Supplementary Material:

Adaptive Color Display via Perceptually-driven Factored Spectral Projection

Isaac Kauvar[1] Samuel J. Yang[1] Liang Shi[1] Ian McDowall[2] Gordon Wetzstein[1]
[1]Stanford University [2]Intuitive Surgical/Fakespace Labs

In this document we provide additional discussion and results in support of the primary text.

A Additional Results

In this section, we show all processed example scenes with three and also four primaries as well as comparisons between NMF, PNMF, and NMF with the largest fixed gamut for the 3-primary case. For each example, we show the simulated sRGB image, the error maps in CIE Lab space, the computed gamut (2D and 3D view), convergence plots for ΔE_{76}, ΔE_{94}, and CIEDE2000, and also the color-coded factorized patterns.

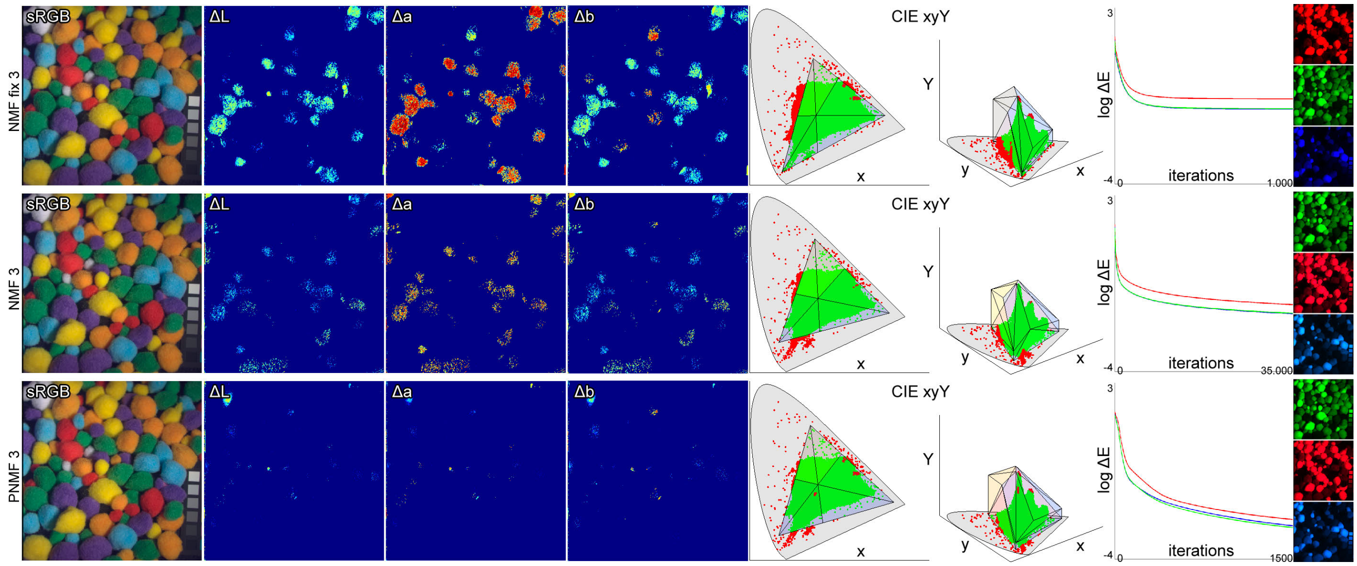

Figure S.1: *“Pom pom” scene for 3 primaries.*

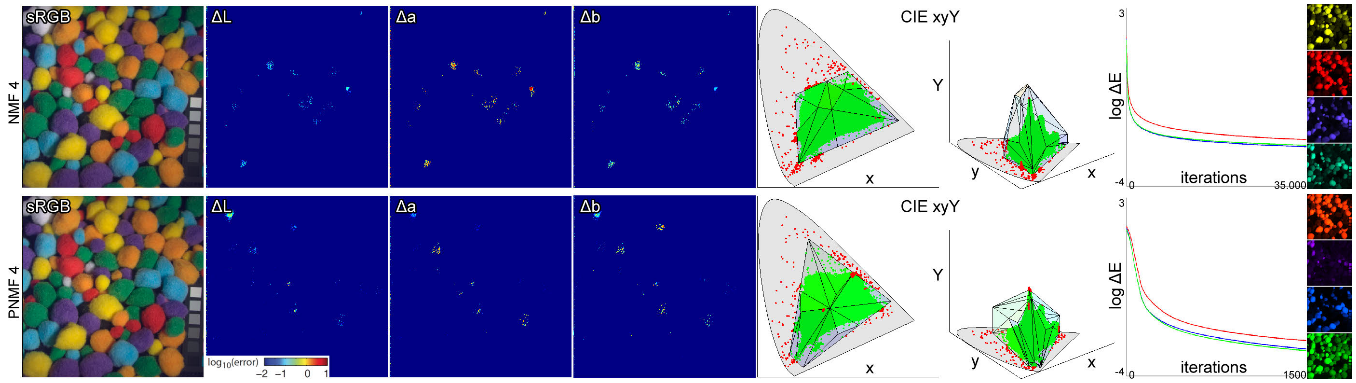

Figure S.2: *“Pom pom” scene for 4 primaries.*

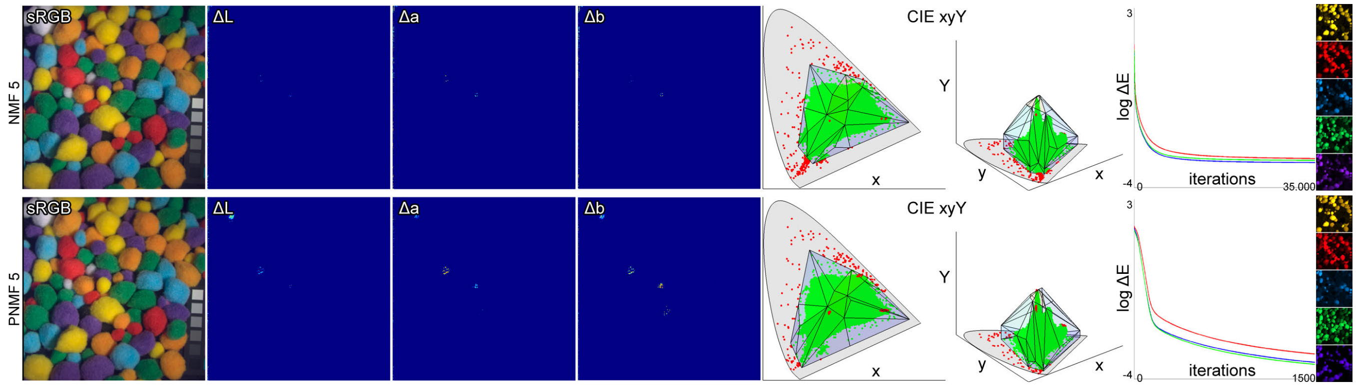

Figure S.3: *“Pom pom” scene for 5 primaries.*

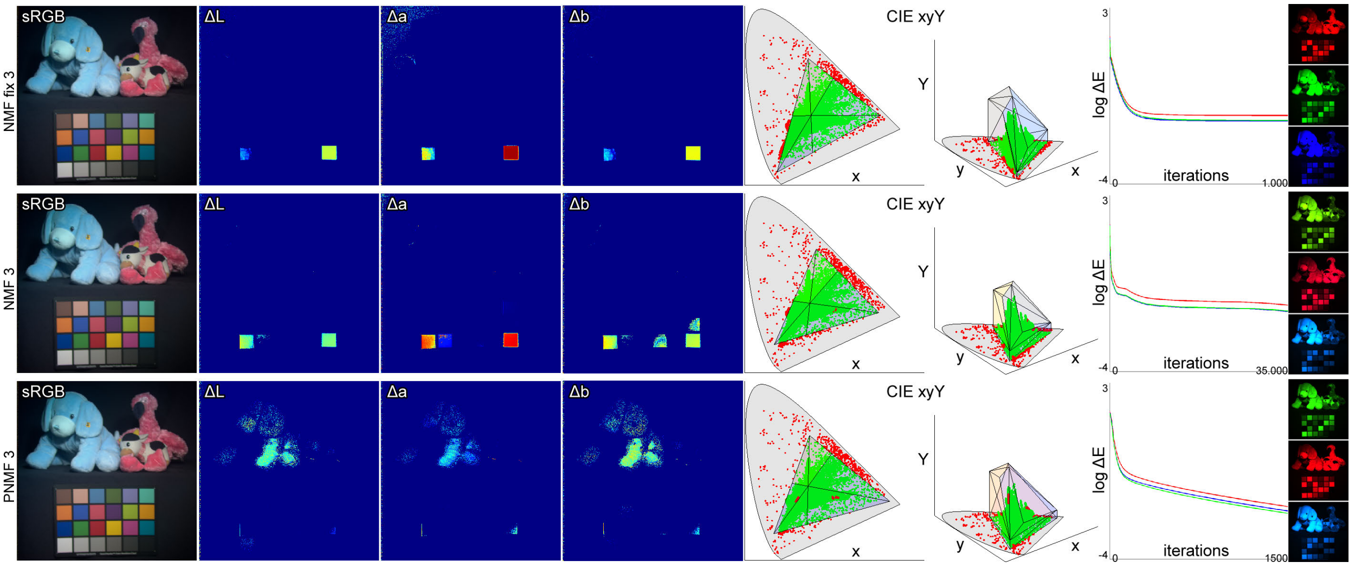

Figure S.4: *"Dog" scene for 3 primaries.*

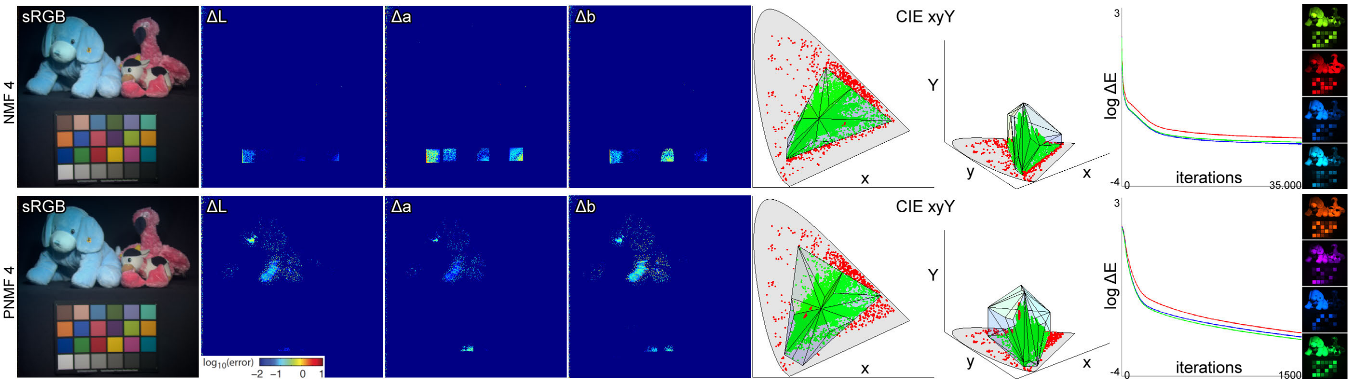

Figure S.5: *"Dog" scene for 4 primaries.*

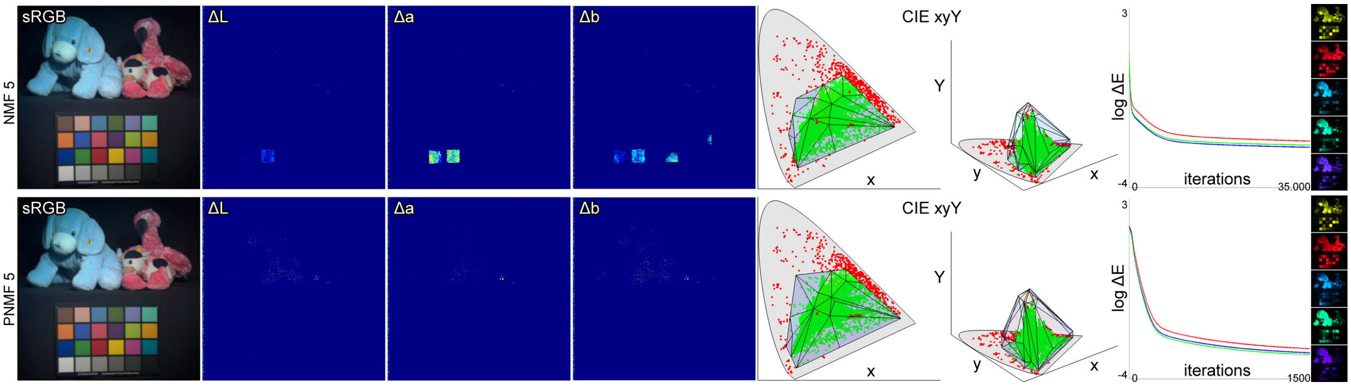

Figure S.6: *"Dog" scene for 5 primaries.*

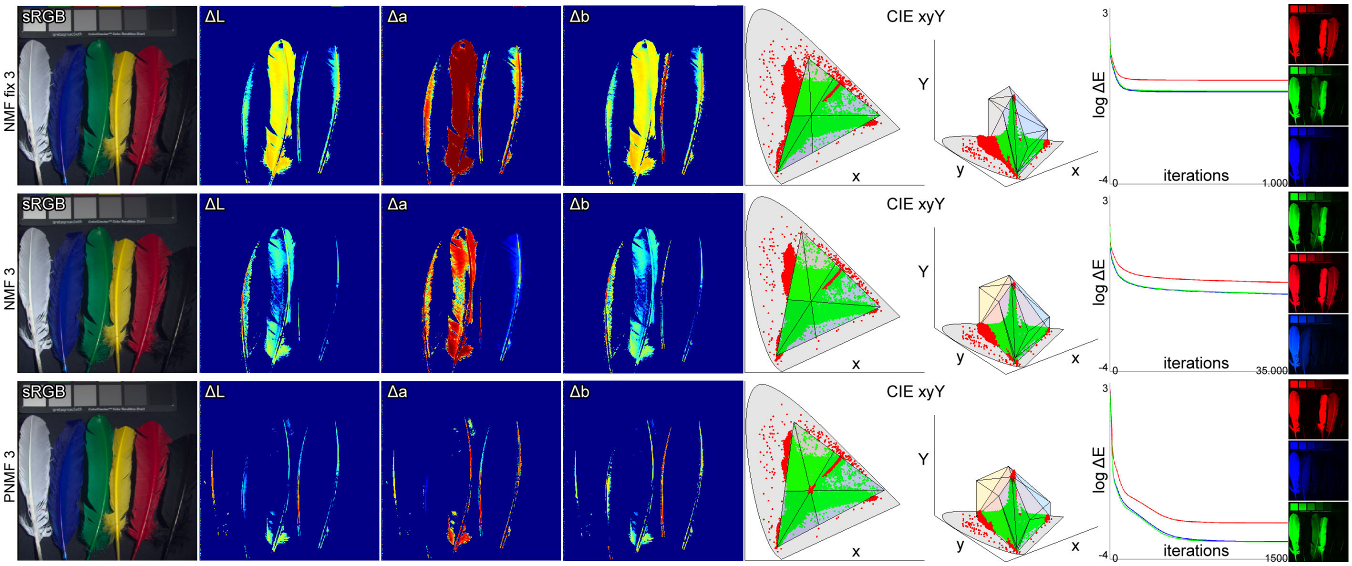

Figure S.7: *"Feather" scene for 3 primaries.*

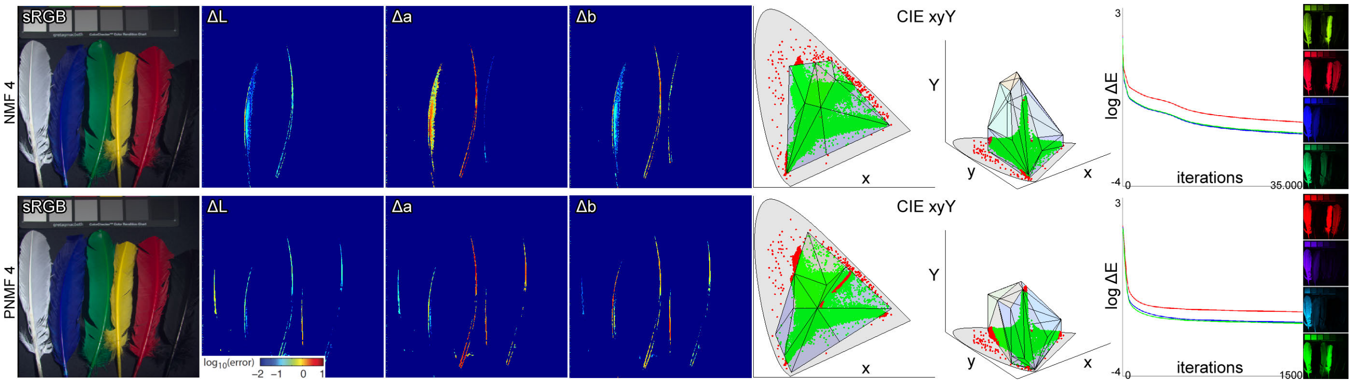

Figure S.8: *"Feather" scene for 4 primaries.*

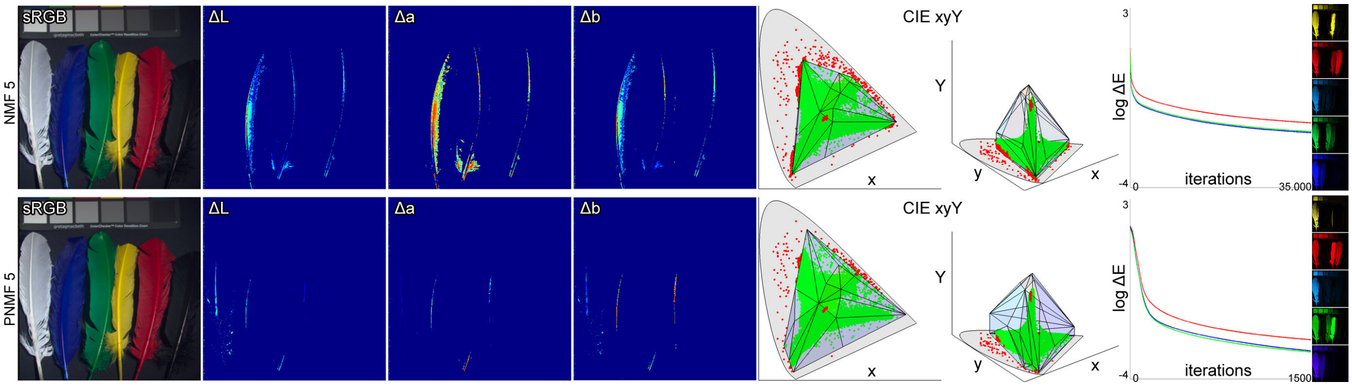

Figure S.9: *"Feather" scene for 5 primaries.*

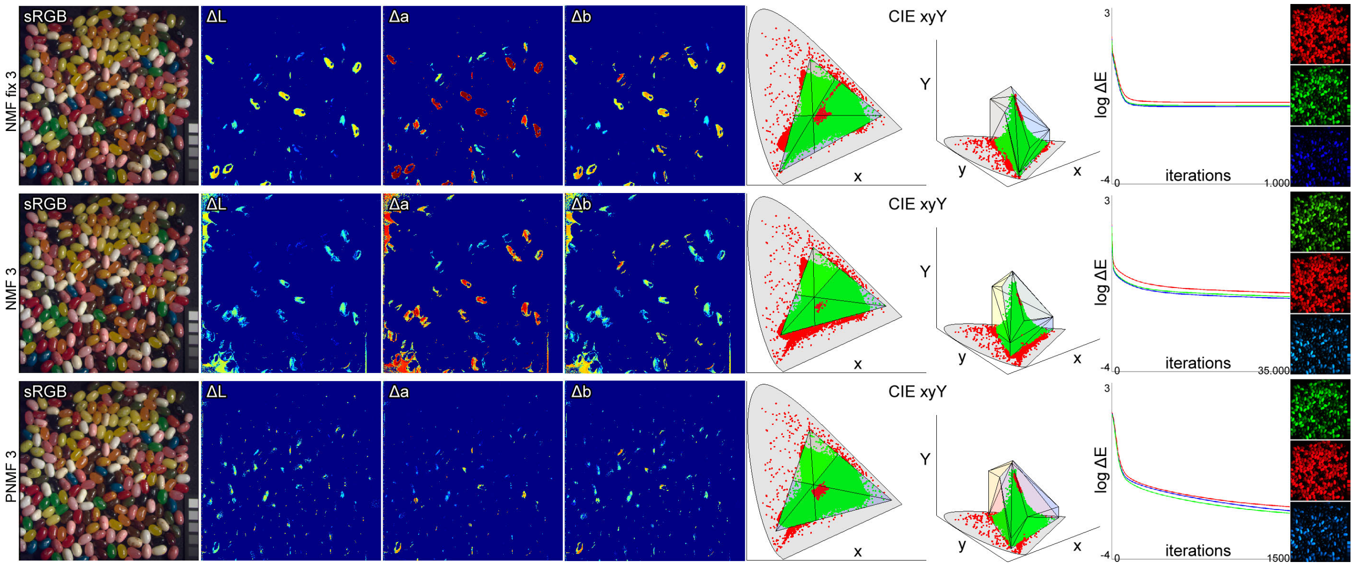

Figure S.10: *"Beans" scene for 3 primaries.*

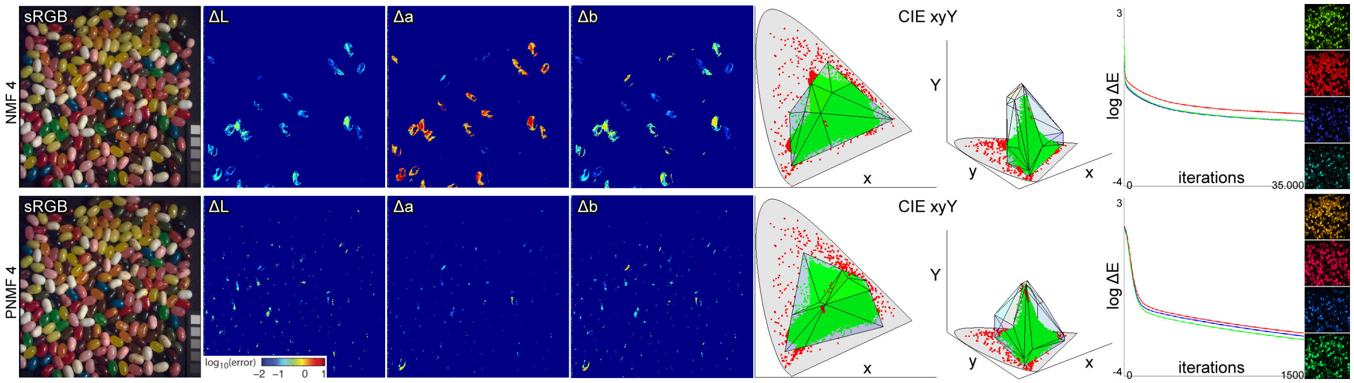

Figure S.11: *"Beans" scene for 4 primaries.*

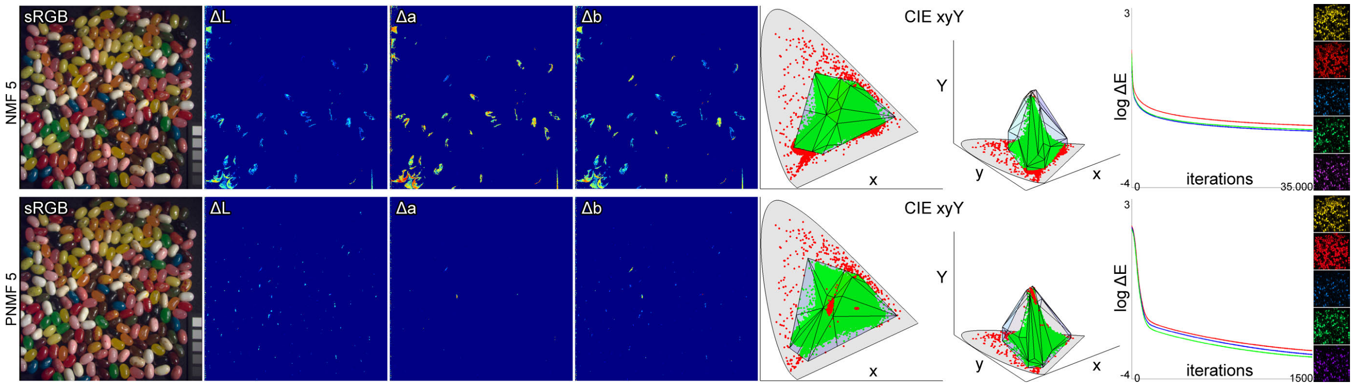

Figure S.12: *"Beans" scene for 5 primaries.*

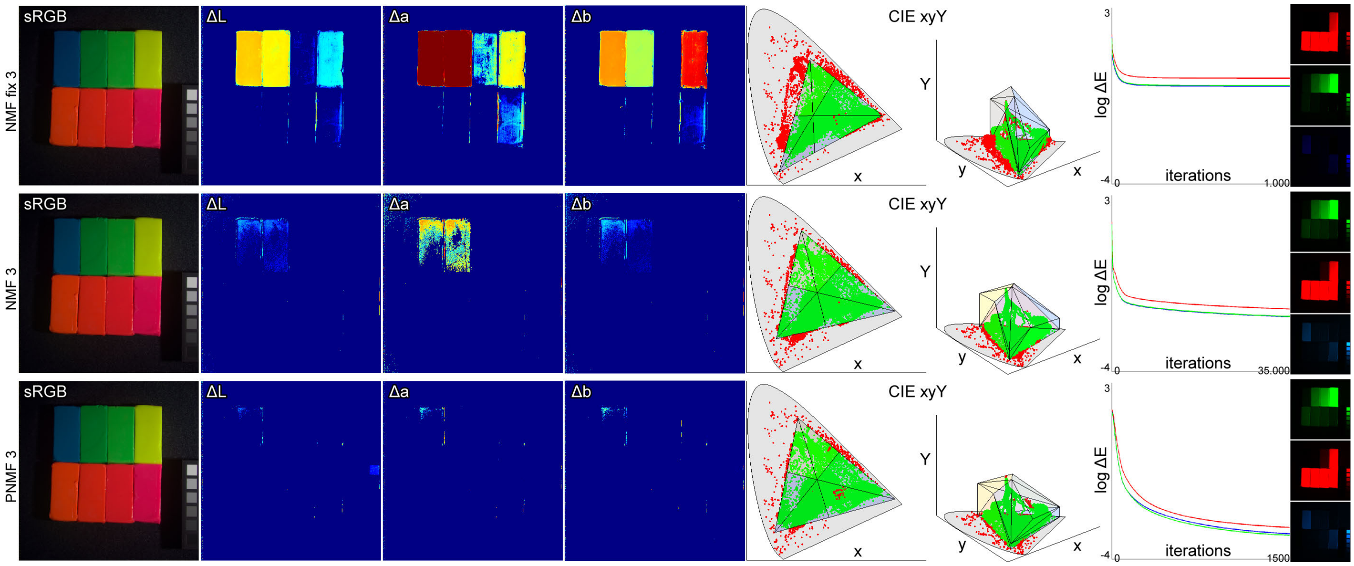

Figure S.13: *"Clay" scene for 3 primaries.*

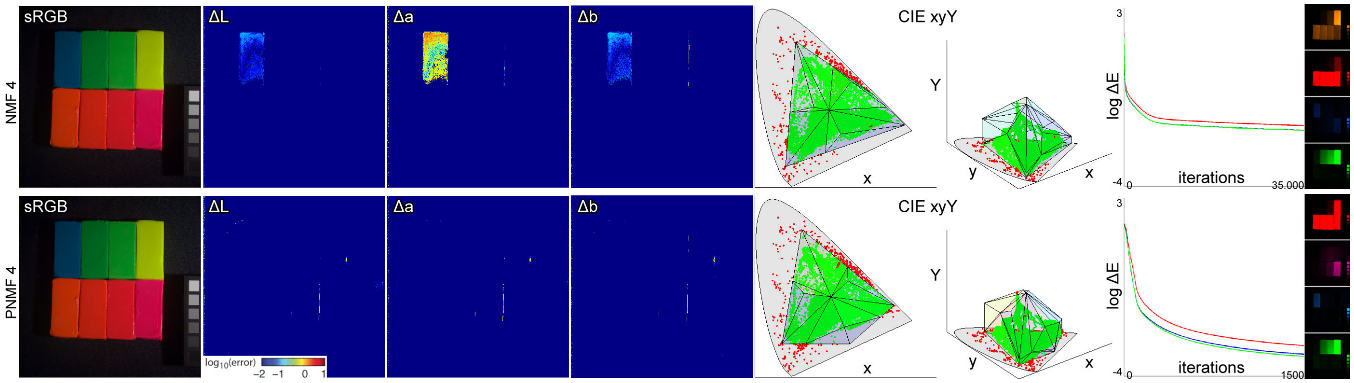

Figure S.14: *"Clay" scene for 4 primaries.*

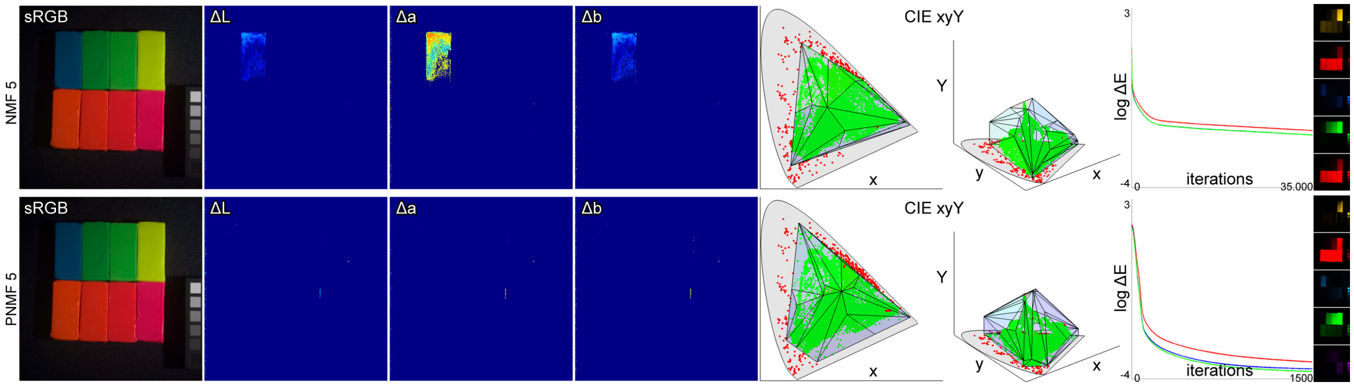

Figure S.15: *"Clay" scene for 5 primaries.*

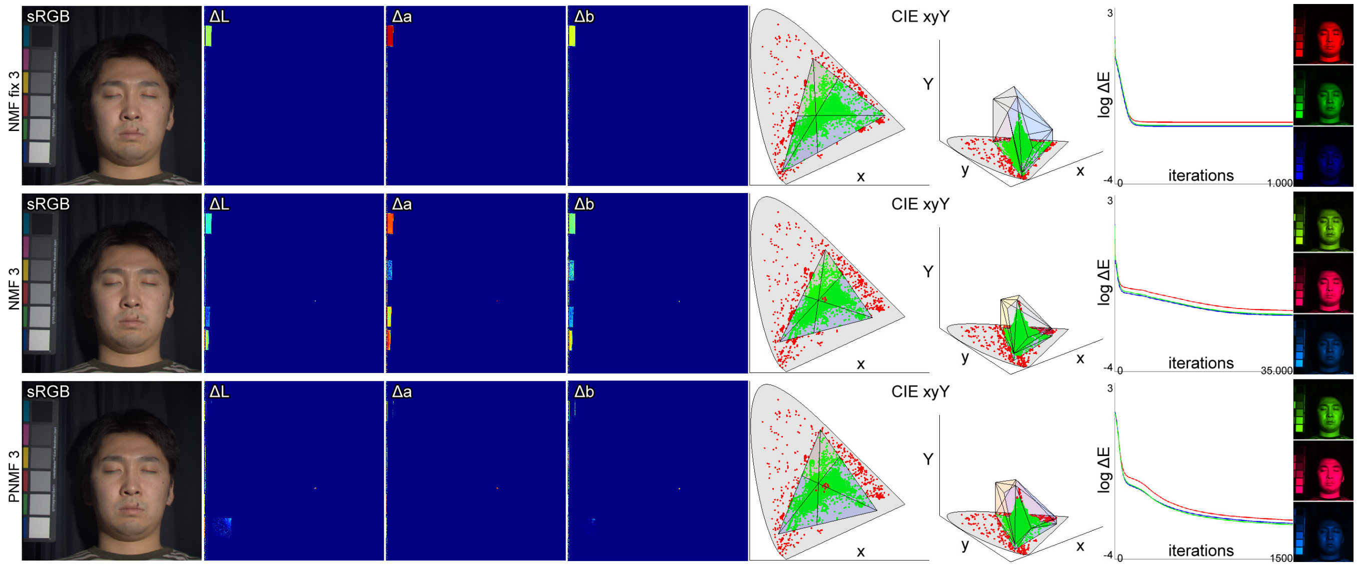

Figure S.16: *"Face" scene for 3 primaries.*

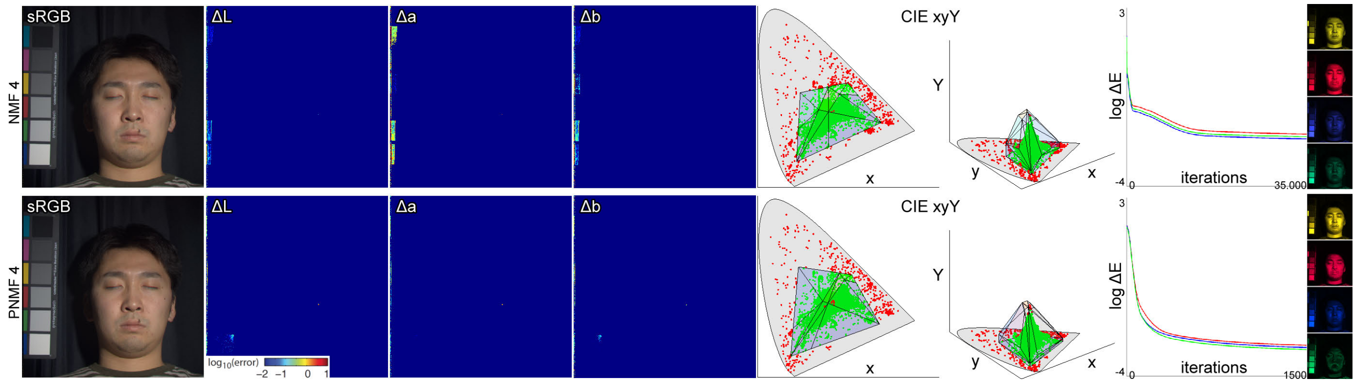

Figure S.17: *"Face" scene for 4 primaries.*

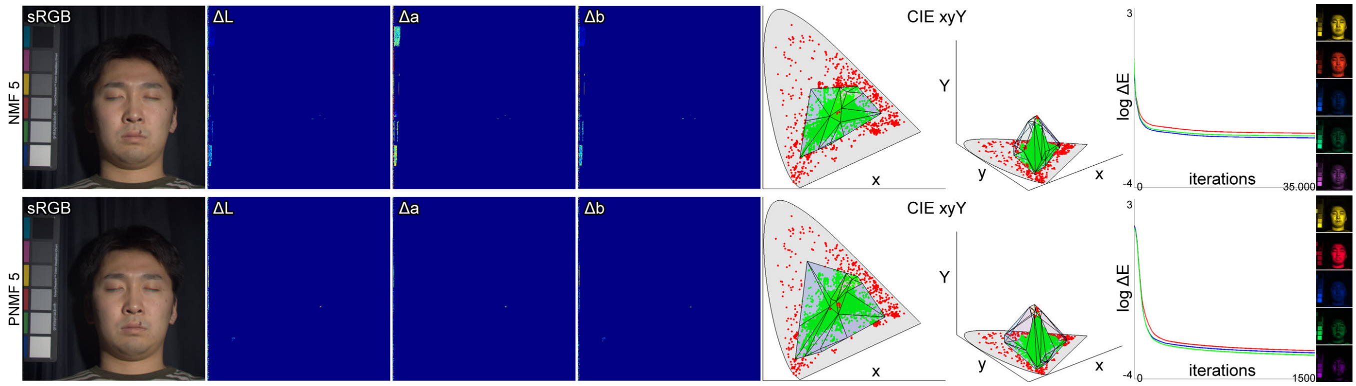

Figure S.18: *"Face" scene for 5 primaries.*

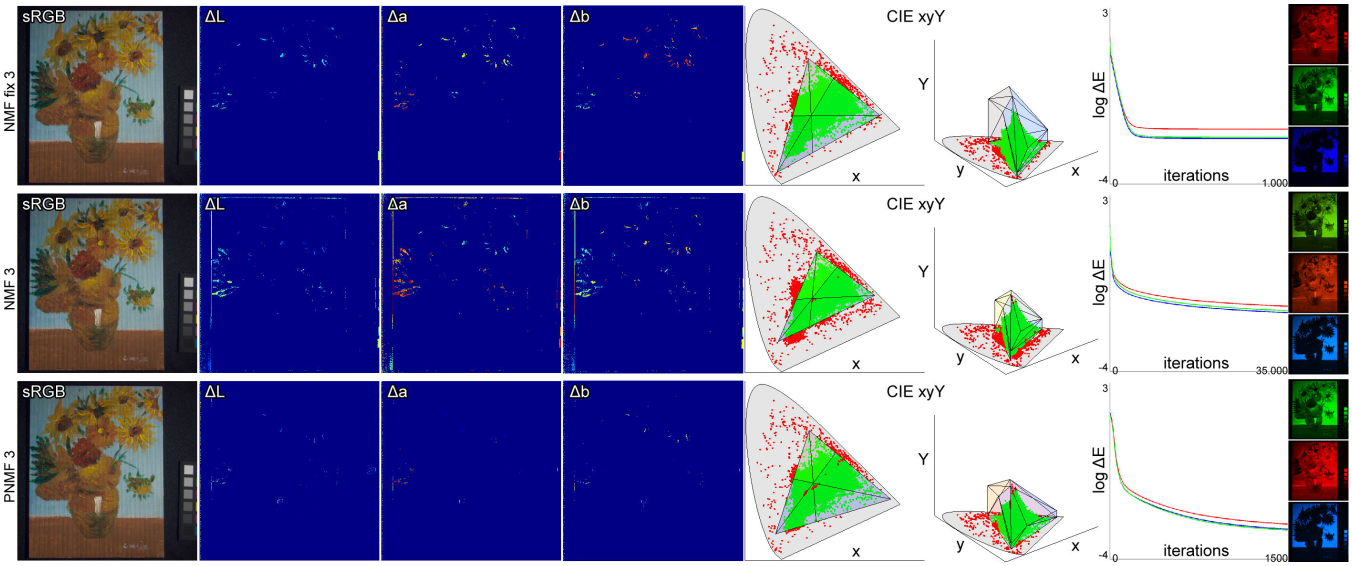

Figure S.19: *"Painting" scene for 3 primaries.*

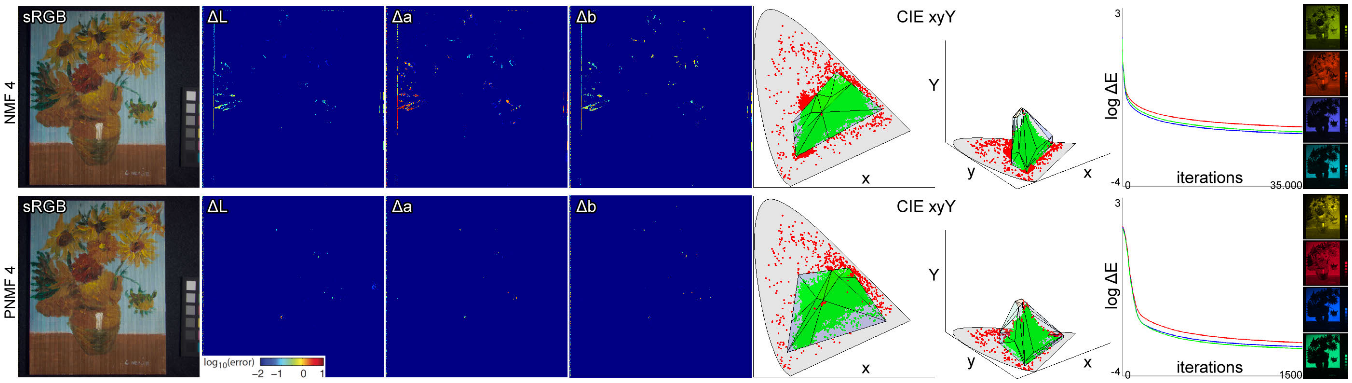

Figure S.20: *"Painting" scene for 4 primaries.*

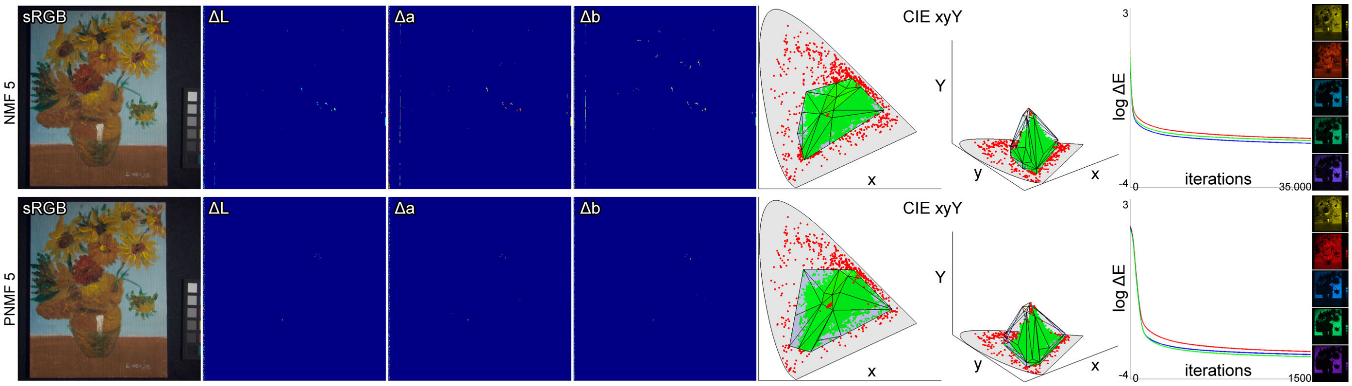

Figure S.21: *"Painting" scene for 5 primaries.*

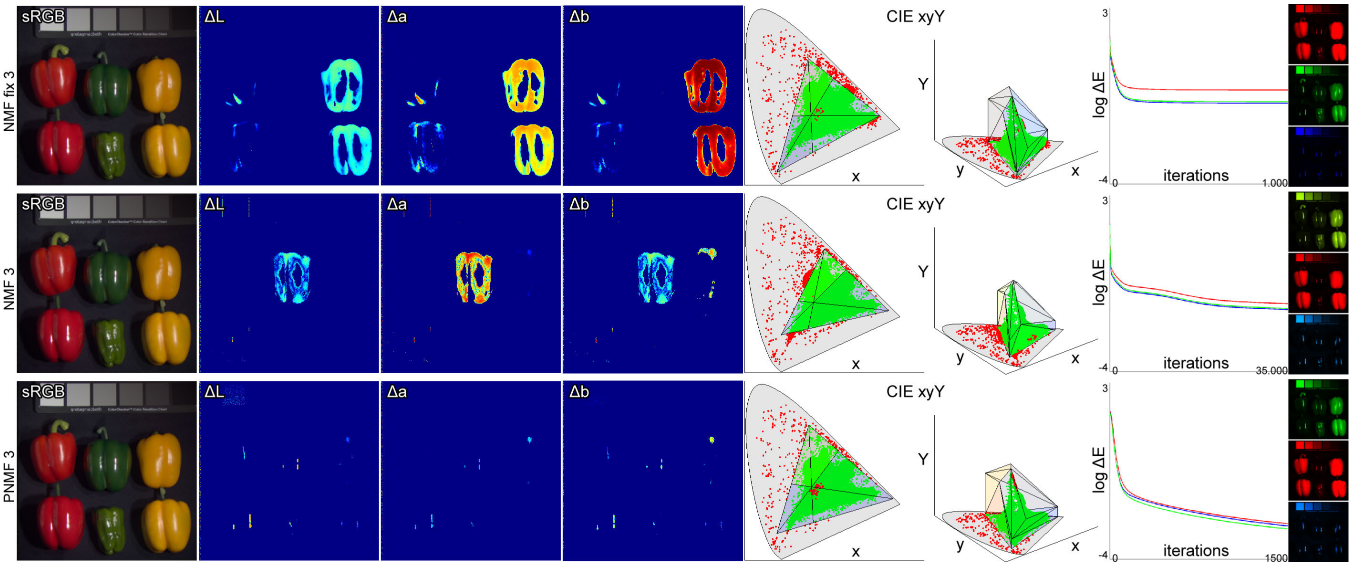

Figure S.22: *"Pepper" scene for 3 primaries.*

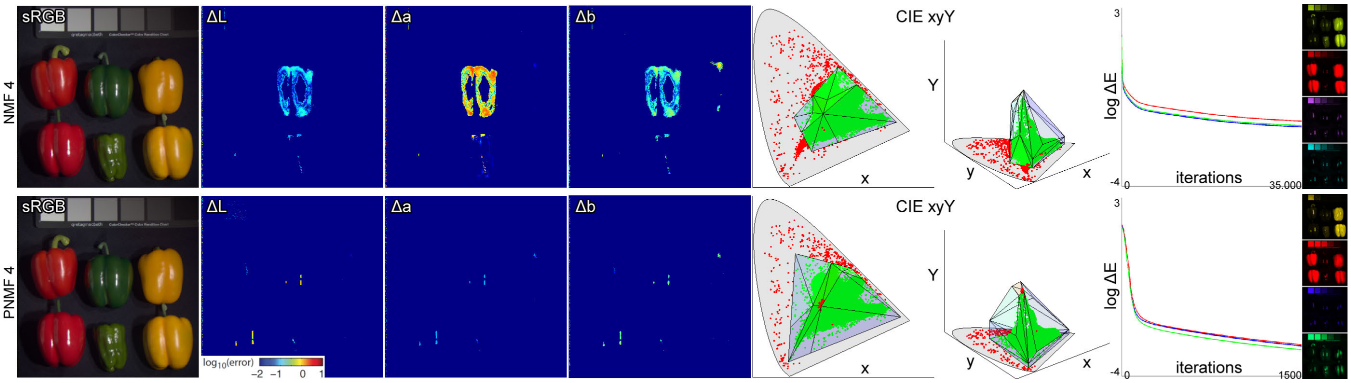

Figure S.23: *"Pepper" scene for 4 primaries.*

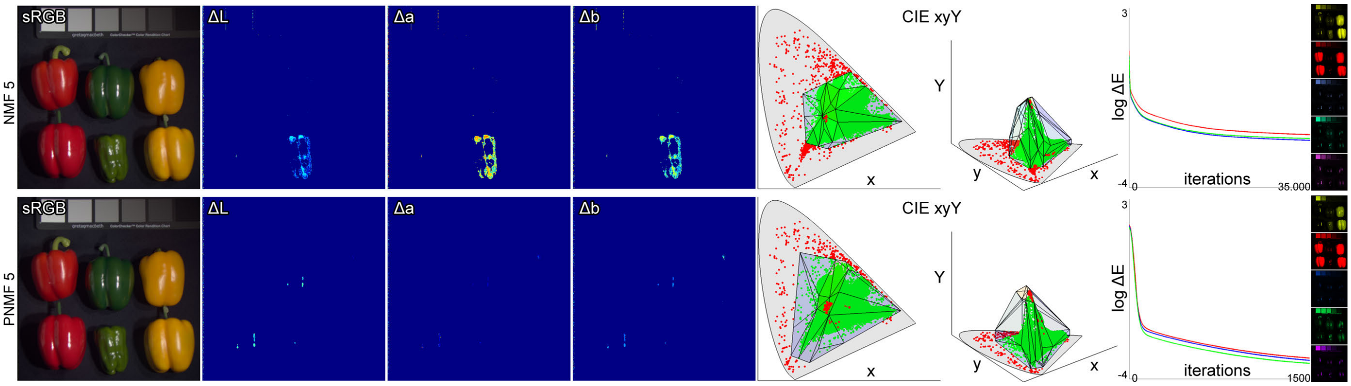

Figure S.24: *"Pepper" scene for 5 primaries.*

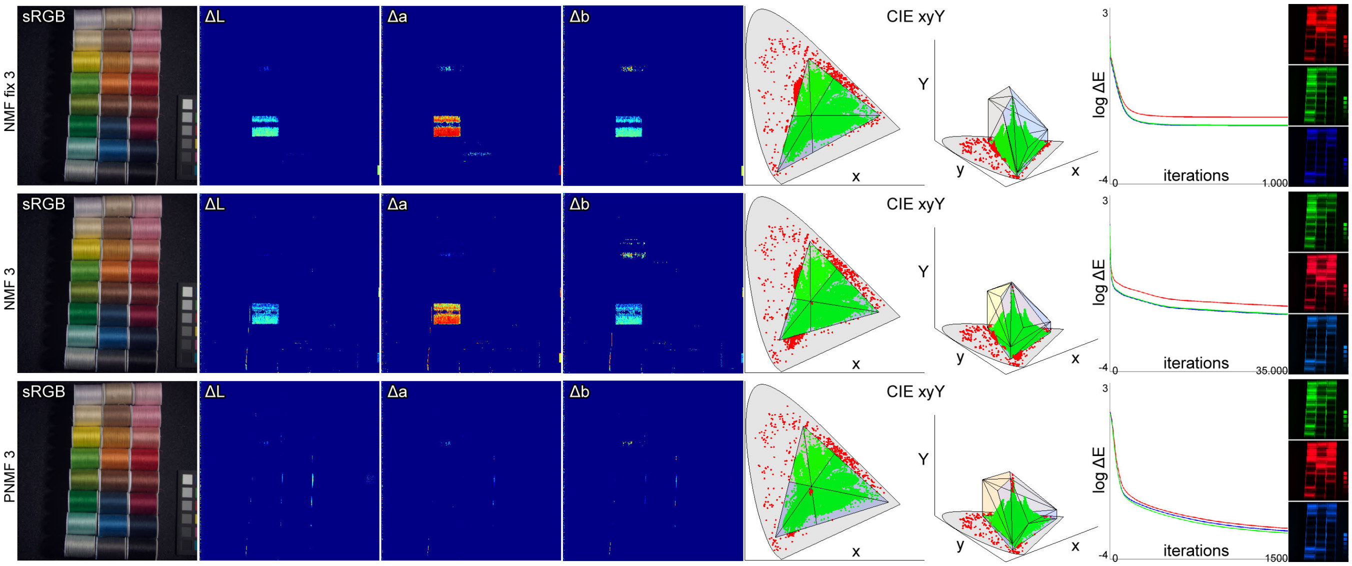

Figure S.25: *"Thread" scene for 3 primaries.*

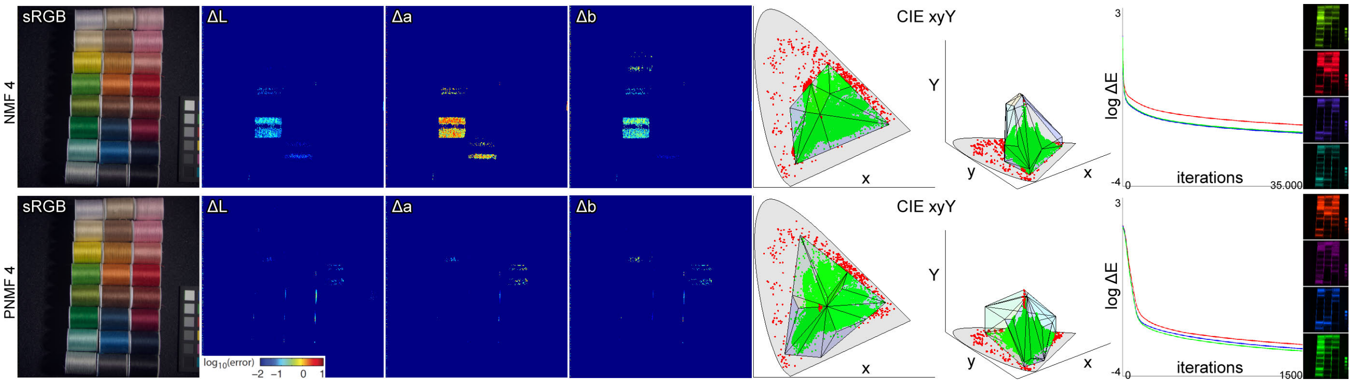

Figure S.26: *"Thread" scene for 4 primaries.*

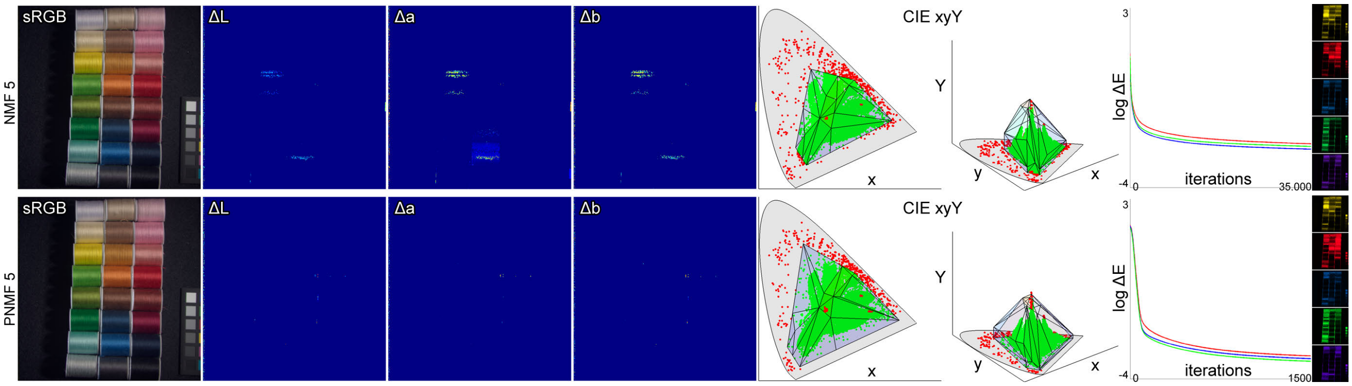

Figure S.27: *"Thread" scene for 5 primaries.*

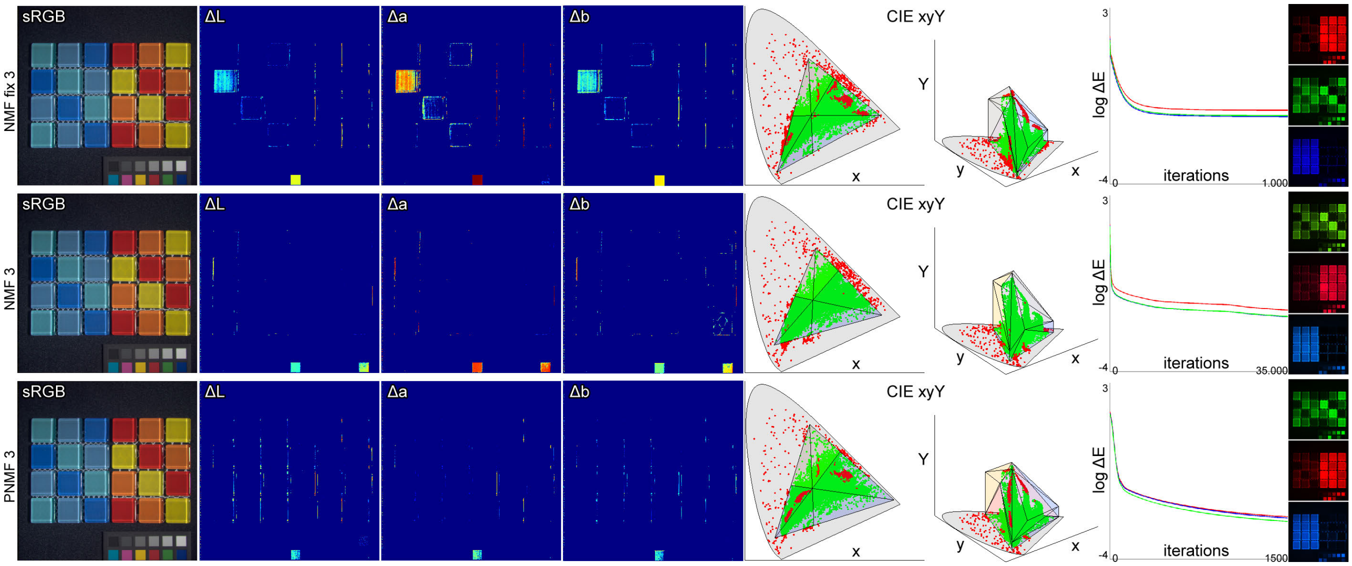

Figure S.28: *"Tile" scene for 3 primaries.*

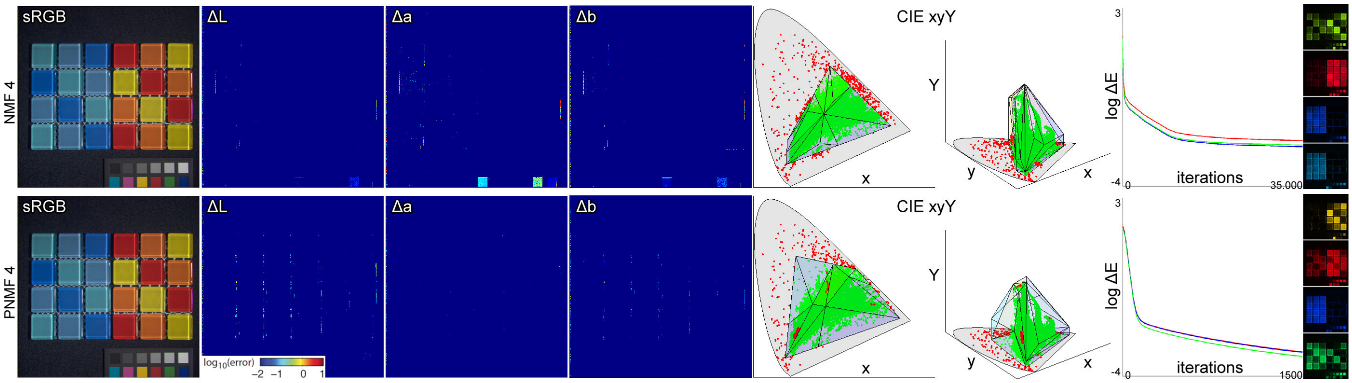

Figure S.29: *"Tile" scene for 4 primaries.*

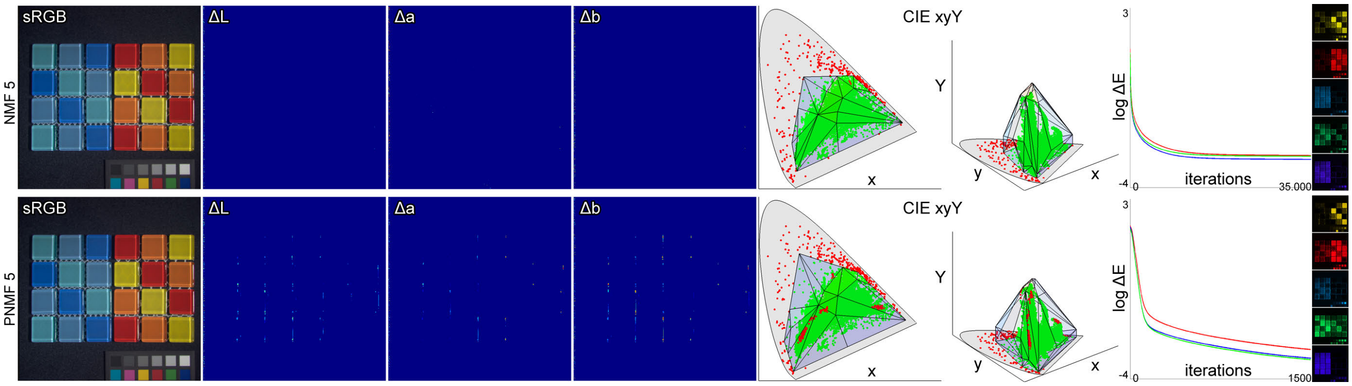

Figure S.30: *"Tile" scene for 5 primaries.*

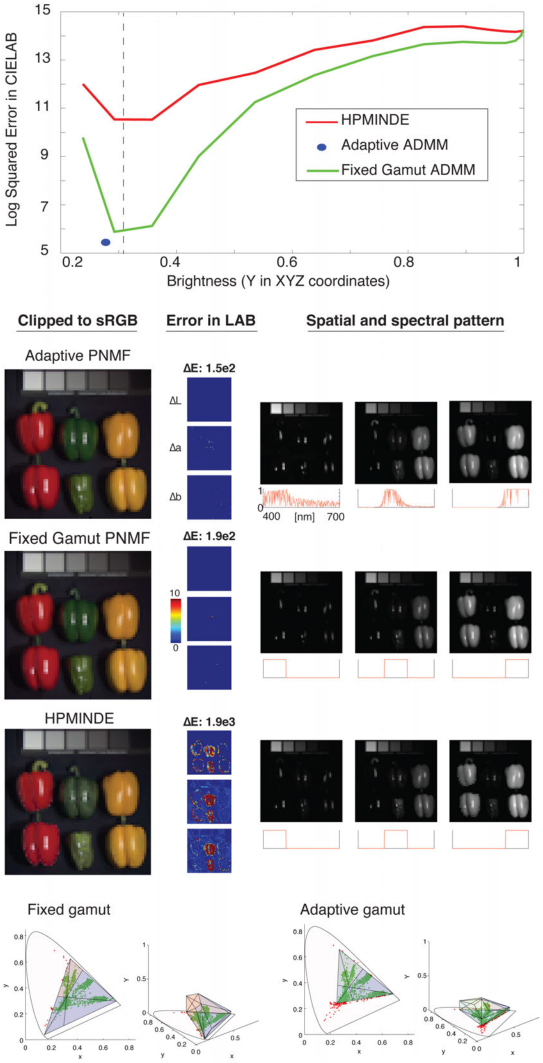

Figure S.31: ΔE_{76} *comparison between gamut mapping algorithms. The topmost figure shows that the minimum* ΔE_{76} *for both fixed gamut algorithms occurred with a gamut brightness (Y value in XYZ coordinates) of around 0.3, while the adaptive gamut performed* better *and yielded a gamut with similar brightness. The middle three rows show the actual images, approximated for print using sRGB primaries; the errors in CIE Lab space when compared to the original multispectral image; as well as the spectral filters and their associated spatial pixel intensity values. The sRGB representation of the HPMINDE result is visibly worse when looking at the green pepper in the middle of the bottom row of the image. The bottom figures display in xyY space the fixed gamut and the final adapted gamut, overlaid on the original multispectral image. Red points are outside of the gamut, but mostly dark.*

B Motivation For Using 3-4 Color Primaries

We review the three basic approaches to displaying color images: temporal multiplexing, optical overlay of multiple devices, and color filter array-type displays. Temporal multiplexing, which is known as “field sequential color”, presents images with different color primaries at high refresh rates. Examples include digital light processing (DLP) projectors that employ high-speed SLMs and a spinning wheel with different color filters that optically modulate the spectrum of the projected images. Optical overlay of multiple SLMs within the same device enclosure is common for liquid crystal display (LCD) projectors; the output of multiple projectors with different color filters could also be overlaid on the screen. Finally, conventional LCDs usually employ spatial multiplexing by interleaving different

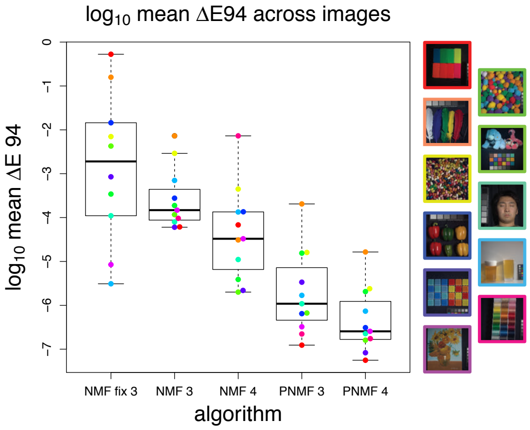

Figure S.32: *PNMF yields lower ΔE than NMF and fixed gamut NMF. Summary of log-transformed mean ΔE across $n = 11$ multispectral images, for fixed gamut NMF with 3 primaries, flexible gamut NMF with 3 and 4 primaries, and flexible gamut PNMF with 3 and 4 primaries. PNMF 3 has significantly lower ΔE_{94} than NMF 3, and PNMF 4 is trending to lower ΔE_{94} than NMF 4 (p = 0.0039 and p = 0.073, paired t-test with Welch correction for unequal variances); PNMF 3 has significantly lower ΔE_{2000} than NMF 3, and PNMF 4 has significantly lower ΔE_{2000} than NMF 4 (p = 0.002 and p = 0.003, respectively, paired t-test with Welch correction). Additionally, PNMF 3 has significantly lower $\Delta E_{sCIELAB}$ than NMF 3, and PNMF 4 has significantly lower $\Delta E_{sCIELAB}$ than NMF 4 (p = 0.024 and p = 0.013, respectively, paired t-test with Welch correction). Colored dots in the boxplots correspond to the images represented in the bottom right.*

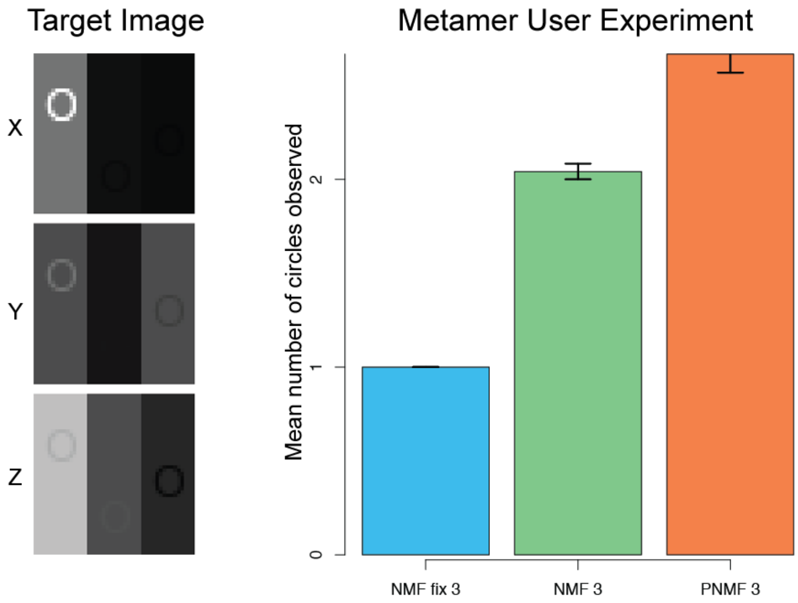

Figure S.33: *Design and results of a user study demonstrating how PNMF differentiates between colors that are mapped to metamers or near-metamers when using NMF or NMF with a fixed gamut. On the left, the target image is displayed in CIEXYZ space. In the middle, the top two rows show two views of the gamut resulting from each factorization, overlaid with the gamut of the target image. The bottom row shows the gamut overlaid with the gamut-mapped image. On the right are the results of the user study: the PNMF factorization enabled significantly higher detection of all three circles than did the NMF factorization ($p < 0.0001$, paired t-test with Welch correction), while the NMF with fixed gamut factorization yielded metamers for two of the color pairs.*

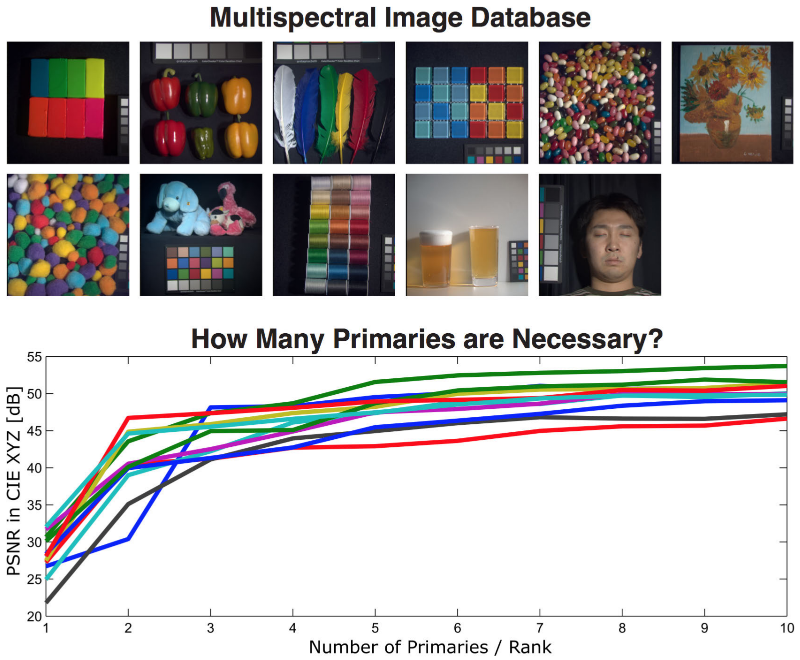

Figure S.34: *Factorization of different hyperspectral images into a set color primaries and corresponding pixel states. This experiments confirms that natural images are usually well-approximated with only about 3-4 color primaries, therefore they are low-rank in the spatio-spectral domain.*

color filters on the same screen. Disregarding the actual technology, all of these displays rely on a simple assumption: high-quality images can be synthesized using only a few color primaries.

Fortunately, natural images are well-known to be smooth. It is generally argued that about 3-4 color primaries are sufficient to represent natural scenes with high fidelity [Dannemiller 1992; Chiao et al. 2000]. We confirm these findings in a simple experiment shown in Figure S.34. Using the factorization method described in Section 3, eleven hyperspectral images[1] are factorized into decompositions with varying rank, which in this application is equivalent to the number of color primaries. As predicted in the literature, all of these images are very well approximated by a few primaries. However, a question that not discussed very often is whether or not the *same (few) primaries* are adequate to represent a wide range of different natural images well.

C Color Difference Metrics

In this section, we review several color difference metrics that are relevant for adaptive color display. The metric dominantly discussed in the primary text is ΔE_{76}, which is based in a sum-of-squared differences in the conventional CIE76 space. we briefly review this and then discuss more sophisticated color difference metrics that are used in the primary text.

C.1 CIE76

The 2-norm of the color difference computed on the conventional CIE Lab (see primary text, Section 3) is usually referred to as ΔE_{76}:

$$\Delta E_{76} = \sqrt{(L_1 - L_2)^2 + (a_1 - a_2)^2 + (b_1 - b_2)^2}, \tag{S.1}$$

[1] cs.columbia.edu/CAVE/databases/multispectral/

where L_1, a_1, b_1 and L_2, a_2, b_2 are the CIE Lab values for two images. The objective function used in the primary text uses the sum-of-squared ΔE_{76} values over all the pixels in an input images as the residual, but we usually plot the average ΔE_{76} in all convergence plots. ΔE_{76} is intuitive, but also somewhat deprecated and mostly replaced by CIE94 and CIEDE2000.

C.2 CIE94

In 1995, the CIE published a recommended practice CIE 116-1995 for industrial color difference evaluation. It includes a new color difference evaluation model designated as CIE94, with symbol ΔE^*_{94}. In this model, color is represented by CIE lightness(L^*), chroma(C^*), and hue(h^*). Instead of directly calculating the Euclidean distance in CIE LCh space, weights are incorporated to correct for the variation in preceived color difference magnitude in each color component.

$$\Delta E^*_{94} = \sqrt{\left(\frac{\Delta L^*}{k_L S_L}\right)^2 + \left(\frac{\Delta C^*_{ab}}{k_C S_C}\right)^2 + \left(\frac{\Delta H^*_{ab}}{k_H S_H}\right)^2} \tag{S.2}$$

Where detailed definition of each term can be found in Wikipedia. In our implementation of PNMF, we establish the objective function with CIE94,

$$\begin{aligned} \underset{G,H}{\text{minimize}} \quad & \left|\left|\Delta E_{94}\left(\beta I^{(lab)}, \varphi(X)\right)\right|\right|_F^2 + \sigma||\Delta_\lambda G||_F^2 \\ \text{subject to} \quad & X - PGH^T = 0 \\ & 0 \le G_{ik}, H_{jk} \le 1, \forall i, j, k \end{aligned} \tag{S.3}$$

The main difference comparing to use CIE76 is in the $\mathbf{X}$-update, where we need to calculate the Jacobian matrix with respect to the new optimization function. Given a single pixel $X_j^{(xyz)}$, its corresponding value in CIE Lab color space $[L^*_{2j}, a^*_{2j}, b^*_{2j}]$ and its reference pixel value $[L^*_{1j}, a^*_{1j}, b^*_{1j}]$. To derive the update formulae, we rewrite the Jacobian matrix when using CIE76 as important intermediate result

$$J^{76}_{X_j} = \begin{bmatrix} j_{11} & j_{12} & j_{13} \\ j_{21} & j_{22} & j_{23} \\ j_{31} & j_{32} & j_{33} \end{bmatrix} = \begin{bmatrix} 0 & 116\phi'(X_j^{(y)}, W_y) & 0 \\ 500\phi'(X_j^{(x)}, W_x) & -500\phi'(X_j^{(y)}, W_y) & 0 \\ 0 & 200\phi'(X_j^{(z)}, W_z) & -200\phi'(X_j^{(z)}, W_z) \end{bmatrix} \tag{S.4}$$

Here, we neglect the last three rows of the original $J^{76}_{X_j}$ since they come from the Lagrange multiplier term and will remain unchanged in $J^{94}_{X_j}$. With $J^{76}_{X_j}$, we can represent the first three rows of $J^{94}_{X_j}$ as

$$J^{94}_{X_j} = \begin{bmatrix} 0 & j_{11} & 0 \\ c_1(a^*_{2j}j_{21} + b^*_{2j}j_{31}) & c_1(a^*_{2j}j_{22} + b^*_{2j}j_{32}) & c_1(a^*_{2j}j_{23} + b^*_{2j}j_{33}) \\ c_2(c_3 j_{21} + c_4 j_{31}) & c_2(c_3 j_{22} + c_4 j_{32}) & c_2(c_3 j_{23} + c_4 j_{33}) \end{bmatrix} \tag{S.5}$$

where

$$\phi'(x, w) = \begin{cases} \dfrac{1}{3w^{\frac{1}{3}}} x^{-2/3} & \text{if } \dfrac{x}{w} > (\dfrac{6}{29})^3 \\ \dfrac{1}{3w}(\dfrac{29}{6})^2 & \text{otherwise} \end{cases} \tag{S.6}$$

$$c_1 = \frac{1}{\left(1 + 0.045\sqrt{a^{*2}_{1j} + b^{*2}_{1j}}\right)\sqrt{a^{*2}_{2j} + b^{*2}_{2j}}} \tag{S.7}$$

$$c_2 = \frac{1}{\left(1+0.015\sqrt{a_{1j}^{*2}+b_{1j}^{*2}}\right)\sqrt{2\sqrt{a_{1j}^{*2}+b_{1j}^{*2}}\sqrt{a_{2j}^{*2}+b_{2j}^{*2}}-2(a_{1j}^{*}a_{2j}^{*}+b_{1j}^{*}b_{2j}^{*})}} \tag{S.8}$$

$$c_3 = \left(a_{1j}^{*} - \frac{a_{2j}^{*}\sqrt{a_{1j}^{*2}+b_{1j}^{*2}}}{\sqrt{a_{2j}^{*2}+b_{2j}^{*2}}}\right) \tag{S.9}$$

$$c_4 = \left(b_{1j}^{*} - \frac{b_{2j}^{*}\sqrt{a_{1j}^{*2}+b_{1j}^{*2}}}{\sqrt{a_{2j}^{*2}+b_{2j}^{*2}}}\right) \tag{S.10}$$

By filling in the last three rows of $J_{X_j}^{94}$, we can apply Gauss-Newton method to iteratively approach the solution for pixel X_j:

$$X_j \leftarrow X_j - \left(\mathbf{J}_{X_j}^{94\,T}\mathbf{J}_{X_j}^{94}\right)^{-1}\mathbf{J}_{X_j}^{94\,T} f\left(X_j\right). \tag{S.11}$$

C.3 CIEDE2000

CIEDE2000 is first proposed by CIE Technical Committee 1-47 in 2001, and later becomes a standard in 2013. Comparing with CIE94, it accommodates 4 major revisions to adequately resolve the perceptual uniformity issue.

In CIE94, for the majority of color centers close to neutral(gray), the color discrimination contours are ellipses with main axis oriented to 90°, not the expected constant-diameter circles. Therefore, CIEDE2000 introduces a new coordinate system that transforms the near-achromatic ellipses into circles by elongating the length of the a^*-coordinate.

In CIE94, the lightness weighting function S_L is neglected, which later proved to be necessary. CIEDE2000 incorporates a V-shape weighting function, of which the minimum occurs at L^* = 50 that is identical to the assumed background lightness. This is consistent with the so-called "crispening effect" ,which states the color difference perception is most sensitive when the background color is close to the color of object sample pair.

For the hue weighting function S_H, a T-function is included in CIEDE2000 to take the hue angle dependence into account which neglected by CIE94. Finally, to cope with the anomalies that the main axes of all ellipses in blue region do not point to the coordinate origin, a rotation term R_T is added to better estimate the difference of blue color pairs.

As a considerably more sophisticated formulae, CIEDE2000 introduces two noticeable source of mathematical discontinuities.

- Mean hue computation
- Hue-difference computation

The mean hue($\overline{H}'$) discontinuity occurs when the absolute difference between two hue angles($|h'_1 - h'_2|$) exceeds $180°$, where an additional $180°$ is added to arithmetic mean. It directly leads to the discontinuity of T and $\Delta\theta$ whose calculation involve the mean hue. Following figures show the volumes of T and $\Delta\theta$ with respect to h'_1 and h'_2. To further manifest the extent of discontinuity, we also show the 2D slice where the maximum discontinuity occurs. As shown in the figures, the discontinuity of T is approximately 60% of the maximum T, and the discontinuity of $\Delta\theta$ is approximately 100% of the maximum $\Delta\theta$.

The hue-difference($\Delta h'$) discontinuity also occurs when the absolute difference between two hue angles($|h'_1 - h'_2|$) exceeds $180°$, where an additional $360°$ is added to the arithmetic hue difference. Since computing $\Delta H'$ calculates $\sin(\Delta h'/2)$, the discontinuity reflects on the sign reversal of $\Delta H'$, and subsequently on the sign of R_T. Notice, there

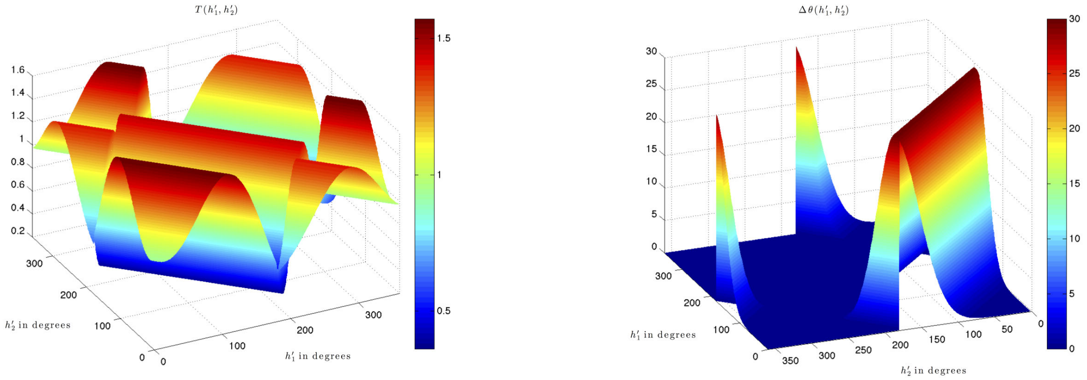

Figure S.35: *Surface plots of* $T(h'_1, h'_2)$ *(left) and for* $\Delta\theta(h'_1, h'_2)$ *(right).*

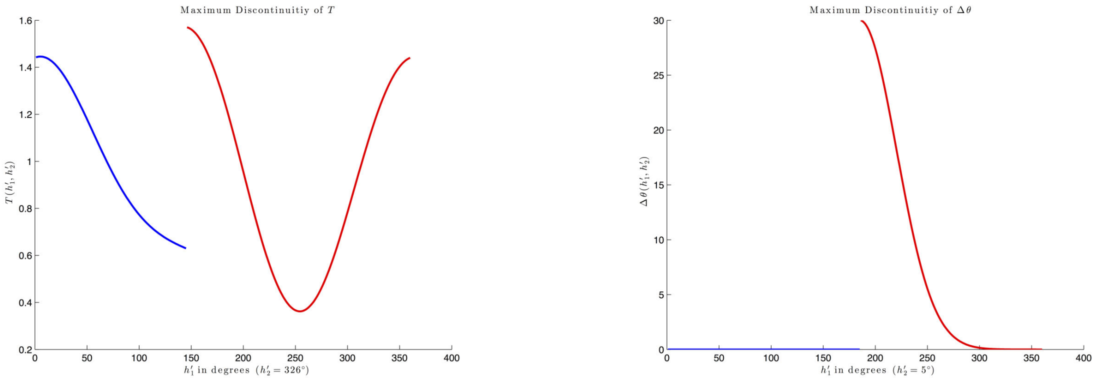

Figure S.36: $T(h'_1, h'_2)$ *at maximum discontinuity (left) and* $\Delta\theta(h'_1, h'_2)$ *at maximum discontinuity (right).*

are two important facts of R_T. First, its discontinuity is simultaneously decided by the mean hue discontinuity and hue-difference discontinuity. Second, when the chroma values of the color pair are same, R_T will equal to 0 due to $\Delta C^* = 0$. Thus, it is hard to quantitatively measure this discontinuity under a meaningful circumstance. However, the impact of hue-difference discontinuity is minor comparing to the mean hue discontinuity.

Overall, we show how the mean hue discontinuity affects the ΔE_{00}. We evaluate ΔE_{00} at $R = 10$ and $R = 60$, where R represents the chroma radius of the color pair. The R_T is zero in both cases, thus eliminates the influence of hue-difference discontinuity.

The discontinuity at $R = 10$ is already noticeable, when R increases to 60 it becomes more apparent. The surface will be more rugged when the chroma radius of the color pair become different, which further introduces the hue-difference discontinuity.

The discontinuities in the CIEDE2000 may not be a major concern in most industrial applications, where other sources of experimental variation are much larger. However, the discontinuities do preclude the use of the formula in analysis based on Taylor series approximations and in design techniques using gradient based optimization, which not only require continuity of the function but also continuity of the first derivative [Sharma et al. 2005]. As a result, our PNMF implementation is based on CIE94 but not CIEDE2000, however, experiments show error of CIEDE2000 is equivalently minimized when we optimize based on CIE94.

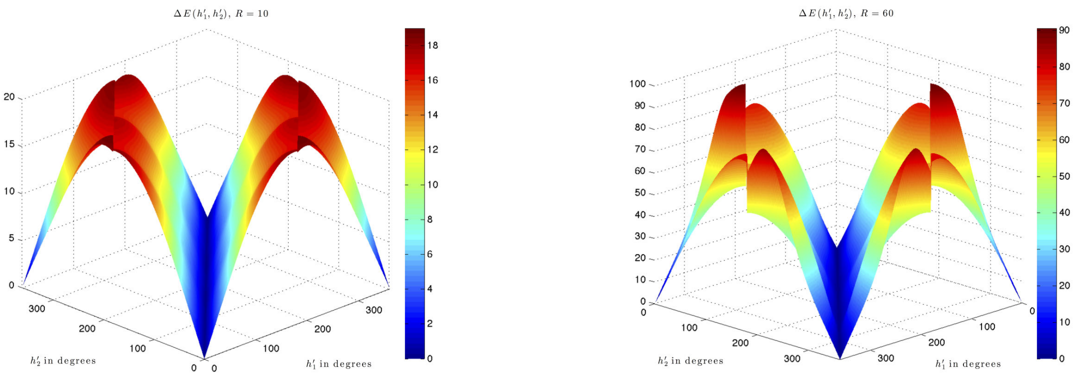

Figure S.37: *Surface plot of* $\Delta E_{00}(h'_1, h'_2), R = 10$ *(left) and* $\Delta E_{00}(h'_1, h'_2), R = 60$ *(right).*

C.4 s-CIE Lab

None of the three color metrics discussed above account for spatial image variation. Effectively, these metrics are valid for comparing the color appearance of two large, uniformly-colored patches. Complex images consist of many many spatial frequencies — the interplay of their color appearances and varying sensitivities to these are not modeled by CIE76, CIE94, or CIEDE2000. Zhang and Wandell [Zhang et al. 1996] introduced s-CIE Lab as a spatial extension of CIE76 that computes the error in scale space to account for varying color sensitivities of these frequency bands. This error metric could easily be optimized for with our algorithm, but it would be computationally much more demanding. Using CIE76 has the advantage of being completely separable into small-scale problems that are solved via Newton's method on a per-pixel basis (see Sec. 3 in the paper). Nevertheless, we evaluate our results using s-CIE Lab in the primary text and demonstrate that our approach consistently drives down the s-CIE Lab error as well.

C.5 SSIM

Traditional similarity measurement methods like peak signal-to-noise ratio (PSNR) and mean squared error (MSE) have been proved to be inconsistent with human eye perception. To address this problem, Wang and Bovik [Wang et al. 2004] introduce the structural similarity (SSIM) index to accurately estimate the preceived errors, which considers image degradation as perceived change in structural information. Mathematically, it computes three terms, namely the illuminance term, the contrast term and the structural term. The overall index is a multiplicative combination of the three terms and will reach 1 when two images are identical. We evaluate our results on the SSIM and demonstrate that our approach will lead SSIM almost equal to 1 across all the dataset.

D Hardware Prototypes

In addition to the LightCrafter prototype described in the primary text, we also built a second prototype that is based imaging spectrometers; we call it *inverse imaging spectrometer (IIS) projector*. The optical design is closely related to recent proposals [Rice et al. 2007, Mohan et al. 2008]. Unfortunately, the color gamut we could achieve with this prototype was limited by the contrast of the employed spatial light modulators. Because it is also much more complex, we decided to outline our efforts on IIS projection only in the supplement while focusing the primary text on the immediately applicable LightCrafter prototype.

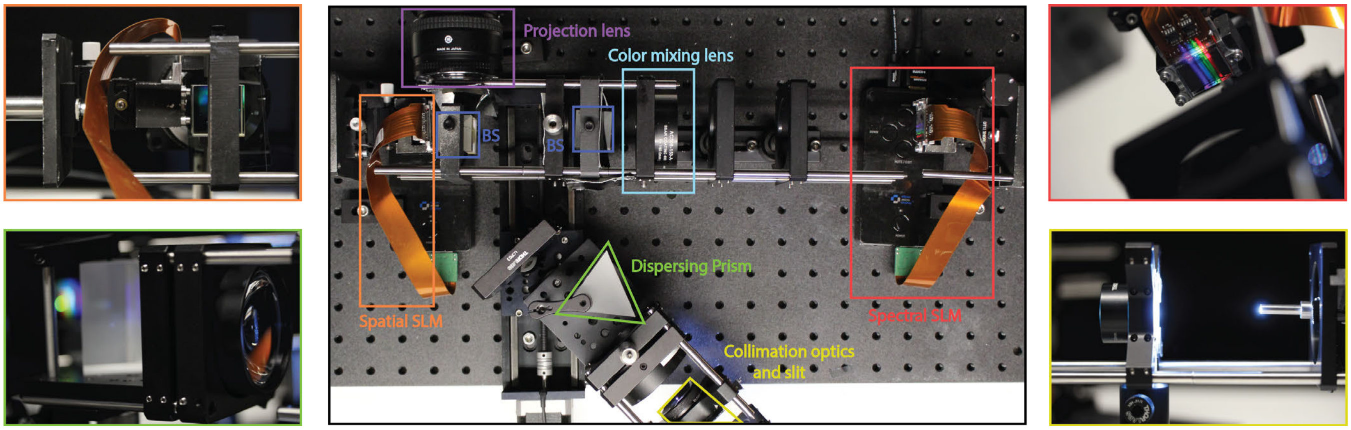

Figure S.38: *Overview of IIS projector prototype. We build an agile spectral projector from scratch, comprising a fiber-coupled broadband light source with collimating optics (lower right), a dispersing prism (lower left), an optical train with magnifying, beam splitting, and relay elements (center) in a Thorlabs 60 mm cage system, and two spatial light modulators (SLMs, SMD ST-1080) that provide a display refresh rate of up to 240 Hz at full high-definition (1080p) resolution. The optical assembly focuses the spectrum of the light source on one of the SLMs (top right), which allows it to be programmatically controlled. The second SLM (top left) is in focus with the projection screen (not shown). Together, these SLMs facilitate high-resolution spatio-spectral image synthesis with the equivalent of up to four primaries per displayed image, assuming a critical flicker fusion of 60 Hz. Note that this projector design allows for arbitrary spectral distributions to be produced, which is not the case for standard multi-primary displays. Hence, this device offers unprecedented degrees of freedom for spectral factorization algorithms.*

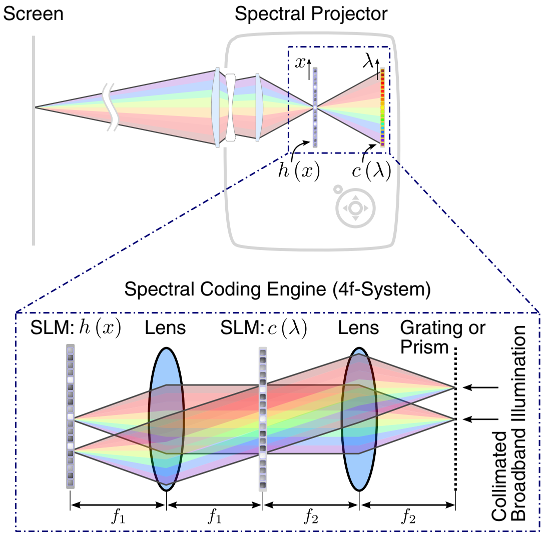

Figure S.39: *Illustration of IIS projection system. Two high-speed spatial light modulators control the illumination inside the device. At any given point in time, one of them shows a spatial image whereas the other adaptively controls the spectrum. Spectral modulation can be optically implemented with a dispersive 4f-system (bottom).*

D.1 IIS Prototype Design

The IIS projector design is exhibited in Figure S.38. We use a liquid light guide-coupled Lumencor Spectra X as the light source, which provided higher power, particularly in the blue end of the spectrum, than the halogen and broadband LED sources we initially tried. The light was collimated using a 45 mm achromatic doublet (Thorlabs), focused onto a slit (cut with a scalpel into a piece of black paper), and then collimated using a 2-inch diameter

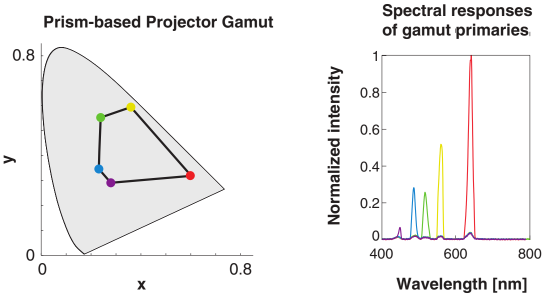

Figure S.40: *Gamut of the IIS IIS projector prototype. In our implementation, the gamut area was limited by the black level of the SLMs, the width of the slit used to collimate light into the prism, and the brightness and spectrum of the light source. A Lumencor Spectra X was use as the light source, as it provided more blue light than other tested sources.*

40 mm focal length aspherical condenser (Thorlabs ACL5040) onto a equilateral dispersing prism (Thorlabs PS854). The prism was placed one focal length away from a 150 mm achromatic doublet (Thorlabs), with a polarizing beamsplitter, and an additional polarizer in between. The spectral SLM (Silicon Microdisplay ST1080 HD liquid crystal on silicon (LCoS) spatial light modulator) was placed one focal length past the 150 mm lens, supported by a structure that allowed a full six degrees of kinematic freedom. The spatial SLM was placed one focal length away on the opposite side of the lens, with more polarizers and a polarizing beam splitter in between. A Nikon 50 mm F-mount lens projected the image that was displayed on the spatial SLM. Images were displayed and synchronized on the SLMs by configuring them as external monitors in MOSAIC display mode, and then using either IrfanView or an OpenGL script in fullscreen mode. For calibration, a fiber-coupled spectrometer (Thorlabs CCS200) was used. For result captures, a Canon T5 with a 50 mm lens directly imaged the output of the projector. The performance of our setup was degraded by a number of factors: the spectral purity and resolution (and thus gamut size) was limited by the width of the slit and by the contrast of the SLMs, the power was limited by our light source and by the fact that we were using polarization based SLMs as opposed to much more efficient digital micromirror devices (DMDs), and the homogeneity of our output could have been improved through the use of an integrating sphere or other homogenizer between the spectral and spatial SLMs.

D.2 IIS Prototype Results

Preliminary results of the IIS prototype are shown in Figure S.41. Note that the error metric here is the sum of squared ΔE_{76} errors of all pixels in the image. The actual number may not be intuitive to a color scientist, but it is the residual optimized by our implementation of the proposed algorithms for this prototype.

E More Detailed ADMM derivation using ΔE_{76}

Here we describe the derivation of an ADMM formulation to solve the following problem. The spectral smoothness term is omitted in this derivation, but the same derivation applies and the final expressions with that term are listed in the primary text.

$$\begin{aligned} \underset{\{\mathbf{G},\mathbf{H}\}}{\text{minimize}} \quad & \left\|\beta \mathbf{I}^{(lab)} - \varphi\left(\mathbf{X}\right)\right\|_F^2 \\ \text{subject to} \quad & \mathbf{X} - \mathbf{P}\mathbf{G}\mathbf{H}^T = 0 \\ & 0 \le \mathbf{G}_{ik}, \mathbf{H}_{jk} \le 1, \quad 0 \le \mathbf{X}_{ij}, \quad \forall i, j, k \end{aligned} \tag{S.12}$$

Addition of a penalty term to the objective leads to formulation of an augmented Lagrangian:

$$L_\rho(\mathbf{X}, \mathbf{G}, \mathbf{H}, \mathbf{Y}) = \left\|\beta \mathbf{I}^{(lab)} - \varphi\left(\mathbf{X}\right)\right\|_F^2 + \sum \left(\mathbf{Y} \circ \left(\mathbf{X} - \mathbf{P}\mathbf{G}\mathbf{H}^T\right)\right) + \left(\frac{\rho}{2}\right) \left\|\mathbf{X} - \mathbf{P}\mathbf{G}\mathbf{H}^T\right\|_F^2 \tag{S.13}$$

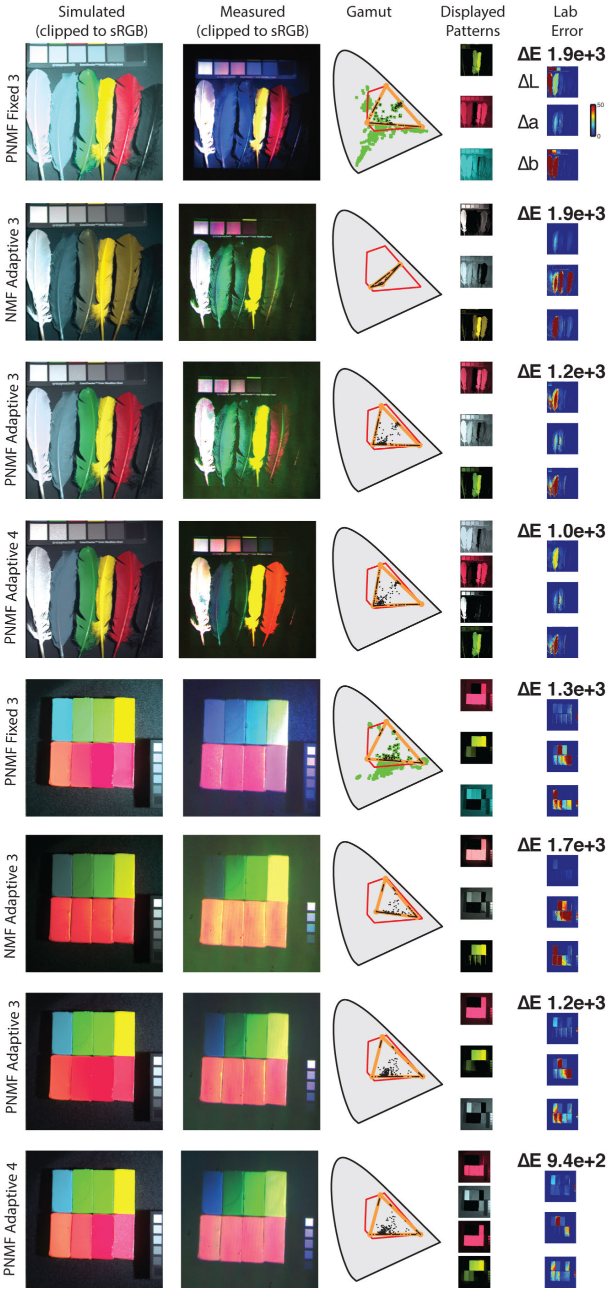

Figure S.41: *Simulation and photographed results for two multispectral datasets on the IIS projector. Note that the image colors have been clipped for printing from the gamut of the projector to that of sRGB. As shown in the left column, in general, ΔE is less for the adaptive gamut PNMF than for NMF and fixed gamut PNMF. Additionally, the photographed results match the simulations to within calibration error. In the center column, the red gamut represents the gamut of the IIS projector, the orange represents the solution gamut, black points represent the solution gamut-mapped pixel chromaticities, and green points represent the target pixel chromaticities.*

We can write the ADMM update rules for this Lagrangian in standard form:

$$\mathbf{X}^{k+1} := \underset{\{0\leq\mathbf{X}\}}{\arg\min}\ L_\rho\left(\mathbf{X},\mathbf{G}^k,\mathbf{H}^k,\mathbf{Y}^k\right) \tag{S.14}$$

$$\left\{\mathbf{G}^{k+1},\mathbf{H}^{k+1}\right\} := \underset{\{0\leq\mathbf{G},\mathbf{H}\leq 1\}}{\arg\min}\ L_\rho\left(\mathbf{X}^{k+1},\mathbf{G},\mathbf{H},\mathbf{Y}^k\right) \tag{S.15}$$

$$\mathbf{Y}^{k+1} := \mathbf{Y}^k + \rho\left(\mathbf{X}^{k+1} - \mathbf{P}\mathbf{G}^{k+1}(\mathbf{H}^{k+1})^T\right) \tag{S.16}$$

Using the scaled form of the augmented Lagrangian, we arrive at the final ADMM update rules by setting $\mathbf{U} = (1/\rho)\mathbf{Y}$:

$$\mathbf{X} \leftarrow \underset{\{0\leq\mathbf{X}\}}{\arg\min} \left\|\beta\mathbf{I}^{(lab)} - \varphi\left(\mathbf{X}\right)\right\|_F^2 + (\rho/2)\left\|\mathbf{X} - \mathbf{PGH}^T + \mathbf{U}\right\|_F^2 \tag{S.17}$$

$$\{\mathbf{G}, \mathbf{H}\} \leftarrow \underset{\{0\leq\mathbf{G},\mathbf{H}\leq 1\}}{\arg\min} \ (\rho/2)\left\|\mathbf{X} - \mathbf{PGH}^T + \mathbf{U}\right\|_F^2 \tag{S.18}$$

$$\mathbf{U} \leftarrow \mathbf{X} - \mathbf{PGH}^T + \mathbf{U} \tag{S.19}$$

The first update step is a constrained nonlinear least squares problem. Importantly, this step can be split and solved in parallel across columns of $\mathbf{X}$, as a per-pixel operation. By deriving the Jacobian of the objective function associated with this step, we can determine $\mathbf{X}^{k+1}$ using an iterative algorithm such as Newton-Raphson, applied in parallel across each pixel.

E.1 Solving PNMF via ADMM

E.1.1 Deriving the Jacobian Matrix

Each column of $\mathbf{I}^{(lab)}$ corresponds to the color values of a single pixel. Let $x \in \mathbb{R}^3$, $i^{(lab)}$, u, and a respectively each be the corresponding column of $\mathbf{X}$, $\mathbf{I}^{(lab)}$, $\mathbf{U}$, and $\mathbf{PGH}^T$ for the current iteration. For a single pixel, the objective function of the first update step is

$$f(x) = \left\|\beta i^{(lab)} - \varphi\left(x\right)\right\|_2^2 + (\rho/2)\left\|x - a + u\right\|_2^2 \tag{S.20}$$

This objective can be expanded into a sum of squares of functions of the three components of x.

$$f(x) = \sum_{j=1,2,3} \left(\beta i_j^{(lab)} - \varphi_j\left(x\right)\right)^2 + (\rho/2)\left(x_j - a_j + u_j\right)^2 \tag{S.21}$$

It can further be represented by a vector of functions, r, such that $f(x) = r^T r$.

$$r(x) = \begin{bmatrix} \beta i_1^{(lab)} - \varphi_1\left(x\right) \\ \beta i_2^{(lab)} - \varphi_2\left(x\right) \\ \beta i_3^{(lab)} - \varphi_3\left(x\right) \\ \sqrt{\rho/2}\left(x_1 - a_1 + u_1\right) \\ \sqrt{\rho/2}\left(x_2 - a_2 + u_2\right) \\ \sqrt{\rho/2}\left(x_3 - a_3 + u_3\right) \end{bmatrix} \tag{S.22}$$

The Jacobian is formed by taking the derivative of each entry of $r(x)$ with respect to each of the three components of x. Note that $\varphi_1, \varphi_2, \varphi_3$ correspond to L^*, a^*, b^*, and x_1, x_2, x_3 correspond to X, Y, Z in Primary Equation 8.

$$J\left(X\right) = \begin{bmatrix} 0 & -116\phi'(Y/W_y) & 0 \\ -500\phi'(X/W_x) & 500\phi'(Y/W_y) & 0 \\ 0 & -200\phi'(Y/W_y) & 200\phi'(Z/W_z) \\ \sqrt{\rho/2} & 0 & 0 \\ 0 & \sqrt{\rho/2} & 0 \\ 0 & 0 & \sqrt{\rho/2} \end{bmatrix} \tag{S.23}$$

$$\phi\left(\frac{x}{w}\right)=\begin{cases}\left(\frac{x}{w}\right)^{1/3} & \text{if } \frac{x}{w}>\left(\frac{6}{29}\right)^{3}\\ \frac{1}{3}\left(\frac{29}{6}\right)^{2}\frac{x}{w}+\frac{4}{29} & \text{otherwise}\end{cases} \tag{S.24}$$

Based on Equation S.24, we find that

$$\phi'\left(\frac{x}{w}\right)=\begin{cases}\frac{1}{3w^{1/3}}x^{-2/3} & \text{if } \frac{x}{w}>\left(\frac{6}{29}\right)^{3}\\ \frac{1}{3w}\left(\frac{29}{6}\right)^{2} & \text{otherwise}\end{cases} \tag{S.25}$$

E.1.2 Solving the Nonlinear Problem using Newton's Method

With the per-pixel Jacobian matrix computed, it is straightforward to apply Newton's method to iteratively approach the solution for pixel X_j using the update rules listed in the primary text:

$$X_j \leftarrow X_j-\left(\mathbf{J}_{X_j}^T\mathbf{J}_{X_j}\right)^{-1}\mathbf{J}_{X_j}^T f\left(X_j\right). \tag{S.26}$$

E.2 Solving PNMF via Alternating Least Squares and Levenberg-Marquardt

Instead of ADMM, we can solve the PNMF problem in many other ways as well. Here, we derive the alternating least squares solution. Again, the objective of PNMF is

$$\begin{aligned}\underset{\{\mathbf{G},\mathbf{H}\}}{\text{minimize}} \quad & \left\|\beta\mathbf{I}^{(lab)}-\varphi\left(\mathbf{P}\mathbf{G}\mathbf{H}^T\right)\right\|_F^2\\ \text{subject to} \quad & 0\le\mathbf{G}_{ik},\mathbf{H}_{jk}\le 1, \quad \forall i,j,k\end{aligned} \tag{S.27}$$

Alternating least squares (ALS) basically alternates between solving for $\mathbf{H}$ with fixed $\mathbf{G}$ and solving for $\mathbf{G}$ with fixed $\mathbf{H}$. This is the same idea as used in most conventional NMF solvers. The pseudo-code for such an update is listed in Algorithm 1.

Algorithm 1 Perceptual Nonnegative Matrix Factorization

```
function $[\mathbf{G},\mathbf{H}] = PNMF$
init $\mathbf{G}=rand(M\times K)$, $\mathbf{H}=rand(N\times K)$
for each ALS iteration
    $\mathbf{H}\leftarrow\underset{\{1\ge\mathbf{H}\ge 0\}}{\text{minimize}}\left\|\beta\mathbf{I}^{(lab)}-\varphi\left(\mathbf{P}\mathbf{G}\mathbf{H}^T\right)\right\|_F^2$
    $\mathbf{G}\leftarrow\underset{\{1\ge\mathbf{G}\ge 0\}}{\text{minimize}}\left\|\beta\mathbf{I}^{(lab)}-\varphi\left(\mathbf{P}\mathbf{G}\mathbf{H}^T\right)\right\|_F^2$
end
```

To solve each substep, a nonlinear optimization routine must be employed — we use Levenberg-Marquard. For this purpose, we first need to compute the Jacobian matrix, for example to solve for $\mathbf{G}\in\mathbb{R}^{C\times K}$ will take the form $\mathbf{J}\in\mathbb{R}^{3N\times CK}$. We will reshape $\mathbf{G}$ so that it is a vector $\mathbf{g}\in\mathbb{R}^{CK}$.

$$J\left(\mathbf{g}\right)=\begin{bmatrix}\partial L_1/\partial g_{11} & \cdots & \partial L_1/\partial g_{1K} & \cdots & \partial L_1/\partial g_{C1} & \cdots & \partial L_1/\partial g_{CK}\\ \partial a_1/\partial g_{11} & \cdots & \partial a_1/\partial g_{1K} & \cdots & \partial a_1/\partial g_{C1} & \cdots & \partial a_1/\partial g_{CK}\\ \partial b_1/\partial g_{11} & \cdots & \partial b_1/\partial g_{1K} & \cdots & \partial b_1/\partial g_{C1} & \cdots & \partial b_1/\partial g_{CK}\\ \vdots & & \vdots & & \vdots & & \vdots\\ \partial L_N/\partial g_{11} & \cdots & \partial L_N/\partial g_{1K} & \cdots & \partial L_N/\partial g_{C1} & \cdots & \partial L_N/\partial g_{CK}\\ \partial a_N/\partial g_{11} & \cdots & \partial a_N/\partial g_{1K} & \cdots & \partial a_N/\partial g_{C1} & \cdots & \partial a_N/\partial g_{CK}\\ \partial b_N/\partial g_{11} & \cdots & \partial b_N/\partial g_{1K} & \cdots & \partial b_N/\partial g_{C1} & \cdots & \partial b_N/\partial g_{CK}\end{bmatrix} \tag{S.28}$$

We can explicitly derive these derivatives by writing out the product

$$\mathbf{PGH}^T = \begin{bmatrix} x_1 & \cdots & x_N \\ y_1 & \cdots & y_N \\ z_1 & \cdots & z_N \end{bmatrix} = \begin{bmatrix} p_{x1} & p_{x2} & \cdots & p_{xC} \\ p_{y1} & p_{y2} & \cdots & p_{yC} \\ p_{z1} & p_{z2} & \cdots & p_{zC} \end{bmatrix} \begin{bmatrix} g_{11} & g_{12} & g_{13} & g_{1K} \\ \vdots & & & \vdots \\ g_{C1} & g_{C2} & g_{C3} & g_{CK} \end{bmatrix} \begin{bmatrix} h_{11} & \cdots & h_{1N} \\ h_{21} & \cdots & h_{2N} \\ h_{31} & \cdots & h_{3N} \\ h_{K1} & \cdots & h_{KN} \end{bmatrix} \tag{S.29}$$

And we can thus write for pixel t

$$x_t = \sum_i^C \sum_j^K p_{xi} g_{ij} h_{jt} \tag{S.30}$$

and

$$y_t = \sum_i^C \sum_j^K p_{yi} g_{ij} h_{jt} \tag{S.31}$$

Thus, we obtain that

$$\partial x_t / \partial g_{ij} = p_{xi} h_{jt} \tag{S.32}$$

And if

$$p_x = \begin{bmatrix} p_{x1} \\ \vdots \\ p_{xC} \end{bmatrix} \tag{S.33}$$

$$h_j = \begin{bmatrix} h_{1j} & \cdots & h_{kj} \end{bmatrix} \tag{S.34}$$

Then for pixel t

$$\partial x_t / \partial g_{1:C,1:K} = repmat(p_x, 1, K). * repmat(h_t, C, 1) \tag{S.35}$$

Finally, we can write out the partial derivative we are looking for:

$$\frac{\partial L_t}{\partial g_{11}} = -116 \frac{\partial \varphi_L(y_t)}{\partial y} \frac{\partial y_t}{\partial g_{11}} = -116 \frac{\partial \varphi_L(y_t)}{\partial y} p_{y1} h_{1t} \tag{S.36}$$

and

$$\frac{\partial a_t}{\partial g_{11}} = -500 (\frac{\partial \varphi_a(x_t)}{\partial x} \frac{\partial x_t}{\partial g_{11}} - \frac{\partial \varphi_a(y_t)}{\partial y} \frac{\partial y_t}{\partial g_{11}}) = -500 \frac{\partial \varphi_a(x_t)}{\partial x} p_{x1} h_{1t} + 500 \frac{\partial \varphi_a(y_t)}{\partial y} p_{y1} h_{1t} \tag{S.37}$$

Given $\mathbf{J} \in \mathbb{R}^{3N \times CK}$, and $f \in \mathbb{R}^{3N}$ our goal is to find $\Delta \in \mathbb{R}^{CK}$ such that $\mathbf{g}^{(n+1)} = \mathbf{g}^{(n)} - \Delta$. We find Δ by solving

$$\mathbf{J}\Delta = f = reshape(I^{(lab)} - \varphi(\mathbf{PGH}')) \tag{S.38}$$

The last equation can then be solved iteratively with either Newton's method or Levenberg-Marquard. We implemented this method in Matlab, constructing the Jacobian matrix for the *global* problem explicitly and then solved tested both Newton and LM updates. This method is extremely slow and did not allow us to process images with a resolution larger than 64×64 pixels. In addition, the convergence is not monotonic. Overall, we don't believe this method to be practical. ADMM exhibits significantly better convergence properties, results in better solutions, it is significantly faster, and also more modular (one per-pixel update step and one global NMF update) so that real-time GPU implementations are in reach.

References

CHIAO, C.-C., CRONIN, T. W., AND OSORIO, D. 2000. Color signals in natural scenes: characteristics of reflectance spectra and effects of natural illuminants. *OSA JOSA A 17*, 2, 218–224.

DANNEMILLER, J. L. 1992. Spectral reflectance of natural objects: how many basis functions are necessary? *OSA JOSA A 9*, 4, 507–515.

SHARMA, G., WU, W., AND DALAL, E. N. 2005. The ciede2000 color-difference formula: Implementation notes, supplementary test data, and mathematical observations. *Color Research and Application 30*, 1.

WANG, Z., BOVIK, A. C., SHEIKH, H. R., AND SIMONCELLI, E. P. 2004. Image quality assessment: from error visibility to structural similarity. *Image Processing, IEEE Transactions on 13*, 4, 600–612.

ZHANG, X., WANDELL, B. A., AND W, B. A., 1996. A spatial extension of cielab for digital color image reproduction.