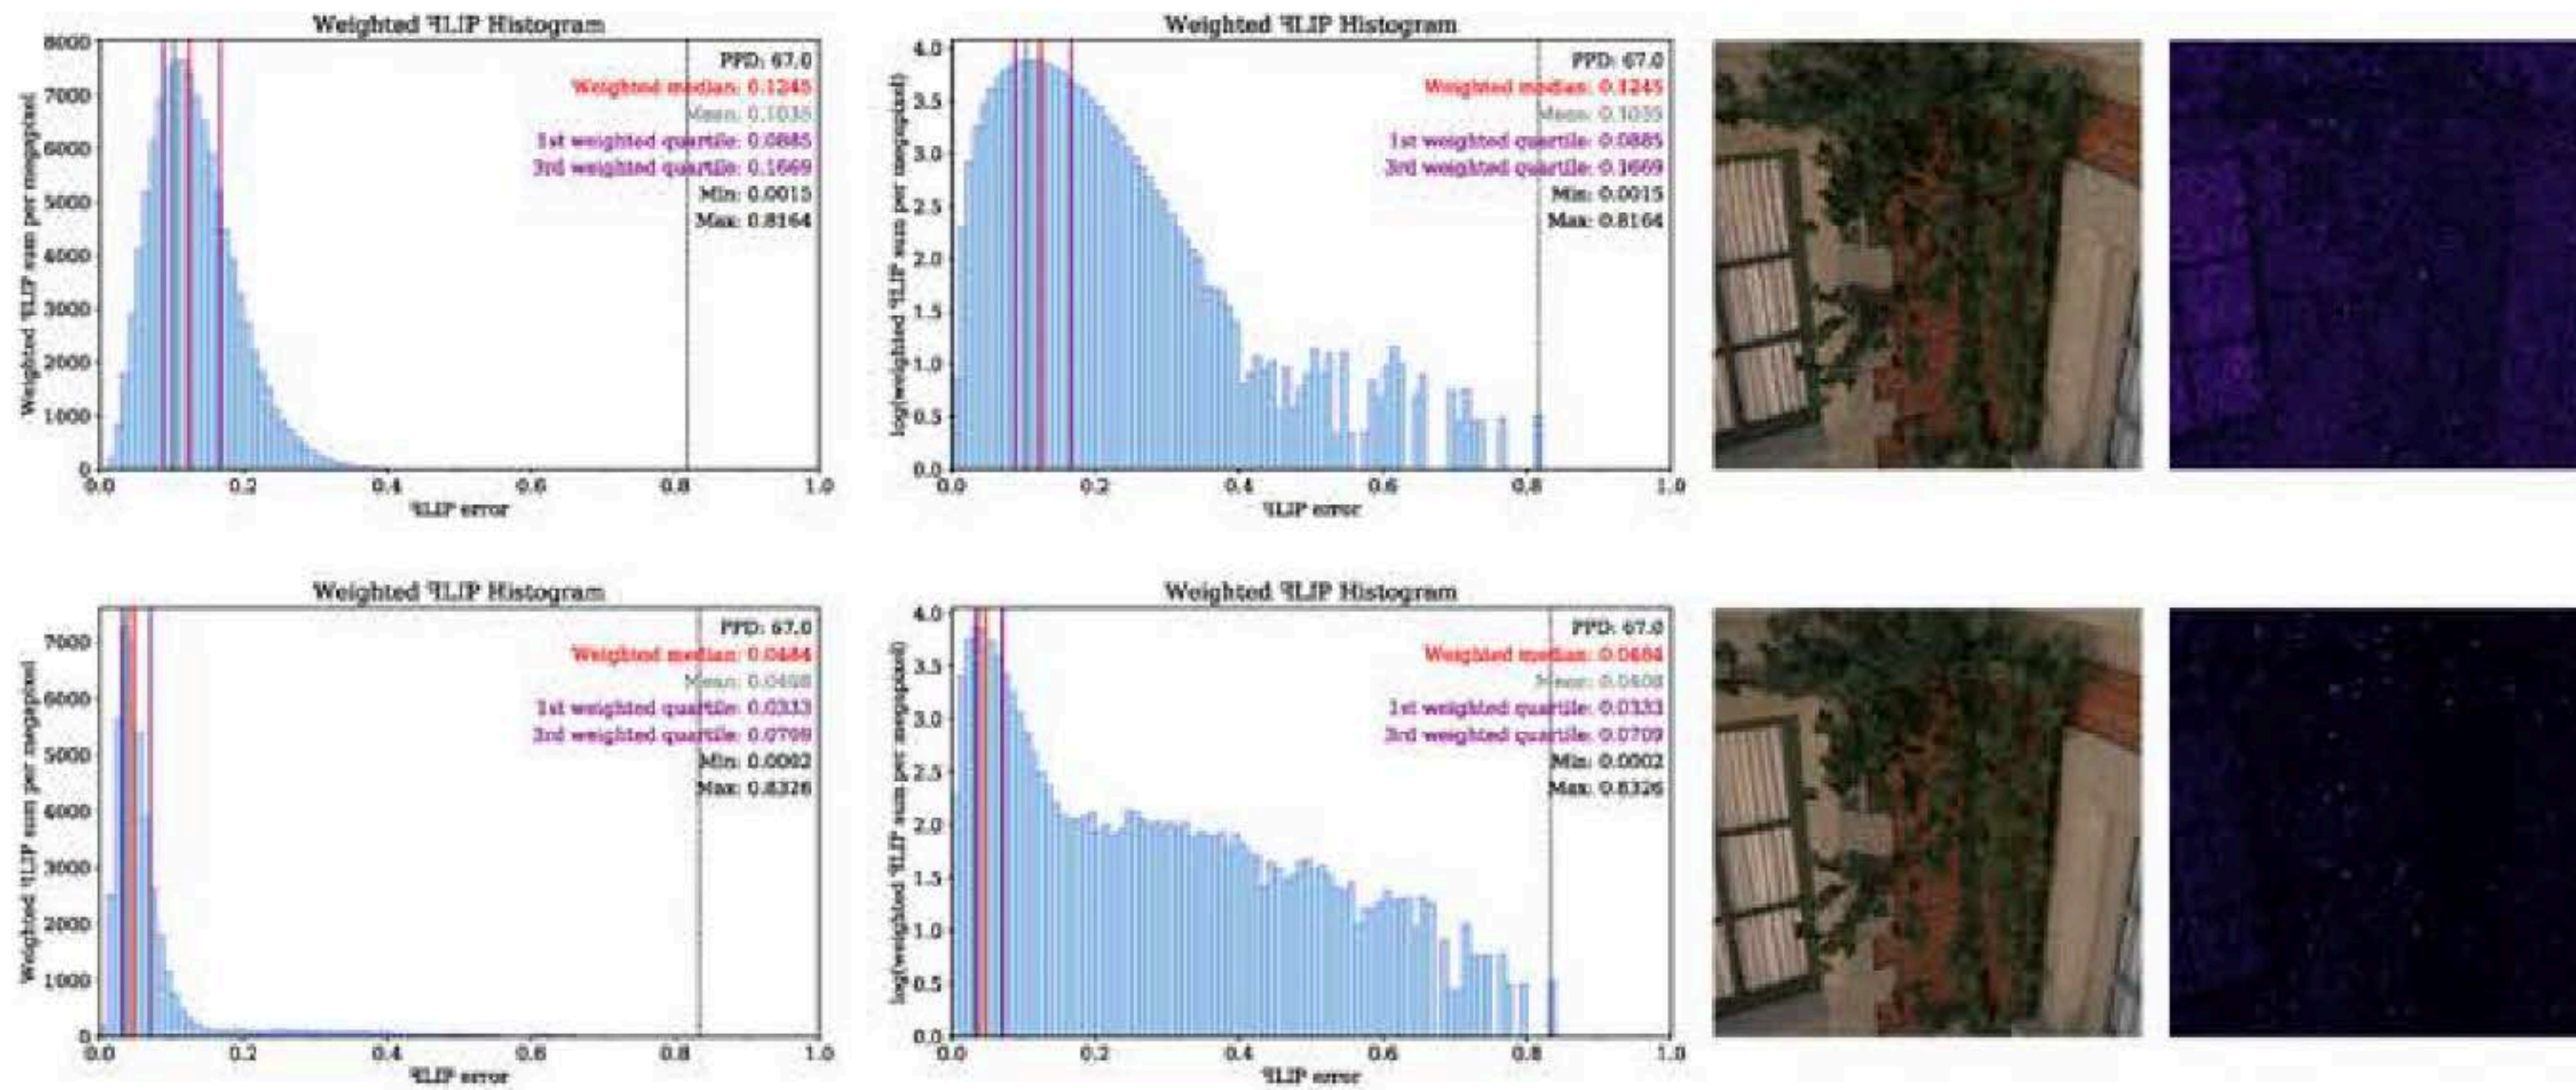


Fig. 7. Top row, from left to right: 1) a weighted FLIP histogram, generated using the third image on this row, showing the weighted median, the mean, the 1st and 3rd weighted quartiles, and the minimum and the maximum value. 2) same as to the left, except that the y -axis is in log space, resulting in the larger errors (with low counts) becoming more pronounced, which may be beneficial in some cases. In the left diagram, it can be seen that the area to the left of the weighted median is the same as the area to the right. 3) the path traced test image with 32 samples per pixel (SPP). 4) the FLIP error map, which was computed against a path traced reference image with 2^{22} SPP. Bottom row: same as the top row, except that the test image was path traced using 1,024 SPP. Note that the image on the bottom row has more pixels with medium to high errors and that this is clearly visible only in the log diagram.

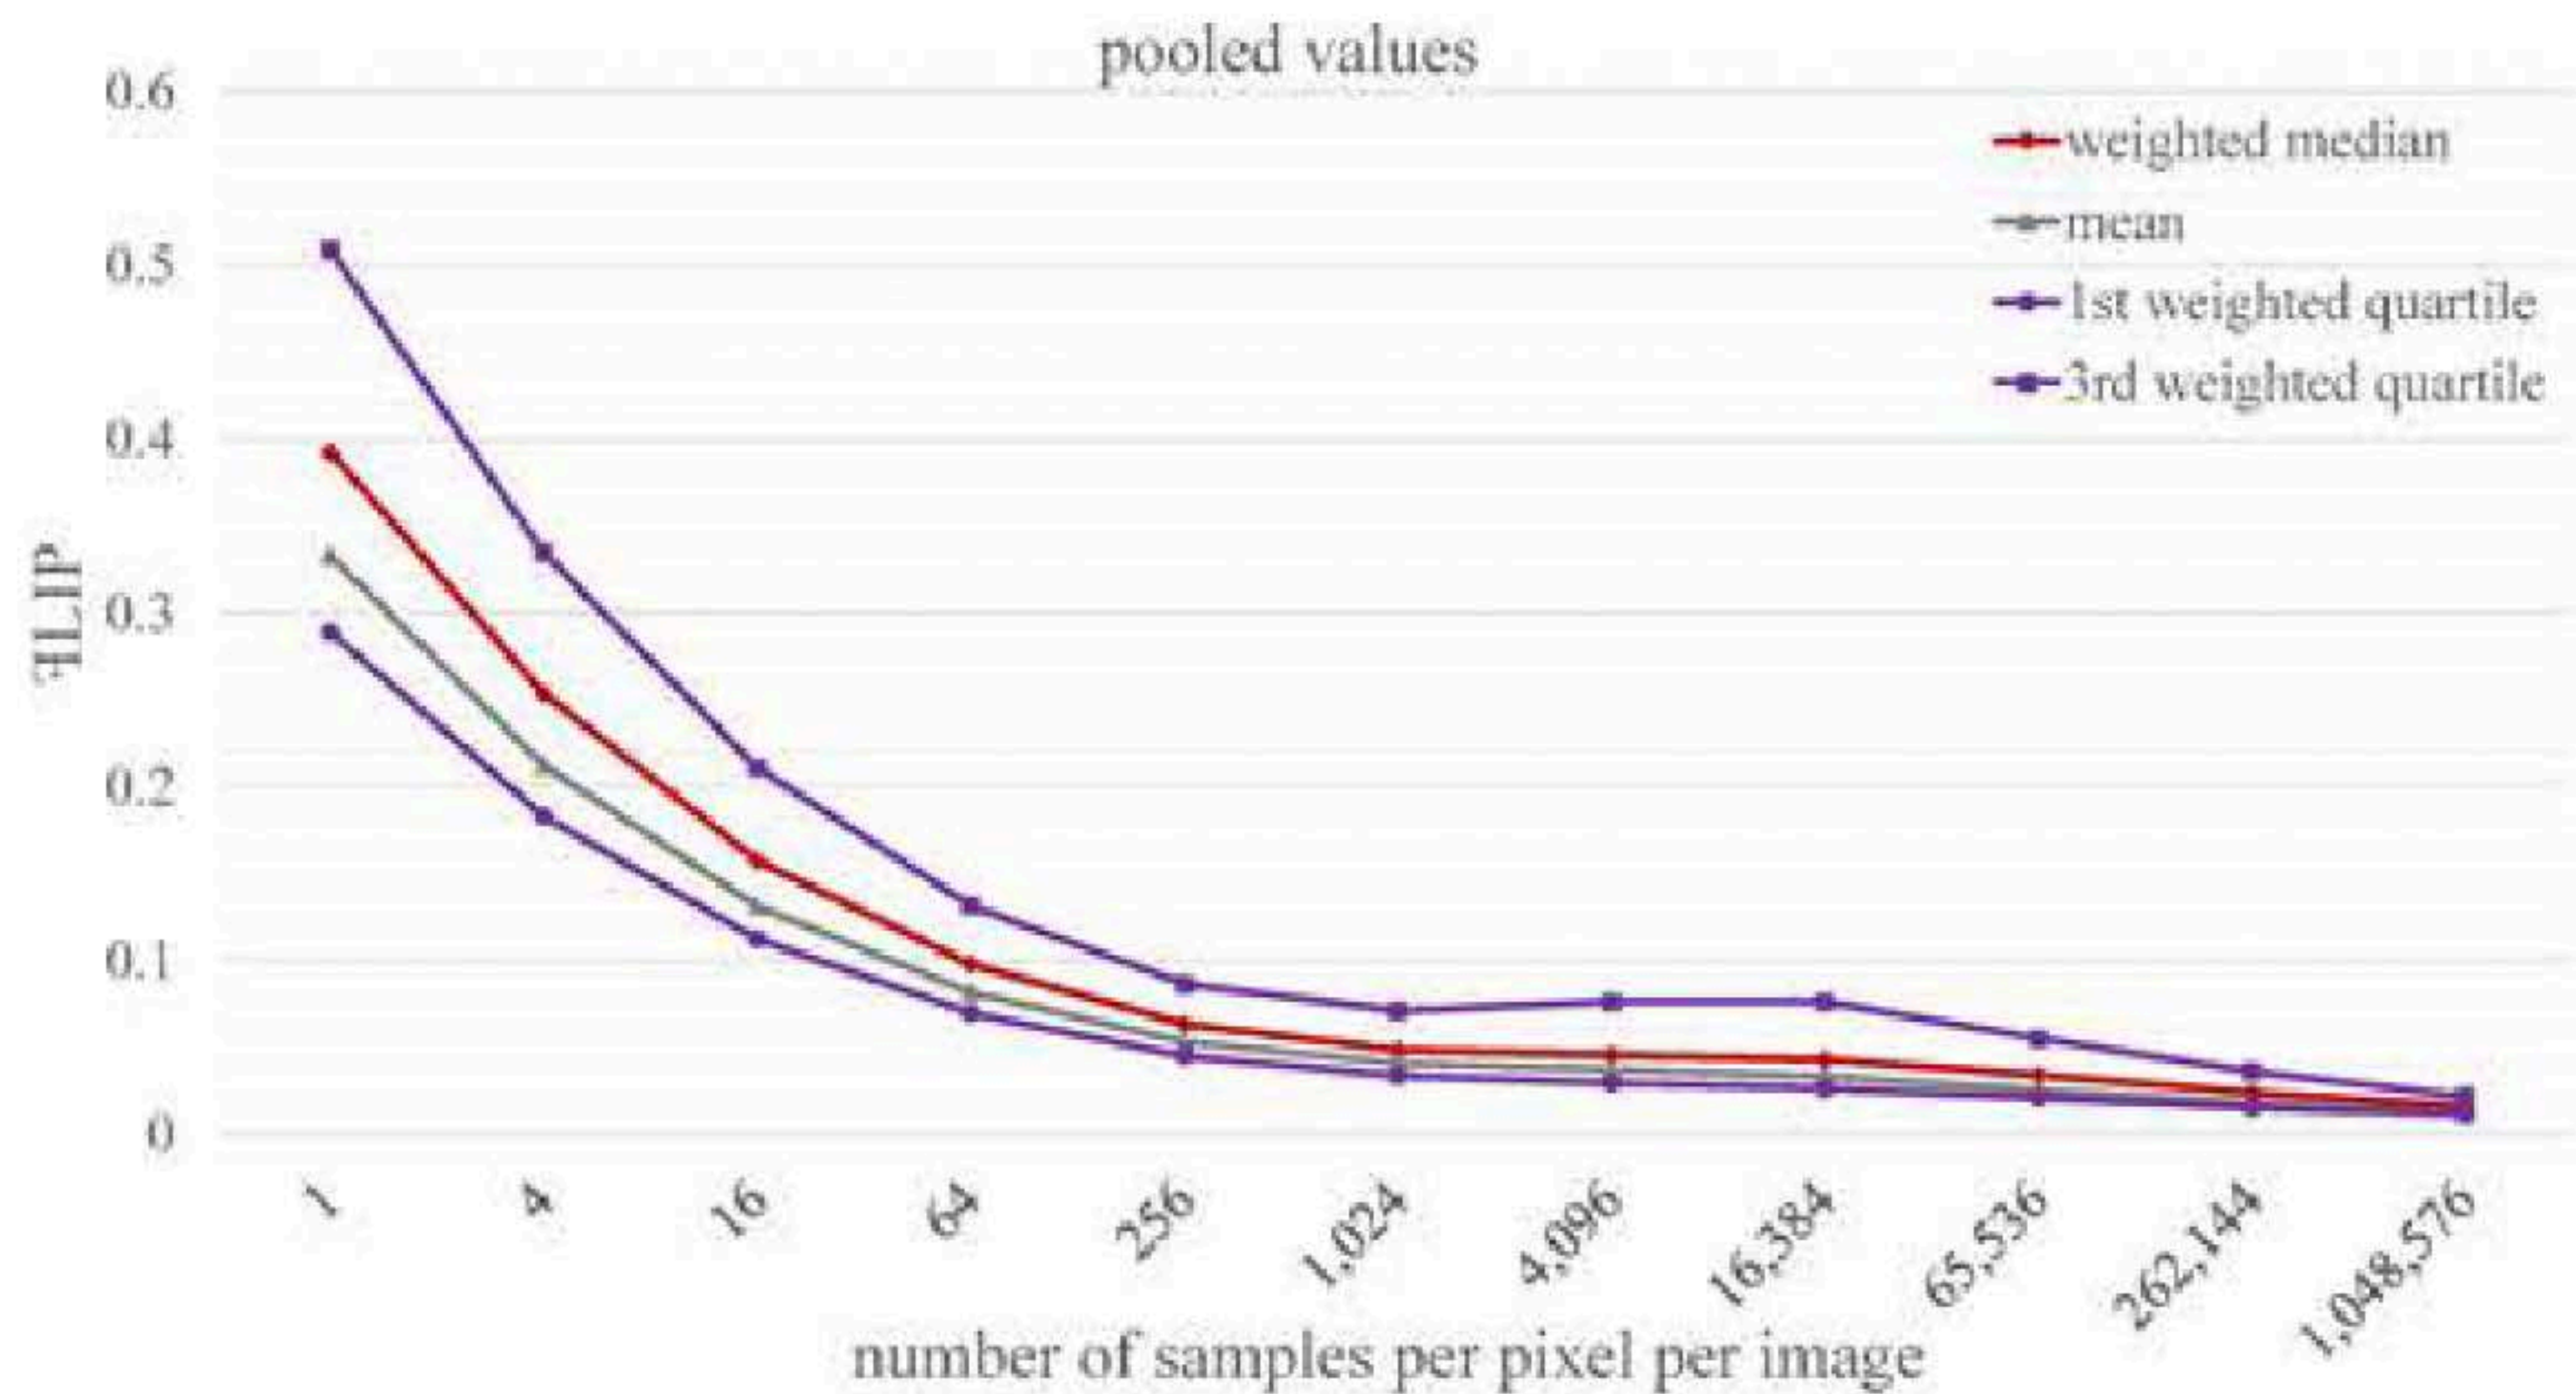


Fig. 8. Pooled FLIP values for the Bistro leaves images, partially shown in Figures 7 and 10. Note that the third quartile has a local peak at 4,096 and 16,348 SPP. The reason for this is that the corresponding images have an increased number of fireflies compared to the lower SPP images, which can be seen when zooming in on the corresponding images in Figure 10.




Fig. 9. A collection of example images and corresponding error maps. Each column contains one reference image, one test image, the FLIP map, and the error map produced by one of the other metrics which FLIP is compared against. Recall that FLIP has been developed for the flip test, so instead of zooming in on the images in the paper, we recommend that the reader looks at the images in our supplemental material at a distance such that $p = 67$ pixels per degree. Here, image pairs 1–5 are rendered, while image pairs 6–8 are natural images. Image pair 6 is courtesy of NASA and image pair 7 of Lambert/Keystone/Getty Images.




Fig. 10. Here, we show an example of the progression of a naive path tracer. In the leftmost column, we show the same reference images in each row (generated with 2^{22} samples per pixel), followed by test images, where each test image has $4\times$ more samples than the one above it. We then present the corresponding error maps, for each test image and each metric, in the remaining columns. Note that there is a substantial amount of firefly detection with \mathcal{FLIP} in the images with 4,096 and 16,384 SPP (visible when zooming in).




Fig. 11. A compressed photograph exemplifying FLIP's tendency to overestimate the perceived error in the presence of masking.

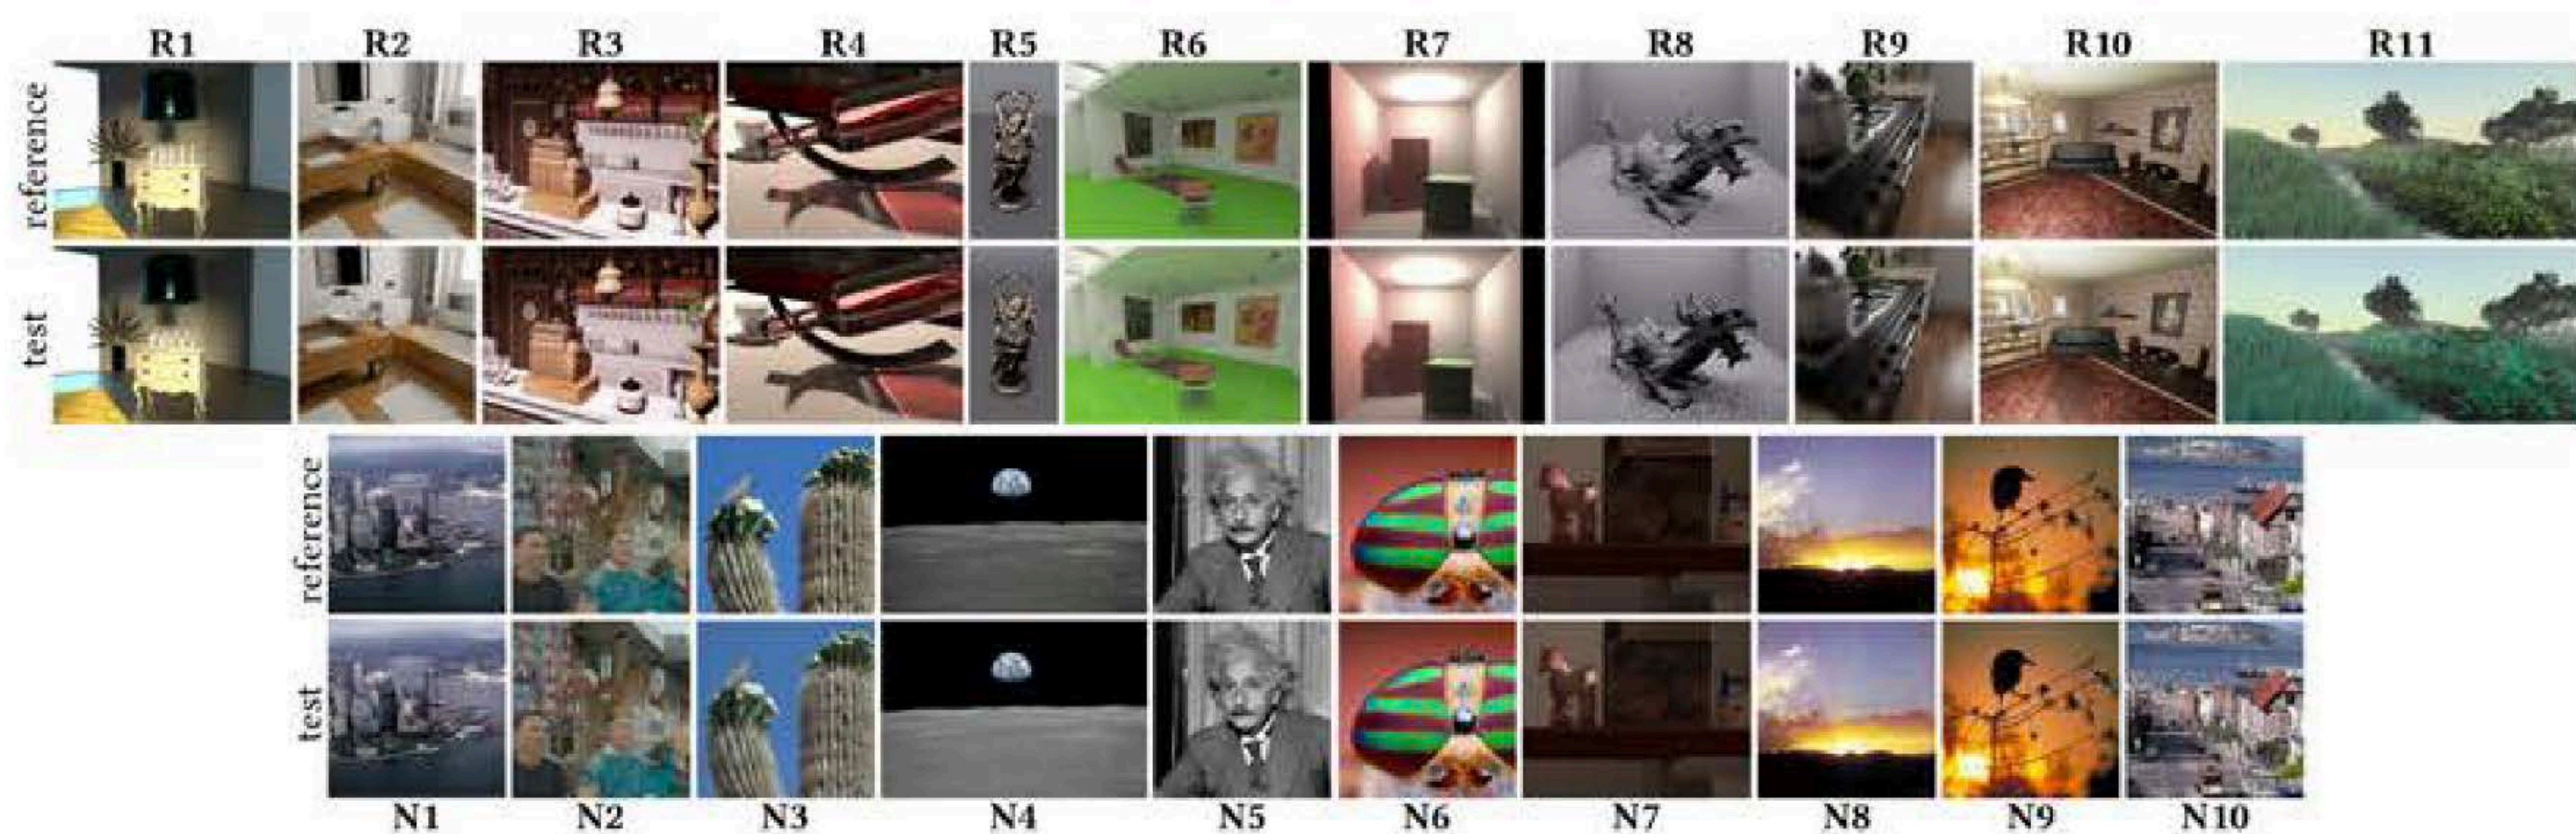


Fig. 12. The image pairs used in our user study. The top row shows rendered images, while the bottom row shows natural images with various distortions. All images, together with the error maps generated by the different metrics used in the study, are included in our supplemental material. **N4** is courtesy of NASA and **N5** is based on a photo by Lambert/Keystone/Getty Images.




Fig. 13. Left: the results from our user study with per-image and per-metric averages. The rendered images (**R_x**) and the natural images (**N_x**) are the ones shown in Figure 12. Right: averages and 95% confidence intervals over all images (gray), only rendered images (orange), and only natural images (blue), respectively, where the x -axis is the score, from 0 (minimum) to 3 (maximum), in the user study.

Not Secure — www.240hz.org/image-metrics/local/results/earthmoon/index.html


ISA PoCS/MCSL mdf.org COS Cal RIT ... Weather News The Correspondent Apple Temp OldFaithful AmezcnSmile MpixPro RW isetbio Grapes & Wines Winemaking | ...al Education

https://www.spectrum.net/api/pub/billing/v1/pdf-statement/MjAyMTU0MDQ3Ny0wMS0wMDFfMjAxOS0xMS0xOHwx?inlin... Image: earthmoon


local: earthmoon

Mouse wheel to zoom in/out. Click and hold to show pixel data. Click + drag to pan while zoomed in. Press '0' to reset zoom.


Reference Image
 Test Image
 Reference/Test Flip Image
 Our
 HyAB (Fairchild)
 RGB Difference
 SSIM
 HDR-VDP-2
 Johnson/Fairchild
 SMAPE




Reference




Test




Our Distance



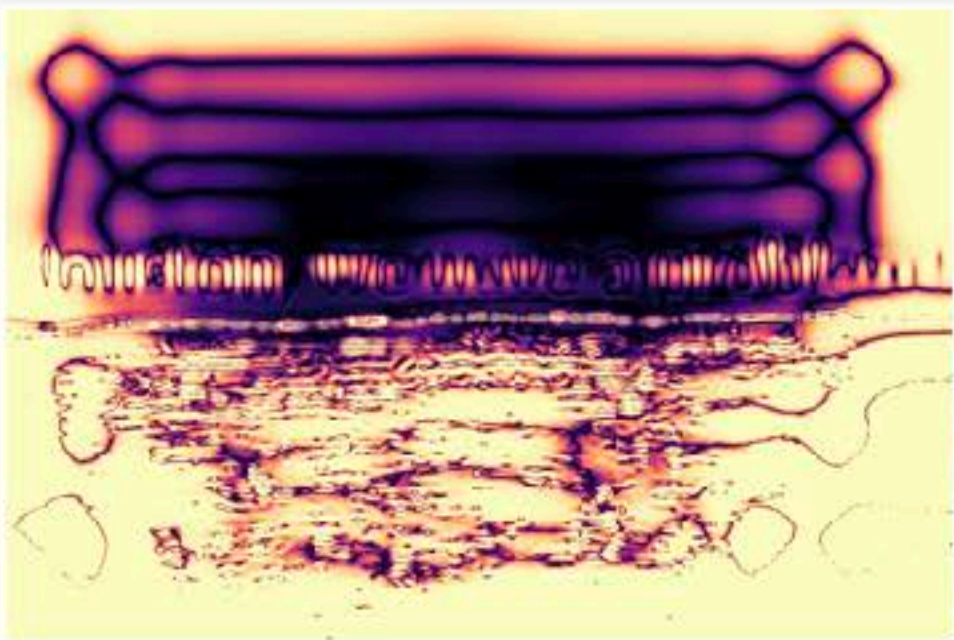
HyAB Distance (Fairchild app. space)




RGB Distance




SSIM



HDR-VDP-2



Johnson/Fairchild Distance



SMAPE Distance

https://research.nvidia.com/publication/2020-07_FLIP

Future: A New Colorimetry

Individual Cone Responsivities

vK20 Chromatic Adaptation Model

Color Appearance Scales & Differences

(Lightness, Saturation, Hue, Brightness, Colorfulness, Chroma)

Image Filtering/Differences Can Be Applied on Scales

Thank You / Gracias

mark.fairchild@rit.edu