

Adaptive Color Display via Perceptually-driven Factored Spectral Projection

Isaac Kauvar¹ Samuel J. Yang¹ Liang Shi¹ Ian McDowall² Gordon Wetzstein¹
¹Stanford University ²Intuitive Surgical/Fakespace Labs

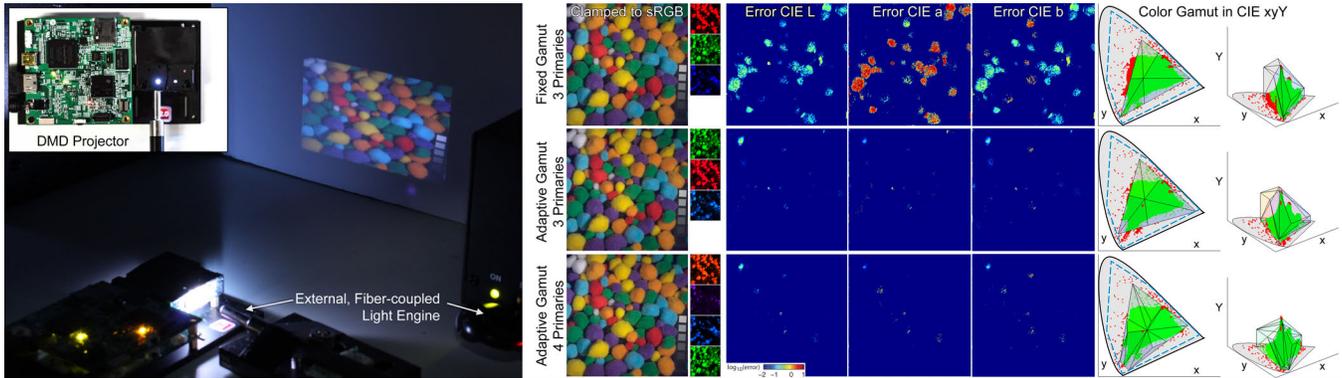


Figure 1: Adaptive color display with our prototype projector (left). The device is a Texas Instruments LightCrafter with its original light engine (three LEDs) replaced by an external, fiber-coupled light engine with six LEDs. The total addressable gamut of the display spans most of the CIE xy space, but the refresh rate of the device only allows for up to four primaries to be used for any image. We develop a perceptually-driven algorithmic framework that computes optimal content-adaptive primaries (each constructed as a linear combination of the LEDs) and corresponding pixel values, as measured in a perceptually uniform space (e.g. CIE Lab). An example hyperspectral image is shown processed for a fixed sRGB gamut (top row) and with our adaptive display system for three (center row) and four (bottom row) primaries. For each case, we show the output image clamped to sRGB along with the color-coded image channels (center left), the error maps in CIE Lab space (center columns), and the resulting gamuts with original colors that are within the optimized gamut in green and those that needed to be gamut mapped in red (right). The total addressable gamut of the light engine, of which the optimal gamuts are a subset, is indicated by a blue dashed line.

Abstract

Fundamental display characteristics are constantly being improved, especially resolution, dynamic range, and color reproduction. However, whereas high resolution and high-dynamic range displays have matured as a technology, it remains largely unclear how to extend the color gamut of a display without either sacrificing light throughput or making other tradeoffs. In this paper, we advocate for adaptive color display; with hardware implementations that allow for color primaries to be dynamically chosen, an optimal gamut and corresponding pixel states can be computed in a content-adaptive and user-centric manner. We build a flexible gamut projector and develop a perceptually-driven optimization framework that robustly factors a wide color gamut target image into a set of time-multiplexed primaries and corresponding pixel values. We demonstrate that adaptive primary selection has many benefits over fixed gamut selection and show that our algorithm for joint primary selection and gamut mapping performs better than existing methods. Finally, we evaluate the proposed computational display system extensively in simulation and, via photographs and user experiments, with a prototype adaptive color projector.

CR Categories: B.4.2 [INPUT/OUTPUT AND DATA COMMUNICATIONS]: Input/Output Devices—Image display; I.3.3 [COMPUTER GRAPHICS]: Picture/Image Generation—Display Algorithms

Keywords: computational displays, computational illumination

1 Introduction

Display resolution standards are moving from ultra high definition to 8K and beyond within the next few years. It is widely antici-

pated, however, that a higher resolution alone will not significantly enhance user experiences unless the dynamic range and color gamut of consumer devices are also improved. High dynamic range display is reliably achieved via dual modulation [Seetzen et al. 2004] and is, after a decade of research and development, entering the market. One of the biggest remaining challenges in display design is that of improving the color gamut while maintaining high light throughput and a sufficient bit depth to encode high-quality images.

Whereas conventional display design strategies determine the tradeoffs between color fidelity, brightness, resolution, bit depth, and other characteristics of a device before it is actually being fabricated and shipped, we advocate for a fundamentally different strategy: *adaptive color display*. Through the co-design of display optics, electronics, and algorithmic processing, we demonstrate how optimal tradeoffs can be made dynamically in a *content adaptive* and *user centric* manner. For example, the spectral power distributions observed in many natural images often do not contain all perceivable colors at once – adapting the display gamut to a specific target image or video clip allows for an optimal tradeoff between brightness and color fidelity to be made. Further, some wide gamut footage may be impossible to be displayed accurately, as it may contain a wide range of colors distributed over the entire perceivable color space. In such cases, the perceptually closest approximation of the target should be presented. Human color vision, however, is a complex and nonlinear process; finding a perceptually optimal solution for the color reproduction problem is therefore challenging. Figures 1 and 2 demonstrate the concept of the proposed approach to adaptive color reproduction for different hyperspectral images processed for a three- and a four-primary display with fixed and adaptive gamuts.

With this work, we propose a computational display system that facilitates adaptive color display. The core contribution of this paper

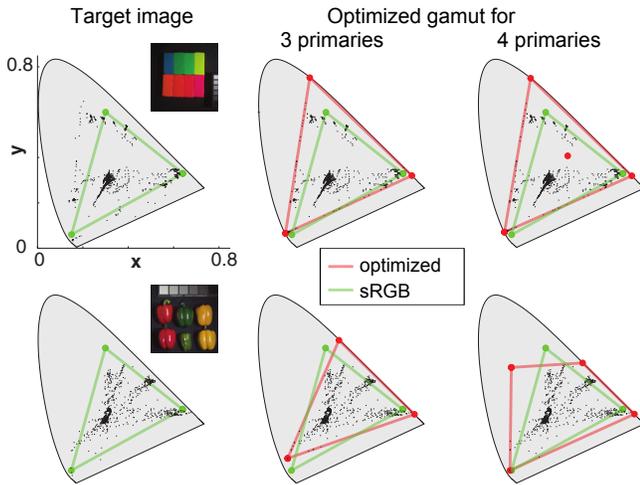


Figure 2: *Optimal color reproduction is content-dependent. For two multispectral images (top and bottom rows), the perceptually optimized three- and four-primary gamuts are computed with our algorithm (red traces), and compared with sRGB (green trace). The optimized gamuts are adaptively defined using primaries that are each a linear combination of at least six LED sources, and are slightly different for each image and also for each display. Notably, in the four-primary case, while the bottom image requires 4 distinct primaries, the top image is well-approximated by 3 primaries and the 4th primary is automatically chosen as a bright color.*

is a *perceptually-driven factorization algorithm* that decomposes a target wide-gamut image into a set of adaptive primaries and corresponding pixel values. The classical problems of primary selection and gamut mapping are thus solved simultaneously and they are robustly optimized in a perceptually uniform color space. We evaluate this algorithm extensively with a custom-built flexible-primary projector. The projector is a modified, low-cost digital light processing (DLP) device offered by Texas Instruments (see Fig. 1). We replace its three light emitting diodes (LEDs) with an external, fiber-coupled light engine providing six distinct primaries covering most of the perceivable colors. As other display systems utilize time-division-multiplexing of the multiple primaries with a single spatial light modulator, our projector is limited to pixel switching speeds of 180-240 grayscale images (8-bit) per second, so given a 60 Hz human flicker fusion threshold, only 3-4 primaries can be supported for any target image. The proposed primary selection algorithm does *not* simply select a subset of the available LEDs but it computes the perceptually-optimal choice of primaries, each of which is a different linear combination of all available LEDs. It is important to note that this choice of optimized primaries results in significantly better color reproduction than simply selecting a few of the LEDs (see Section 6).

One of the main benefits of the proposed display system is its flexibility. Several different display modes are supported by the same device—without mechanically moving parts—simply by switching the software driver. For the proposed display, one could imagine a mode that supports very bright monochrome images, for example text or technical slides, whereas another display mode would support extremely high color fidelity at a lower peak brightness using our algorithm; optimal tradeoffs between brightness and color fidelity are dynamically made by the software. When presenting legacy sRGB content, the device is operated in a conventional three-primary mode without any adjustments. Therefore, the display hardware is flexible enough to support existing content without any tradeoffs while also supporting emerging color spaces.

Overview of Limitations Multi-primary displays have been explored in previous work (see Sec. 2); we built a prototype projector that is based on previously reported optical designs to evaluate our algorithm. The prototype is currently limited to a maximum of 3-4 adaptive primaries via field-sequential color display. The digital micromirror device (DMD) allows for several bit planes to be presented at quick succession that together form an image. Streaming, and therefore video, is not supported by the DMD. With this limitation in mind, we did not attempt to optimize the runtime of the algorithm’s implementation. Currently, our unoptimized Matlab implementation takes a few minutes (for lower resolutions) up to a few hours (for higher resolution) to process each target frame but there is strong evidence that it could be implemented in real time. Finally, color difference metrics are complex, nonlinear, and oftentimes not differentiable, making it difficult to optimize color reproduction for some metrics. We demonstrate robust minimizations of the ΔE_{76} and ΔE_{94} metrics [Witt 2007] but found that CIEDE2000 [CIE 2001], one of the most accurate metrics, is unfortunately too discontinuous to be robustly optimized. Nevertheless, we show that images optimized for ΔE inherently also minimize errors measured with CIEDE2000 and other metrics that also take spatial image variation into account. The latter would be the most appropriate metrics for the discussed application, but also severely increase the computational burden compared to ΔE . Eventually, we hope that our work stimulates follow-on research in the color science community to develop accurate and computationally efficient color difference metrics that are differentiable and possibly even convex, such that they become suitable for content optimization rather than only for evaluation.

2 Related Work

Computational Displays Over the last decades, projector-camera systems and multi-projector displays have been thoroughly explored [Majumder and Brown 2007]. Today, computational display approaches extend virtually every characteristic of displays, including their dynamic range [Seetzen et al. 2004], resolution [Didyk et al. 2010; Sajadi et al. 2012; Berthouzoz and Fattal 2012; Hirsch et al. 2014; Heide et al. 2014], 3D capabilities [Wetzstein et al. 2011; Wetzstein et al. 2012], and refresh rate [Heide et al. 2014]. A comprehensive review of computational displays can be found in the recent survey by Masia et al. [2013]. Our work builds on recent proposals on factored displays [Pauca et al. 2006; Ben-Chorin and Eliav 2007; Lanman et al. 2010; Wetzstein et al. 2012]. In contrast to these approaches, we explore perceptually-driven factorization algorithms for adaptive color display.

Spectral Displays and Algorithms Spectral displays can roughly be classified as multi-primary displays [Teragawa et al. 2012] and hyperspectral displays [Rice et al. 2007; Mohan et al. 2008]. Multi-primary displays usually aim for a wide color gamut, as perceived by a human observer. Related algorithmic problems include selecting the optimal color primaries [Ben-Chorin and Eliav 2007; Long and Fairchild 2011; Li et al. 2015] as well as gamut mapping (e.g., [Banterle et al. 2011]), where pixels of an image are processed to fit within the fixed gamut provided by a display. Gamut expansion can also help to optimize image presentation with large-gamut displays [Majumder et al. 2010]. Hyperspectral displays have the potential to synthesize more complex spectral power distributions than multi-primary displays. Similar to the latter, applications of hyperspectral displays include extended color gamuts, but in addition these types of devices are also useful for hyperspectral imaging, remote sensing, reflectance estimation, and medical imaging [Rice et al. 2012].

In this paper, we target *content-adaptive color display* with multi-

primary displays. We build a custom, multi-primary projector that can dynamically address a large portion of the CIE xy chromaticity diagram. This design is based on similar devices described in the literature (e.g., [Ajito et al. 2000]) but compact and easily built by modifying off-the-shelf hardware. The proposed perceptually-driven algorithm for joint primary selection and gamut mapping is demonstrated with our prototype but also applicable to other displays. For example, we demonstrate a custom-built hyperspectral projection system that is closely related to the design introduced by Rice et al. [2007] and later used by Mohan et al. [2008] in the supplemental document. Hyperspectral projectors could achieve simultaneous color and spectral match of displayed target content, which may be beneficial for some applications [Li et al. 2015].

Our work builds on decades of research in color science; not all of this work can be cited. However, in the last few years a trend towards adaptive color processing is observable. For example, non-negative matrix factorization (NMF) techniques for spectral display have recently been proposed in CIE XYZ space [Pauca et al. 2006; Ben-Chorin and Eliav 2007]. Similar to our algorithm, these approaches solve the joint problem of primary selection and gamut mapping. However, we show that our technique, dubbed *perceptual nonnegative matrix factorization (PNMF)*, outperforms NMF-based techniques in all common perceptual quality metrics. We are the first to demonstrate both NMF and PNMF with a flexible, multi-primary projector; previous work was only simulated. Rodriguez et al. [2012] recently proposed to determine an optimal gamut that maximizes total gamut volume in CIE Luv space. Here, we argue that there is no single optimal gamut in general, and that the best gamut depends on the content to be displayed (see Fig. 2); we offer an algorithm to compute this gamut adaptively, implement several flexible hardware designs to test the performance of this adaptive algorithm, solve the joint problem of primary selection and gamut mapping, and discuss practical applications.

Compressive Spectral Imaging Over the last few years, several compressive hyperspectral imaging systems have been proposed [Gehm et al. 2007; Wagadarikar et al. 2009; Lin et al. 2014]. The goal of these approaches is to recover a 3D hyperspectral image from a coded 2D projection via *sparsity-constrained optimization*. Compressive or factored displays on the other hand usually employ some form of *low-rank factorization* of high-dimensional visual data into the *non-negative* pixel states of the display. Although both types of devices differ in the employed algorithms and hardware, fundamentally they both exploit compressibility of visual data; compressive displays can be interpreted as the dual of compressive cameras. Content-adaptive multi-primary cameras have also been explored and provide benefits for color reproduction in low-light scenarios [Sajadi et al. 2011].

3 Factored Spectral Projection

3.1 Image Formation

Many projectors, in particular those based on digital light processing (DLP) technology, synthesize color images by multiplexing each channel in time (field sequential color). This image formation can be expressed as an additive superposition:

$$i(x, \lambda) = \sum_{k=1}^K c_k(\lambda) \cdot h_k(x), \quad (1)$$

where i is the emitted spectral power distribution and c_k are the temporally-multiplexed primaries. These can, for example, physically be achieved with spectral filters mounted on a quickly rotating

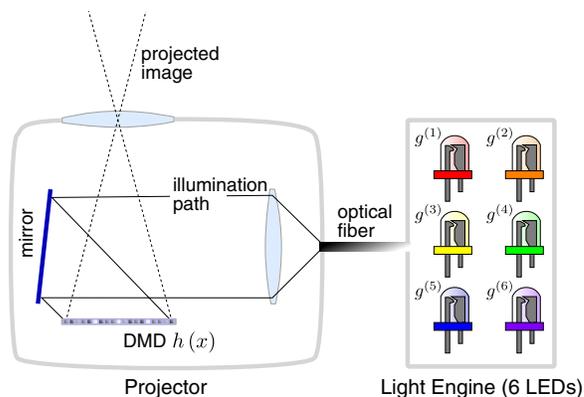


Figure 3: Illustration of multi-primary projection system. A digital micromirror device (DMD) is synchronized with a fiber-coupled light engine. The latter comprises six LEDs that can be individually controlled at high speed.

color wheel in front of a broadband light source. The human visual system averages stimuli presented at refresh rates higher than the critical flicker fusion threshold [Kelly 1979], thus Equation 1 assumes that the refresh rate of the display is K times higher than the critical flicker fusion threshold.

The proposed setup is illustrated in Figure 3 and comprises a standard digital micromirror device (DMD) that rapidly changes pixel states as outlined by Equation 1. Instead of sequentially cycling through the primaries, as done in all conventional time-multiplexed color displays, our light engine allows each color primary to be dynamically chosen as a weighted combination of multiple high-power light emitting diodes (LEDs):

$$c_k(\lambda) = \sum_{m=1}^M g_k^{(m)} \psi^{(m)}(\lambda). \quad (2)$$

Here, the coefficients $g_k^{(m)}$, $m = 1 \dots M$, represent the weights of M individual LEDs at some time k and $\psi^{(m)}(\lambda)$ are the emission spectra of the individual LEDs. Combining Equations 1 and 2 and discretizing them yields

$$\mathbf{I} = \Psi \mathbf{G} \mathbf{H}^T, \quad (3)$$

where the matrix $\Psi \in \mathbb{R}^{L \times M}$ encodes the discrete spectral basis functions in its columns, $\mathbf{G} \in \mathbb{R}^{M \times K}$ are the temporally-varying weights of the LEDs, and $\mathbf{H} \in \mathbb{R}^{N \times K}$ represents the pixel states of a two-dimensional spatial light modulator, such as our DMD, with a total of N pixels. Both the DMD and the LEDs are electronically synchronized and run at a framerate that is K times higher than the critical flicker fusion threshold.

To model the color image perceived by a human observer, we need to take the spectral sensitivity of the receptor cells in the retina into account. Although these slightly vary from person to person, a “standard” observer was characterized by the International Commission on Illumination (CIE) [CIE 1932]. Modeling the spectral sensitivity of the three cone types as $p(\bar{x}, \bar{y}, \bar{z})(\lambda)$, a spectral power distribution is converted into a perceived color image as

$$i^{(x,y,z)}(x) = \int p(\bar{x}, \bar{y}, \bar{z})(\lambda) i(x, \lambda) d\lambda. \quad (4)$$

Here, $i^{(x,y,z)}$ models the spectral power distribution in CIE XYZ color space, representing a color image perceived by the standard

observer. Combining Equations 3 and 4 allows us to write the discretized formation of a perceived color image as

$$\mathbf{I}^{(xyz)} = \mathbf{P}^{(xyz)} \boldsymbol{\Psi} \mathbf{G} \mathbf{H}^T = \mathbf{P} \mathbf{G} \mathbf{H}^T, \quad (5)$$

where the matrix $\mathbf{P}^{(xyz)}$ contains discretized versions of $p(\bar{x}, \bar{y}, \bar{z})(\lambda)$ in its three rows and $\mathbf{P}^{(xyz)} \boldsymbol{\Psi} = \mathbf{P} \in \mathbb{R}^{3 \times M}$ projects a spectral power distribution into CIE XYZ space.

3.2 Factored Image Generation

Similar to [Pauca et al. 2006] and [Ben-Chorin and Eliav 2007], we build on nonnegative matrix factorization to decompose a target image into a set of adaptive color primaries and corresponding pixel values. The former problem is commonly referred to as *primary selection* and the latter as *gamut mapping*. Together, these problems can be formulated as a least-squares problem:

$$\begin{aligned} \arg \min_{\{\mathbf{G}, \mathbf{H}\}} & \left\| \beta \mathbf{I}^{(xyz)} - \mathbf{P} \mathbf{G} \mathbf{H}^T \right\|_2^2 \\ \text{subject to} & \quad 0 \leq \mathbf{G}_{ik}, \mathbf{H}_{jk} \leq 1, \forall i, j, k \end{aligned} \quad (6)$$

Non-negativity constraints enforce physically-feasible pixel states. Equation 6 can be solved iteratively, for example by applying multiplicative update rules [Lee and Seung 1999] in an alternating manner:

$$\mathbf{H} \leftarrow \mathbf{H} \circ \frac{(\mathbf{P} \mathbf{G})^T \beta \mathbf{I}^{(xyz)}}{(\mathbf{P} \mathbf{G})^T (\mathbf{P} \mathbf{G} \mathbf{H}^T) + \epsilon}, \quad \mathbf{G} \leftarrow \mathbf{G} \circ \frac{\mathbf{P}^T \beta \mathbf{I}^{(xyz)} \mathbf{H}}{\mathbf{P}^T (\mathbf{P} \mathbf{G} \mathbf{H}^T) \mathbf{H}^T + \epsilon}. \quad (7)$$

A small value ϵ is often added to the denominator to prevent division by zero. The scalar $\beta \in [0, 1]$ can be useful for digitally trading image brightness for color accuracy. The operators \circ and $//$ denote element-wise multiplication and division, respectively.

Solving the joint problem of primary selection and gamut mapping in CIE XYZ space was proposed in prior work [Pauca et al. 2006; Ben-Chorin and Eliav 2007]. The goal of this paper, on the other hand, is to compute the *closest perceptual approximation* of the target image. Unfortunately, an objective function using a CIE XYZ error metric is ill-formulated to achieve this goal because the Euclidean distances between XYZ coordinates are not perceptually uniform [MacAdam 1942]. A low ℓ_2 error may not be a good perceptual approximation.

4 Towards Perceptually-optimal Factored Spectral Projection

To overcome limitations of previously-proposed spectral factorization methods, as discussed in the previous section, we propose a formulation and practical solutions to solving the factorization problem in a perceptually relevant space. One such color space is CIE Lab, because distances are—at least locally—perceptually linear. This space specifically separates lightness (L) from chromaticity (a,b). A simple nonlinear transformation can be applied to convert from CIE XYZ to CIE Lab:

$$\begin{aligned} L^* &= 116\phi(Y/W_y) - 16 \\ a^* &= 500(\phi(X/W_x) - \phi(Y/W_y)) \\ b^* &= 200(\phi(Y/W_y) - \phi(Z/W_z)) \end{aligned} \quad (8)$$

where $\phi: \mathbb{R}_+ \rightarrow \mathbb{R}_+$ is

$$\phi(x) = \begin{cases} x^{1/3} & \text{if } x > (\frac{6}{29})^3 \\ \frac{1}{3} (\frac{29}{6})^2 x + \frac{4}{29} & \text{otherwise} \end{cases} \quad (9)$$

Following the ICC standard, a reference white of $W_x = 0.9642$, $W_y = 1.0000$, $W_z = 0.8249$ is often used in practice (e.g., for high-definition television), although as described in section 4, we used the calibrated white point of the light engine. An objective function for the spectral factorization problem can be formulated in this perceptually linear space such that a least squared error solution will more closely resemble the best perceptual approximation of a target image. To this end, we write the mapping from CIE XYZ to Lab space as a function $\varphi: \mathbb{R}^{3 \times N} \rightarrow \mathbb{R}^{3 \times N}$ that applies Equations 8 and 9 to each of N image pixels. The objective function is then

$$\begin{aligned} \text{minimize}_{\{\mathbf{G}, \mathbf{H}\}} & \left\| \beta \mathbf{I}^{(lab)} - \varphi(\mathbf{P} \mathbf{G} \mathbf{H}^T) \right\|_F^2 \\ \text{subject to} & \quad 0 \leq \mathbf{G}_{ik}, \mathbf{H}_{jk} \leq 1, \forall i, j, k \end{aligned} \quad (10)$$

We call this *perceptual nonnegative matrix factorization (PNMF)*. PNMF is composed of two subproblems: a bi-convex factorization and the nonlinear term modeling a transformation into the perceptually uniform space.

Note that the squared Frobenius norm in Equation 10 divided by the number of pixels N corresponds to the color difference metric ΔE_{76} averaged over all image pixels. While this is one of the most intuitive, perceptually-motivated color spaces, it is actually neither perfectly perceptually linear nor is it the most appropriate choice for comparing spatially-varying color images. Metrics such as ΔE_{94} and CIEDE2000 model perceptual linearity more accurately and spatial extensions to these spaces (e.g., [Zhang and Wandell 1996]) model the complex interplay of color and spatial image frequencies more appropriately. Unfortunately, CIEDE2000 is too discontinuous to be robustly minimized with any of the evaluated optimization schemes (see Supp. Sec. C.3) and s-CIE Lab would drastically increase computational cost. Therefore, we optimize Equation 10 for ΔE_{94} in practice (see derivation in Suppl. Sec. C.2) but demonstrate in the following sections that this approach inherently also drives down all of the more sophisticated color difference metrics.

An obvious choice for solving Equation 10 would be to apply a nonlinear least squares solver, such as Levenberg-Marquardt (LM). We implemented LM and observe slow and unpredictable convergence in practice (see Supplement). The quality of the solution depends very much on the initial guesses of \mathbf{G} and \mathbf{H} and the runtime is extremely long for high-resolution images. As a practical alternative, we propose to split the problem into multiple subproblems that each can be solved very efficiently. This strategy has recently gained a lot of traction in the optimization community; for example, the alternating direction method of multipliers (ADMM, [Boyd et al. 2011]) is a method that can be applied to solve such problems. In ADMM, Equation 10 is formulated as

$$\begin{aligned} \text{minimize}_{\{\mathbf{G}, \mathbf{H}\}} & \left\| \beta \mathbf{I}^{(lab)} - \varphi(\mathbf{X}) \right\|_F^2 \\ \text{subject to} & \quad \mathbf{X} - \mathbf{P} \mathbf{G} \mathbf{H}^T = \mathbf{0} \\ & \quad 0 \leq \mathbf{G}_{ik}, \mathbf{H}_{jk} \leq 1, \forall i, j, k \end{aligned} \quad (11)$$

where $\mathbf{X} \in \mathbb{R}^{3 \times N}$ is an intermediate variable. Following the general ADMM algorithm design strategy, we derive the Augmented Lagrangian of this problem and derive a sequence of simple update rules as

$$\begin{aligned} \mathbf{X} &\leftarrow \arg \min_{\{\mathbf{X}\}} \left\| \beta \mathbf{I}^{(lab)} - \varphi(\mathbf{X}) \right\|_F^2 + \frac{\rho}{2} \left\| \mathbf{X} - \mathbf{P} \mathbf{G} \mathbf{H}^T + \mathbf{U} \right\|_F^2 \\ \{\mathbf{G}, \mathbf{H}\} &\leftarrow \arg \min_{\{\mathbf{G}, \mathbf{H}\}} \frac{\rho}{2} \left\| \mathbf{X} - \mathbf{P} \mathbf{G} \mathbf{H}^T + \mathbf{U} \right\|_F^2 \\ \mathbf{U} &\leftarrow \mathbf{U} + \mathbf{X} - \mathbf{P} \mathbf{G} \mathbf{H}^T \end{aligned} \quad (12)$$

4.1 Updating X with ΔE_{76}

The first update in Equation 12 is *independent for each pixel*. Although being a small problem, it is still nonlinear and non-convex. We use Newton’s method to solve it. To this end, the Jacobian matrix for each pixel $j = 1 \dots N$ is formulated as

$$\mathbf{J}_{X_j} = \begin{bmatrix} 0 & -116\phi'(X_j^{(y)}, W_y) & 0 \\ -500\phi'(X_j^{(x)}, W_x) & 500\phi'(X_j^{(y)}, W_y) & 0 \\ 0 & -200\phi'(X_j^{(y)}, W_y) & 200\phi'(X_j^{(z)}, W_z) \\ \sqrt{\rho/2} & 0 & 0 \\ 0 & \sqrt{\rho/2} & 0 \\ 0 & 0 & \sqrt{\rho/2} \end{bmatrix} \quad (13)$$

where

$$\phi'(x, w) = \begin{cases} \frac{1}{3w^{1/3}}x^{-2/3} & \text{if } \frac{x}{w} > (\frac{6}{29})^3 \\ \frac{1}{3w}(\frac{29}{6})^2 & \text{otherwise} \end{cases} \quad (14)$$

Newton’s method is an iteration scheme that updates X_j as

$$X_j \leftarrow X_j - \left(\mathbf{J}_{X_j}^T \mathbf{J}_{X_j} \right)^{-1} \mathbf{J}_{X_j}^T f(X_j), \quad (15)$$

where $f : \mathbb{R}^3 \rightarrow \mathbb{R}$ is the objective function evaluating the first term of Equation 12 for a single pixel j . In practice, we compute all results in this paper with the slightly more sophisticated ΔE_{94} metric instead of ΔE_{76} . A detailed derivation of the corresponding X, which builds on Equation 13, can be found in Suppl. Sec. C.2.

4.2 Updating G and H

The factor matrices can be updated using conventional nonnegative matrix factorization [Lee and Seung 1999] by updating **G** and **H** in an alternating manner. We found that the rank-one residue iteration proposed by Ho [2008] improves convergence and usually finds better solutions for the **H** update. Unfortunately, it is unclear how to apply this method to the **G** update due to the projection matrix **P**. We summarize the NMF updates using pseudo-code in Algorithm 1.

Algorithm 1 Nonnegative Matrix Factorization (NMF)

```

1: function [G, H] = NMF (X, G, H,  $\sigma$ )
2: for each NMF iteration
3:   G  $\leftarrow$  [G  $\circ$  ( $\mathbf{P}^T \mathbf{X} \mathbf{H}$ ) // ( $\mathbf{P}^T (\mathbf{P} \mathbf{G} \mathbf{H}^T) \mathbf{H}^T + \epsilon$ )]+
4:   for k=1...K # rank-1 residue update
5:     Rk = X -  $\sum_{i \neq k} (\mathbf{P} \mathbf{G})_i \mathbf{H}_i^T$ 
6:     Hk  $\leftarrow$  [RkT (PG)k]+ // ((PG)kT (PG)k)
7:   end
8: end

```

The overall ADMM updates, implementing Equation 12, comprises three steps in each iteration: the Newton updates of each pixel in X, the NMF update of **G** and **H**, and an update of the slack variable **U**. These steps are outlined in Algorithm 2.

5 Implementation

Software PNMF was implemented in Matlab. The X-update step in ADMM is based on parallel per-pixel operations; we sped up this step considerably by implementing it as a mex module that loops through the pixels for each Newton iteration. For each pixel, we

Algorithm 2 Perceptual Nonnegative Matrix Factorization (PNMF)

```

1: function [G, H] = PNMF
2: init G = rand(M  $\times$  K), H = rand(N  $\times$  K)
3:   X = zeros(O  $\times$  N), U = zeros(O  $\times$  N)
4: for each ADMM iteration
5:   for each pixel j
6:     for each Newton iteration
7:       Xj  $\leftarrow$  Xj - ( $\mathbf{J}_{X_j}^T \mathbf{J}_{X_j}$ )-1  $\mathbf{J}_{X_j}^T f(X_j)$ 
8:     end
9:   end
10: [G, H]  $\leftarrow$  NMF (X+U, G, H,  $\frac{\rho}{\sigma}$ )
11: U  $\leftarrow$  U + X - PABT
12: end

```

calculate the gradient, and we apply a simple backtracking scheme to determine the step length that yielded the largest decrease in the residual. For all results, we use the Newton updates for ΔE_{94} described in Supplemental Section C.2, which are slightly more complex than the updates for ΔE_{76} outlined in Section 4.1. The ρ parameter in ADMM is chosen as 1.5E+06, which heuristically yielded consistent and fast convergence.

We convert all target multispectral images to CIE Lab space using the white point specified by the Rec. 709 standard. For all results shown in this paper, we used a scaling parameter $\beta = 1$ which never compromises image brightness; however, choosing $\beta < 1$ may be useful to trade brightness for higher color accuracy. For PNMF, we use 1,500 total ADMM iterations, where within each iteration there are 5 Newton iterations for the first update step, and 10 NMF iterations for the second step. For the comparisons to NMF, we use 35,000 NMF iterations, which was adequate to yield convergence to a delta between iterations smaller than at least 1e-6. On an Intel i7-4790 3.6 GHz processor with 8 GB RAM, for 1,500 iterations on a 512 \times 512 pixel image, PNMF took 9,500 seconds and for 35,000 iterations on the same image, NMF took 13,500 seconds. While the processing times seem prohibitive, PNMF actually converges in less time than NMF. Note that the per-pixel operations are all independent and can be further sped up through multi-core concurrency or a GPU-based implementation, and NMF has been shown to run in real-time using efficient GPU implementations. We did not attempt to optimize runtimes and believe there is a clear path towards optimizing these significantly, as discussed in Section 7.

Projector Hardware The prototype consists of a modified Texas Instruments LightCrafter 0.3 WVGA DLP chipset projector, a six LED fiber-coupled analog light engine, control electronics, and software for synchronization. The existing LEDs and dichroic mirrors from the DLP projector were removed. A fiber optic light guide from an external light engine was mounted in place of the LED nearest to the DMD. The light engine is a Lumencor Spectra X that contains six individually controllable high-power LEDs which are cooled and optically combined into the light guide. To enable display of three custom primaries at a frame rate faster than the human flicker fusion rate of 60 Hz, we require the ability to configure all six LED channels to a different analog state every 5.56 ms or less. To achieve this, we utilized our light engine’s high speed TTL-based electronic shutter (up to 5 kHz) in combination with analog control of each LED sequentially through a serial port protocol at baud rates up to 115,200. More specifically, we first shutter all LEDs, and then sequentially set each of the six LEDs to their respective analog values using serial commands, which takes 3 ms in total. We then unshutter the light source for 2 ms before repeating the procedure for the next primary, enabling us the ability to display 3-primary images at 66.7 Hz. This lower 40% duty

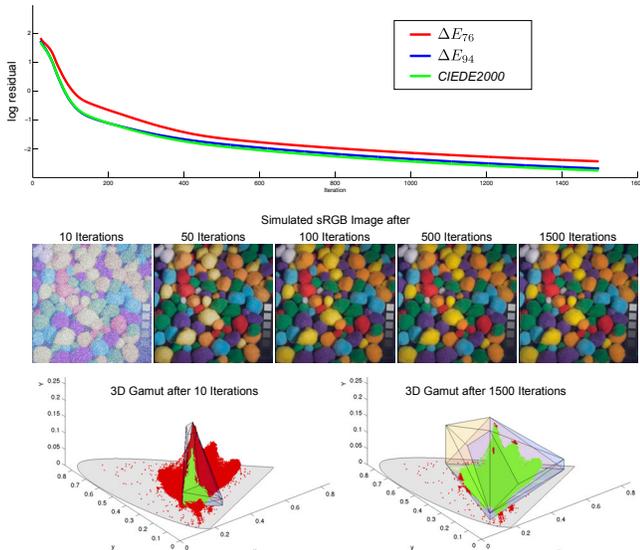


Figure 4: Convergence of a dataset optimized for the average ΔE_{94} error metric. In addition to plotting the ΔE_{94} log residual, we also plot errors in ΔE_{76} and CIEDE2000 and observe strong correlations between these metrics (top). Intermediate results, printed in sRGB, are shown in the center. Finally, the 3D gamuts are visualized in CIE xyY in the bottom. Here, green points represent target colors that are within the optimized gamut and red points those outside. After convergence, only a few very bright and very dark points are outside the optimized gamut.

cycle, which reduces the maximum brightness of the projected images, could be much higher if we had had higher speed or fully parallel analog control over all six LEDs, which we are considering for future implementations. To operate the prototype, we first preload a sequence of image frames onto the DMD, which is set to change frames upon receiving a TTL trigger signal. Then, to precisely synchronize the light source with the DMD, we use a National Instrument DAQ (USB-6343) to generate a series of digital voltage waveforms, which are pre-computed and preloaded onto the device. Here, the first channel controls the DMD trigger, the second channel controls the electronic shutter of the Lumencor, and the third channel we used to send the binary serial command sequences to the Lumencor, using a TTL-to-serial voltage level shifter as the interface. Software in MATLAB was written to preload images via a mouse emulator macro onto the LightCrafter graphical interface, and to compute and preload the digital voltage waveforms onto the DAQ over USB.

Calibration Software calibration was needed to determine the color gamut and gamma curves of the projector as well as non-linear intensity curves of the LEDs. To obtain the gamut, a white background was displayed on the DMD, while all LEDs of the light engine were turned on sequentially and their spectrum was captured using the spectrometer (see Fig. 6). The gamma curve of the DMD was fitted to intensity measurements made with a fixed LED illumination while varying DMD transmission values. Finally, the nonlinear intensity curve of each LED was captured with a similar procedure. These measurements are used to compute look-up tables that linearize all values sent to the LEDs and the DMD. Given the spectrum of each LED, we could obtain the CIE XYZ color coordinate of each LED. As a baseline for comparison, we found the linear combination of LED intensities that yielded the closest approximation to the sRGB gamut that our projector could achieve.

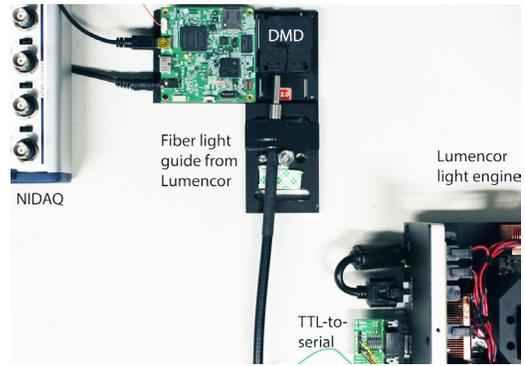


Figure 5: Photograph of prototype projector.

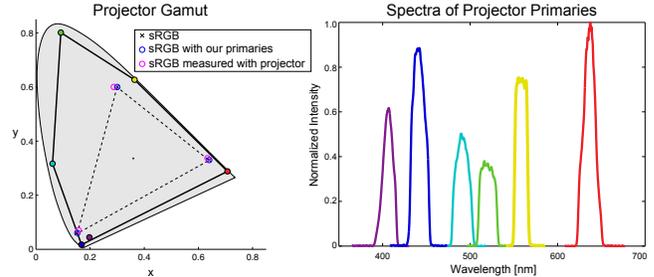


Figure 6: Gamut and spectral profiles of the projector prototype. The light engine provides six additive primaries that, together, cover a large part of the CIE xy chromaticity diagram (left). The spectral emission profiles of each LED are plotted on the right.

In particular, using `fmincon`, we solved for the LED coefficients that yielded the sRGB primary chromaticities as well as the sRGB white point when all primaries were summed together. The result was an sRGB gamut, scaled to the maximum brightness achievable with our projector. We used the white point of this gamut ($x = 0.3127$, $y = 0.3290$, $Y = 0.1774$) as the white point for all computations, such as in the conversion from XYZ to LAB space.

6 Results

Simulation performance We tested our algorithm on 11 different multispectral images. Results for one example are shown in Figure 7; results for all other images can be found in the supplement, and are summarized in Figure 9 and Table 1. We analyzed the factored images using a number of different metrics, as shown in Figure 9: ΔE_{76} , ΔE_{94} , CIEDE2000, SSIM in IPT color space [Bonnier et al. 2006], and PSNR. With all metrics, our algorithm performed better. In particular, PNMF 4 yielded a significantly lower ΔE than NMF 4, and PNMF 4 yielded a significantly lower ΔE_{94} than NMF 3 ($p = 0.002$ and $p = 0.003$, respectively, paired t-test with Welch correction for multiple comparisons).

Hardware performance Our projector prototype was able to successfully display adaptive gamut images. During the calibration, we verified using a single pixel spectrometer that the projector was able to reproduce target spectra within its range. Further, we were able to display artifact-free images. Because our projector can create a larger gamut than any of the capture or print devices available to us, we cannot directly show the image quality in the paper. As a proxy, to demonstrate that our projector generates images that look similar in content to the original images, in Fig. 8, we present photographs

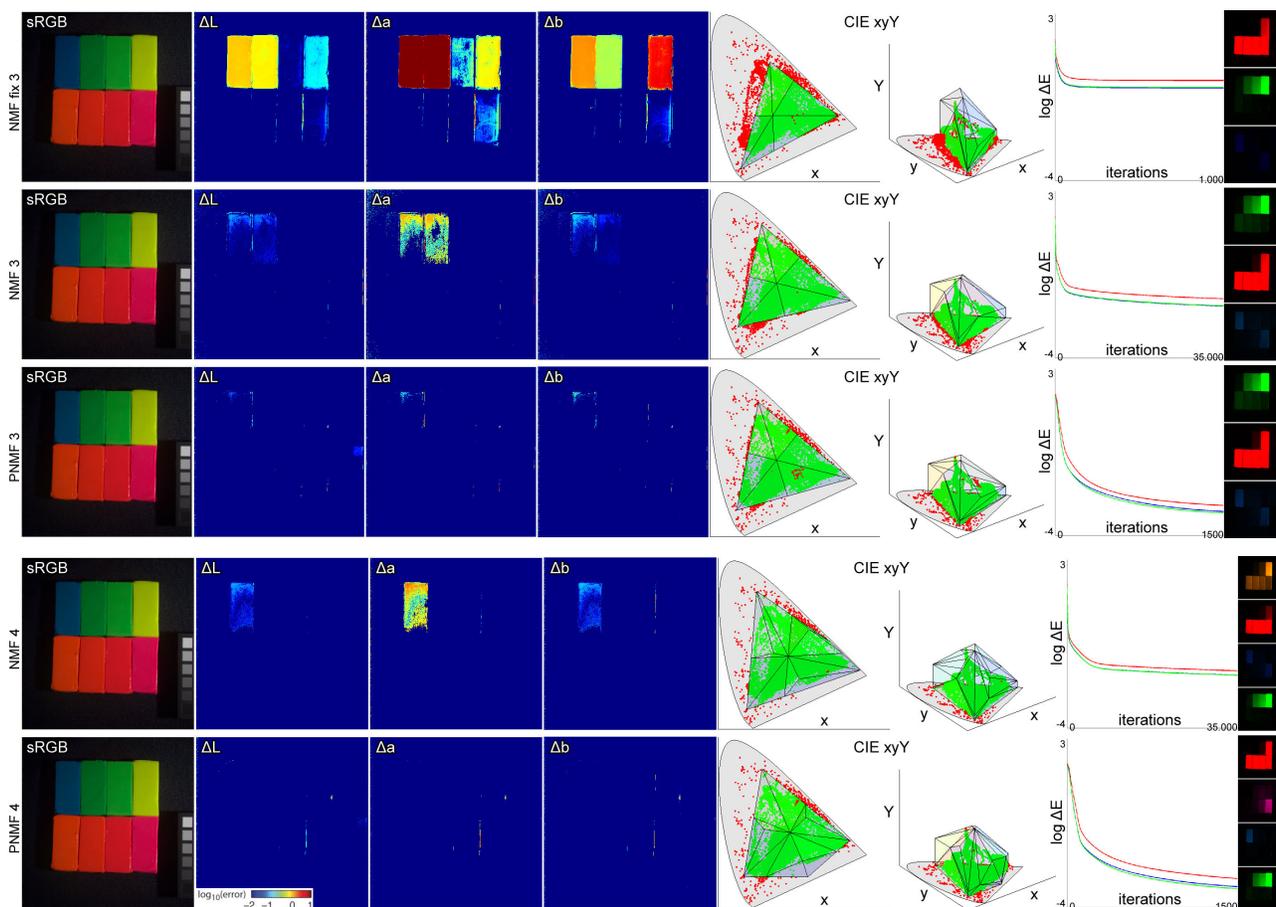


Figure 7: PNMF yields lower ΔE error than NMF and NMF with fixed primaries. Here we show for one example multispectral image the results of the five algorithms tested (NMF with 3 primaries fixed sRGB, NMF with 3 and 4 adaptive primaries, PNMF with 3 and 4 adaptive primaries). From left to right: simulation of the factored result image, clipped for the sake of printing to the sRGB color gamut; the per-pixel error in CIE Lab space between the target image and factored image, shown in log-space; the result gamut in CIE xyY space (top view and side view), overlaid with the CIE xyY coordinates of the target image—red points are pixels that fall outside of the gamut; the convergence trajectory; the images associated with each primary.

of the image displayed on the projector. Obviously, the camera has clipped the gamut of the image to be smaller than the gamut of the actual, displayed image; however, this still shows that our prototype is capable of producing artifact free images.

User Experiment For our user experiment, we aimed to support three claims: 1) A flexible gamut has the ability to adapt more readily in situations which cause a fixed gamut to fail; 2) PNMF can maintain perceptual color differences better than NMF; 3) Our hardware prototype is capable of conveying the results of the flexible gamut algorithm. We limited the scope of this user experiment; a more in depth user study will be the subject of future work.

We selected three pairs of similar colors, specified in CIE XYZ space. We additionally selected a fixed gamut that had the largest area in CIE xy chromaticity space possible within the gamut of our projector. Two of the color pairs were chosen such that the pairs were metamers of this fixed gamut: when we applied NMF with a fixed gamut as a simple gamut mapping technique, the colors in the pair were mapped to the same color. We note that with this gamut mapping technique, it is possible to find metamers for any reasonable fixed gamut. We also acknowledge that other gamut mapping techniques, such as a relative colorimetric approach, could mitigate

this metamer effect, however this would still lead to a loss in color fidelity and desaturation associated with shrinking the entire image’s gamut. An adaptive gamut, on the other hand, has enough degrees of freedom that it can always choose a gamut that eliminates at least some of the metamers observed in the fixed gamut case. Further, gamut mapping in a perceptually uniform space, using PNMF, spreads the error equally across all of the perceived colors and tries to retain color differences that are reflective of the original colors, which is more likely to remove metamers.

We generated an image with these three color pairs: assigned one member of each pair as the background color of a stripe, and the other member as the color of an overlaid circle. We factored this image using NMF3 with a fixed gamut, NMF3 with a flexible gamut, and PNMF3 with an adaptive gamut. The ΔE errors for each factorization were, respectively, 52, 32, and 13. The errors were large because of the premeditated difficult gamut of the target image. Using the prototype projector, $n = 6$ human subjects were each presented with each of the three factorizations a total of four times, in a counter-balanced, pseudo-random order. Each image was displayed for 600 ms, short enough such that detection of all three circles was at the perceptual threshold. The perceptual visibility of each circle was thus one indication of how well the algorithm and projector paired together faithfully represent the color differences that were

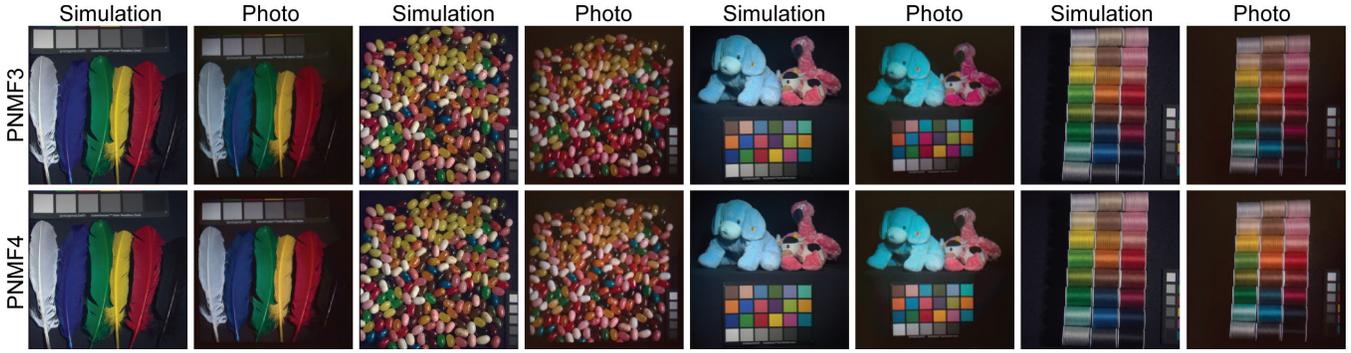


Figure 8: Simulation and photographed results for three multispectral datasets as displayed on the prototype projector. Note that the image colors have been clipped for printing from the gamut of the projector to that of sRGB. The photographed results match the simulations to within calibration error but the color space of the camera is not calibrated. Hence, there are slight differences between the simulated sRGB images and the captured photographs.

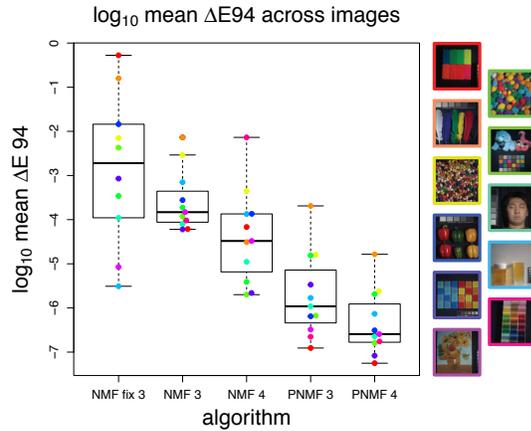


Figure 9: PNMf yields lower ΔE_{94} than NMF and fixed gamut NMF. Summary of \log -transformed mean ΔE across $n = 11$ multispectral images, for fixed gamut NMF with 3 primaries, flexible gamut NMF with 3 and 4 primaries, and flexible gamut PNMf with 3 and 4 primaries. PNMf 3 has significantly lower 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.

inherent to the target image. As shown in Figure 10, PNMf enabled significantly better detection of all three circles than NMF. Further, as expected, two of the color pairs were metamERICALLY invisible when mapped to the fixed gamut.

While this user experiment is by no means sufficient to prove that PNMf is better than NMF in general, it is an example that demonstrates the potential of the algorithm and projector design for faithfully representing the color of images that possess large and high contrast gamuts. A full user study will be the subject of future work.

Dataset	Algorithm	ΔE_{70}^*	ΔE_{94}^*	CIEDE2000	S-CIELab	SSIM	PSNR(dB)
Clay	NMF fix 3	1.56E+00	7.58E-01	8.00E-01	1.53E+00	9.89E-01	45.62
	NMF 3	2.99E-02	1.48E-02	1.52E-02	2.41E-02	1.00E+00	70.36
	NMF 4	2.35E-02	1.55E-02	1.54E-02	2.13E-02	1.00E+00	73.56
	PNMF 3	1.79E-03	1.00E-03	8.58E-04	2.52E-03	1.00E+00	75.45
	PNMF 4	1.53E-03	7.09E-04	5.81E-04	1.97E-03	1.00E+00	75.60
Feathers	NMF fix 3	1.30E+00	4.49E-01	4.86E-01	1.31E+00	9.91E-01	46.72
	NMF 3	3.43E-01	1.18E-01	1.22E-01	3.26E-01	9.98E-01	56.64
	NMF 4	3.21E-02	1.10E-02	1.20E-02	2.34E-02	1.00E+00	66.47
	PNMF 3	6.54E-02	2.50E-02	2.42E-02	7.90E-02	9.99E-01	62.24
	PNMF 4	2.47E-02	8.36E-03	8.23E-03	3.17E-02	9.99E-01	65.17
Beans	NMF fix 3	1.74E-01	1.16E-01	1.28E-01	1.80E-01	9.97E-01	55.94
	NMF 3	1.30E-01	7.92E-02	9.42E-02	1.10E-01	9.95E-01	61.01
	NMF 4	6.77E-02	3.51E-02	3.59E-02	5.42E-02	1.00E+00	64.49
	PNMF 3	1.19E-02	8.27E-03	6.62E-03	1.77E-02	1.00E+00	69.83
	PNMF 4	4.78E-03	3.62E-03	2.65E-03	7.47E-03	1.00E+00	73.04
Pompon	NMF fix 3	2.35E-01	9.35E-02	9.70E-02	2.28E-01	9.98E-01	58.10
	NMF 3	4.48E-02	1.96E-02	2.04E-02	2.71E-02	9.99E-01	67.18
	NMF 4	6.68E-03	3.35E-03	4.00E-03	4.53E-03	1.00E+00	72.75
	PNMF 3	3.69E-03	2.08E-03	1.77E-03	3.92E-03	1.00E+00	75.35
	PNMF 4	2.39E-03	1.12E-03	9.66E-04	2.56E-03	1.00E+00	76.79
Dog	NMF fix 3	5.16E-02	3.13E-02	3.30E-02	5.48E-02	1.00E+00	62.18
	NMF 3	4.27E-02	2.41E-02	2.51E-02	4.29E-02	1.00E+00	65.90
	NMF 4	7.69E-03	4.47E-03	5.27E-03	6.28E-03	1.00E+00	71.08
	PNMF 3	1.20E-02	8.14E-03	6.55E-03	1.26E-02	1.00E+00	69.78
	PNMF 4	4.89E-03	3.39E-03	2.70E-03	5.11E-03	1.00E+00	71.86
Face	NMF fix 3	2.77E-02	1.91E-02	2.06E-02	3.01E-02	1.00E+00	65.29
	NMF 3	2.60E-02	1.66E-02	1.81E-02	2.42E-02	1.00E+00	67.38
	NMF 4	1.08E-02	7.04E-03	8.60E-03	8.39E-03	1.00E+00	68.49
	PNMF 3	3.47E-03	2.57E-03	2.29E-03	3.40E-03	1.00E+00	70.15
	PNMF 4	1.59E-03	1.30E-03	1.08E-03	1.62E-03	1.00E+00	71.87
Beer	NMF fix 3	6.28E-03	4.05E-03	4.48E-03	6.54E-03	1.00E+00	71.23
	NMF 3	7.01E-02	4.27E-02	4.79E-02	6.86E-02	9.99E-01	58.59
	NMF 4	3.80E-02	2.08E-02	2.33E-02	3.68E-02	9.99E-01	62.78
	PNMF 3	4.57E-03	3.11E-03	2.65E-03	4.82E-03	1.00E+00	71.46
	PNMF 4	3.04E-03	2.17E-03	1.73E-03	3.08E-03	1.00E+00	71.79
Pepper	NMF fix 3	5.34E-01	1.59E-01	1.85E-01	4.59E-01	9.98E-01	50.50
	NMF 3	4.88E-02	2.85E-02	3.18E-02	4.07E-02	9.99E-01	66.91
	NMF 4	3.53E-02	2.09E-02	2.35E-02	2.93E-02	9.99E-01	67.57
	PNMF 3	2.52E-03	2.05E-03	1.57E-03	3.89E-03	1.00E+00	73.59
	PNMF 4	1.68E-03	1.49E-03	1.06E-03	2.86E-03	1.00E+00	74.74
Tiles	NMF fix 3	8.34E-02	4.64E-02	5.08E-02	9.73E-02	9.98E-01	56.86
	NMF 3	2.69E-02	1.47E-02	1.51E-02	2.52E-02	1.00E+00	66.17
	NMF 4	6.03E-03	3.48E-03	4.04E-03	5.02E-03	1.00E+00	71.46
	PNMF 3	4.75E-03	4.20E-03	3.12E-03	1.32E-02	1.00E+00	72.15
	PNMF 4	8.81E-04	8.41E-04	5.79E-04	2.02E-03	1.00E+00	77.73
Painting	NMF fix 3	1.47E-02	6.27E-03	7.03E-03	1.42E-02	1.00E+00	68.08
	NMF 3	3.89E-02	2.17E-02	2.57E-02	2.91E-02	9.99E-01	64.40
	NMF 4	2.10E-02	1.13E-02	1.34E-02	1.62E-02	9.99E-01	66.80
	PNMF 3	2.41E-03	1.52E-03	1.38E-03	2.46E-03	1.00E+00	73.50
	PNMF 4	1.98E-03	1.37E-03	1.16E-03	2.10E-03	1.00E+00	71.45
Thread	NMF fix 3	4.40E-02	2.04E-02	2.12E-02	4.36E-02	1.00E+00	66.18
	NMF 3	3.88E-02	1.80E-02	1.90E-02	3.18E-02	1.00E+00	66.96
	NMF 4	2.48E-02	1.18E-01	1.28E-02	1.78E-02	1.00E+00	69.42
	PNMF 3	1.65E-03	1.29E-03	1.07E-03	2.17E-03	1.00E+00	73.46
	PNMF 4	1.71E-03	1.16E-03	9.53E-04	2.24E-03	1.00E+00	74.34

Table 1: Evaluation of NMF with a fix three-primary gamut as well as adaptive NMF and PNMf for three and four-primary displays using 11 datasets. PNMf in all cases results in a better quality than both NMF and fix-gamut NMF. For most examples, SSIM did not calculate a significant difference between the results (light green).

7 Discussion

In summary, we introduce a new paradigm for color display: adaptive, perceptually-driven factored spectral projection. We demon-

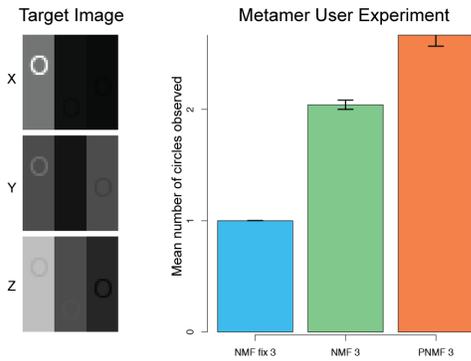


Figure 10: Design and results of a user experiment demonstrating how PNMf differentiates between colors that are mapped to metamers or near-metamers when using NMF or NMF with a fixed gamut. On the left, the target image is displayed in CIE XYZ space. On the right are the results of the user study: the 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.

strate with many examples that adaptive gamuts have strong benefits over fixed gamuts and that the proposed algorithm yields the best quality in all common quality metrics. Finally, we evaluate the proposed display system with a prototype projector, through photographs and a user experiment that validates our simulations.

Limitations and Future Work The solver is currently implemented in Matlab and does not run in real-time. Although real-time processing may not be needed in many applications, such as movies, it would be desirable for interactive applications. We are convinced that this is possible. Specifically, the discussed ADMM update scheme could be replaced by more efficient primal-dual methods, which would improve overall convergence (e.g., [Chambolle and Pock 2011]). Further, the Newton updates of each pixel in the \mathbf{X} -update are independent of each other and could be parallelized on the GPU. Finally, the NMF updates could also be ported to the GPU, which has been shown to provide a speed-up of approx. $15\times$ [Platos et al. 2010]. Other types of nonnegative factorization algorithms with a similar data size have been demonstrated to run in real time [Wetzstein et al. 2012; Heide et al. 2014]. We believe that there is a clear path to efficiently implementing the proposed method, possibly in real time, but that was not the goal of this project and is left for future work. For applications in video processing we could further reduce the required iterations by warm-starting the optimization of a frame with the output of the previous frame or just compute a fixed gamut for a video sequence or keyframe-segmented clips rather than for every frame. Imposing constraints on temporal consistency between successively-displayed video frames may be necessary for future video applications and could easily be incorporated into the optimization. It would be interesting to analyze the behavior of quantization in more detail. Large gamuts and high dynamic range content require a larger bit depth than what is available currently. Adaptive color display should ideally also consider the available bit depth for representing colors. We do observe that the gamuts produced by our algorithm are usually a tight fit around the target colors, which would suggest that the available bit depth is used very efficiently.

In the future, we are most interested in exploring new display devices that could benefit from adaptive color reproduction. Any display hardware that supports dynamic adjustment of the color pri-

maries would be supported by the proposed algorithm, the only adjustment that has to be made is in the basic image formation (Eq. 1). We would like to optimize for other color metrics, especially those specifically modeling spatial image variation. We demonstrate that our optimization scheme consistently drives down s-CIE Lab, but more sophisticated metrics could further improve color reproduction with our system. Eventually, we hope to stimulate the color science community to develop color metrics that are both accurate but also useful for optimization procedures including conventional gamut mapping techniques and algorithms for joint primary selection and gamut mapping, such as the proposed. Finally, we always optimize color fidelity for the “standard” observer—personalizing these kinds of display for the particular characteristics of a person would be most interesting.

8 Conclusion

Trends in the display industry are clear: higher resolution, extended contrast, and wider color gamuts will become commonplace in the near future. Increasing resolution will be provided by advanced nano-fabrication technology and high dynamic range displays have matured to the point where they can be deployed to the market. With the flexible computational color display system proposed in this paper, we provide a viable solution for adaptive, wide color gamut display that could be seamlessly integrated into cinematic projection systems, home theaters, and office projectors.

Acknowledgements

The authors would like to thank Brian Wandell, Felix Heide, and Logan Grosenick for early feedback; Michael Jin from Silicon Micro Display and Matthew Hirsch for help with the opto-mechanical parts of an early prototype; and Karl Deisseroth for lending us the Lumencor. Isaac Kauvar was supported by an NSF graduate fellowship. Samuel Yang was supported by a National Defense Science and Engineering Graduate Fellowship, and Gordon Wetzstein was supported by a Terman Faculty Fellowship.

References

- AJITO, T., OBI, T., YAMAGUCHI, M., AND OHYAMA, N. 2000. Expanded color gamut reproduced by six-primary projection display. In *Proc. SPIE 3954*, 130–137.
- BANTERLE, F., ARTUSI, A., AYDIN, T. O., DIDYK, P., EISEMANN, E., GUTIERREZ, D., MANTIUK, R., AND MYSZKOWSKI, K. 2011. Multidimensional image retargeting. In *SIGGRAPH Asia 2011 Courses*, 15:1–15:612.
- BEN-CHORIN, M., AND ELIAV, D. 2007. Multi-primary design of spectrally accurate displays. *Journal of the SID* 15, 9, 667–677.
- BERTHOUSOZ, F., AND FATTAL, R. 2012. Resolution enhancement by vibrating displays. *ACM Trans. Graph.* 31, 2, 15:1–15:14.
- BONNIER, N., SCHMITT, F., BRETTEL, H., AND BERCHE, S. 2006. Evaluation of spatial gamut mapping algorithms. In *Color and Imaging Conference*, vol. 2006, Society for Imaging Science and Technology, 56–61.
- BOYD, S., PARIKH, N., CHU, E., PELEATO, B., AND ECKSTEIN, J. 2011. Distributed optimization and statistical learning via the alternating direction method of multipliers. *Found. Trends Mach. Learn.* 3, 1, 1–122.

- CHAMBOLLE, A., AND POCK, T. 2011. A first-order primal-dual algorithm for convex problems with applications to imaging. *J. Math. Imaging and Vision*, 40, 120145.
- CIE. 1932. *Proceedings*. Cambridge University Press.
- CIE. 2001. *Improvement to industrial colour-difference evaluation*. CIE Central Bureau.
- DIDYK, P., EISEMANN, E., RITSCHER, T., MYSZKOWSKI, K., AND SEIDEL, H.-P. 2010. Apparent display resolution enhancement for moving images. *ACM Trans. Graph. (SIGGRAPH)* 29, 4, 113:1–113:8.
- GEHM, M. E., JOHN, R., BRADY, D. J., WILLETT, R. M., AND SCHULZ, T. J. 2007. Single-shot compressive spectral imaging with a dual-disperser architecture. *OSA Opt. Express* 15, 21, 14013–14027.
- HEIDE, F., LANMAN, D., REDDY, D., KAUTZ, J., PULLI, K., AND LUEBKE, D. 2014. Cascaded displays: Spatiotemporal superresolution using offset pixel layers. *ACM Trans. Graph. (SIGGRAPH)* 33, 4, 60:1–60:11.
- HIRSCH, M., WETZSTEIN, G., AND RASKAR, R. 2014. A compressive light field projection system. *ACM Trans. Graph. (SIGGRAPH)* 33, 4, 58:1–58:12.
- HO, N. 2008. *Nonnegative Matrix Factorization Algorithms and Applications*. PhD thesis, Universite catholique de Louvain.
- KELLY, D. H. 1979. Motion and vision. ii. stabilized spatiotemporal threshold surface. *Journal of the Optical Society of America* 69, 1340–1349.
- LANMAN, D., HIRSCH, M., KIM, Y., AND RASKAR, R. 2010. Content-adaptive parallax barriers: Optimizing dual-layer 3d displays using low-rank light field factorization. *ACM Trans. Graph. (SIGGRAPH)* 29, 6, 163:1–163:10.
- LEE, D. D., AND SEUNG, S. 1999. Learning the Parts of Objects by Non-negative Matrix Factorization. *Nature* 401, 788–791.
- LI, Y., MAJUMDER, A., LU, D., AND GOPI, M. 2015. Content-independent multi-spectral display using superimposed projections. *Computer Graphics Forum (Eurographics)*.
- LIN, X., LIU, Y., WU, J., AND DAI, Q. 2014. Spatial-spectral encoded compressive hyperspectral imaging. *ACM Trans. Graph. (SIGGRAPH Asia)* 33, 6.
- LONG, D., AND FAIRCHILD, M. 2011. Optimizing spectral color reproduction in multiprimary digital projection. In *Color and Imaging Conference*.
- MACADAM, D. L. 1942. Visual sensitivities to color differences in daylight. *OSA JOS A* 32, 5, 247–273.
- MAJUMDER, A., AND BROWN, M. S. 2007. *Practical Multi-projector Display Design*. AK Peters.
- MAJUMDER, A., BROWN, R., AND EL-GHOROURY, H. 2010. Display gamut reshaping for color emulation and balancing. In *Proc. CVPR Workshops*, 17–24.
- MASIA, B., WETZSTEIN, G., DIDYK, P., AND GUTIERREZ, D. 2013. A survey on computational displays: Pushing the boundaries of optics, computation, and perception. *Computers & Graphics* 37, 8, 1012 – 1038.
- MOHAN, A., RASKAR, R., AND TUMBLIN, J. 2008. Agile spectrum imaging: Programmable wavelength modulation for cameras and projectors. *Computer Graphics Forum* 27, 2, 709–717.
- PAUCA, V. P., PIPER, J., AND PLEMMONS, R. J. 2006. Non-negative matrix factorization for spectral data analysis. *Linear Algebra and its Applications* 416, 1, 29 – 47.
- PLATOS, J., GAJDOS, P., KROMER, P., AND SNASEL, V. 2010. Non-negative matrix factorization on gpu. *Comm. Computer and Information Science Volume 87*, 21–30.
- RICE, J., BROWN, S., NEIRA, J., AND BOUSQUET, R. 2007. A hyperspectral image projector for hyperspectral imagers. In *Proc. SPIE* 6565, 65650C.
- RICE, J. P., BROWN, S. W., ALLEN, D. W., YOON, H. W., LITORJA, M., AND HWANG, J. C. 2012. Hyperspectral image projector applications. vol. 8254, 82540R–82540R–8.
- RODRÍGUEZ-PARDO, C. E., SHARMA, G., FENG, X.-F., SPEIGLE, J., AND SEZAN, I. 2012. Optimal gamut volume design for three primary and multiprimary display systems. In *Proc. SPIE Electronic Imaging*, 82920C–82920C.
- SAJADI, B., MAJUMDER, A., HIWADA, K., MAKI, A., AND RASKAR, R. 2011. Switchable primaries using shiftable layers of color filter arrays. *ACM Trans. Graph. (SIGGRAPH)* 30, 4.
- SAJADI, B., GOPI, M., AND MAJUMDER, A. 2012. Edge-guided resolution enhancement in projectors via optical pixel sharing. *ACM Trans. Graph. (SIGGRAPH)* 31, 4, 79:1–79:122.
- SEETZEN, H., HEIDRICH, W., STUERZLINGER, W., WARD, G., WHITEHEAD, L., TRENTACOSTE, M., GHOSH, A., AND VOROZCOVS, A. 2004. High dynamic range display systems. *ACM Trans. Graph. (SIGGRAPH)* 23, 3, 760–768.
- TERAGAWA, M., YOSHIDA, A., YOSHIYAMA, K., NAKAGAWA, S., TOMIZAWA, K., AND YOSHIDA, Y. 2012. Review paper: Multi-primary-color displays: The latest technologies and their benefits. *Journal of the SID* 20, 1, 1–11.
- WAGADARIKAR, A. A., PITSIANIS, N. P., SUN, X., AND BRADY, D. J. 2009. Video rate spectral imaging using a coded aperture snapshot spectral imager. *OSA Opt. Express* 17, 8, 6368–6388.
- WETZSTEIN, G., LANMAN, D., HEIDRICH, W., AND RASKAR, R. 2011. Layered 3d: Tomographic image synthesis for attenuation-based light field and high dynamic range displays. *ACM Trans. Graph. (SIGGRAPH)* 30, 4.
- WETZSTEIN, G., LANMAN, D., HIRSCH, M., AND RASKAR, R. 2012. Tensor displays: Compressive light field synthesis using multilayer displays with directional backlighting. *ACM Trans. Graph. (SIGGRAPH)* 31, 4, 80:1–80:11.
- WITT, K. 2007. *CIE Color Difference Metrics*. Wiley.
- ZHANG, X. M., AND WANDELL, B. A. 1996. A spatial extension to cielab for digital color image reproduction. In *Proc. SID*.