

Deep Multispectral Painting Reproduction via Multi-Layer, Custom-Ink Printing

LIANG SHI, MIT CSAIL

VAHID BABAEI, MIT CSAIL, MPI Informatik & Saarland University MMCI

CHANGIL KIM, MIT CSAIL

MICHAEL FOSHEY, MIT CSAIL

YUANMING HU, MIT CSAIL

PITCHAYA SITTHI-AMORN, Chulalongkorn University

SZYMON RUSINKIEWICZ, Princeton University

WOJCIECH MATUSIK, MIT CSAIL




Fig. 1. We present a spectral reproduction technique using a 3D printer. Our workflow targets accurate reproduction of paintings and provides faithful color reproductions under varying light sources. Above, we show three printed replicas of oil paintings with different image statistics, generated by our method, next to the original. On the right, we show cropped regions (three water lilies) from big water lily replica with the original under varying lighting sources. Paintings ©Azadeh Asadi.

We propose a workflow for spectral reproduction of paintings, which captures a painting's spectral color, invariant to illumination, and reproduces it using multi-material 3D printing. We take advantage of the current 3D printers' capabilities of combining highly concentrated inks with a large number of layers, to expand the spectral gamut of a set of inks. We use a data-driven method to both predict the spectrum of a printed ink stack and optimize for the stack layout that best matches a target spectrum. This bidirectional mapping is modeled using a pair of neural networks, which

are optimized through a problem-specific multi-objective loss function. Our loss function helps find the best possible ink layout resulting in the balance between spectral reproduction and colorimetric accuracy under a multitude of illuminants. In addition, we introduce a novel spectral vector error diffusion algorithm based on combining color contoning and halftoning, which simultaneously solves the layout discretization and color quantization problems, accurately and efficiently. Our workflow outperforms the state-of-the-art models for spectral prediction and layout optimization. We demonstrate reproduction of a number of real paintings and historically important pigments using our prototype implementation that uses 10 custom inks with varying spectra and a resin-based 3D printer.

CCS Concepts: • **Computing methodologies** → **Computer graphics**;

Additional Key Words and Phrases: spectral reproduction, multi-spectral imaging, 3d printing

ACM Reference Format:

Liang Shi, Vahid Babaei, Changil Kim, Michael Foshey, Yuanming Hu, Pitchaya Sitthi-Amorn, Szymon Rusinkiewicz, and Wojciech Matusik. 2018. Deep Multispectral Painting Reproduction via Multi-Layer, Custom-Ink Printing. *ACM Trans. Graph.* 37, 6, Article 271 (November 2018), 15 pages. <https://doi.org/10.1145/3272127.3275057>

Authors' addresses: Liang Shi, MIT CSAIL, liangs@mit.edu; Vahid Babaei, MIT CSAIL, MPI Informatik & Saarland University MMCI, babaei@mit.edu; Changil Kim, MIT CSAIL, changil@csail.mit.edu; Michael Foshey, MIT CSAIL, mfoshey@mit.edu; Yuanming Hu, MIT CSAIL, yuanming@mit.edu; Pitchaya Sitthi-Amorn, Chulalongkorn University, ps4wd@virginia.edu; Szymon Rusinkiewicz, Princeton University, smr@princeton.edu; Wojciech Matusik, MIT CSAIL, wojciech@csail.mit.edu.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
© 2018 Copyright held by the owner/author(s). Publication rights licensed to Association for Computing Machinery.
0730-0301/2018/11-ART271 \$15.00
<https://doi.org/10.1145/3272127.3275057>

1 INTRODUCTION

Fine art is a crucial element of human culture, but our reliance on museums to exhibit original paintings and sculpture inherently limits access and leaves those precious originals vulnerable to deterioration and damage. The availability of high-quality facsimiles has the potential to not only broaden exposure but also leave the originals unharmed, in such scenarios as restoration practice, conservatory studies and education in museums [Elkhuizen et al. 2014]. Paintings are brilliant examples of fine art artifacts and natural first candidates for an appearance reproduction effort. The advance of both 3D printing and 3D scanning technology has brought an increasing interest in high-quality painting reproduction to both academia and industry. A number of commercial Companies have established relationships with art distributors and museums to provide painting reproductions and are preparing for potential mass production of artwork.

Traditionally, the color reproduction of paintings is carried out using 2D printers. However, they have severe shortcomings for high-fidelity color reproduction as a consequence of two interrelated problems. First, their color gamut is limited because of their fixed set of inks, which usually includes only cyan, magenta, yellow, and black (CMYK). Although the limited gamut can be overcome with a larger number and a wider variety of inks, only a limited amount of ink can be deposited within a given region. Exceeding the so-called total ink limit results in deteriorated image quality, ink blotting, or mechanical malfunction of the printer [Babaei and Hersch 2016]. Second, 2D printers predominantly use *colorimetric color reproduction*, in which the color matching is carried out under a reference illuminant and assumes a set of standard observers. This can lead to *metamerism*, a well-known problem in color reproduction wherein a good reproduction is obtained under one light source, but not under another [Wyszecki and Stiles 1982].

In this work, we focus on an accurate reproduction of spectral color with 3D printing. The recent breakthrough in high-resolution, multi-material 3D printing carries the potential for groundbreaking advances in fine art reproduction. While 3D printers are powerful devices for fabricating objects with custom and complex geometry, they can also create multi-material composites that result in new appearance properties. Among them, the spectral color is often considered as most important appearance attribute for most forms of paintings. The spectral reflectance of a surface encodes the complete information about its color. Therefore, the *spectral color reproduction* is invariant to the color of the light source under which the reproduction is observed.

Good spectral reproduction requires a large spectral gamut. We, therefore, equip our 3D printer with 10 inks and take advantage of its ability to combine highly concentrated inks with a large number of thin layers, inspired by the recently proposed color contoning technique [Babaei et al. 2017]. In doing so, we expand the gamut realized by our set of inks significantly. The use of high-concentration inks in contoning, however, comes at the cost of color quantization artifacts. Therefore, we introduce a novel spectral vector error diffusion halftoning [Kawaguchi et al. 1999] that uses the ink stacks of all possible layer combinations as halftone primaries with an

unprecedented efficiency. This turns the spectral vector error diffusion of potentially billions of primaries into a viable approach, with significantly alleviated artifacts (due to an insufficient number of primaries). Furthermore, the halftoning resolves the problems associated with the discrete nature of ink layers, making even a simple layer rounding reliable in practice.

Accurate modeling of the complex light transport across ink layers, in the presence of printing inaccuracies, is a challenging task if pursued using physical modeling methods. We instead take a data-driven approach, modeling the bidirectional mapping between the spectral reflectance and the ink stack layout using artificial neural networks, which are learned from the actual printed ink stacks. We produce superior performance in both reproduction quality and run-time, compared to previous physically-based or data-driven approaches. We validate and evaluate our approach using historically important pigments in painting and a number of real paintings.

Our main contributions are:

- A complete physical reproduction framework, comprising both acquisition and fabrication, tailored for the re-creation of paintings;
- A data-driven spectral color prediction model based on neural networks that outperforms existing methods for both predicting the spectrum of a stack of ink layers and finding an optimal ink stack for a target spectrum;
- A highly efficient spectral vector error diffusion method that combines ideas from contoning (combining thin layers of inks) and halftoning (employing spatial modulation) to achieve smooth yet accurate color reproduction;
- A dataset, which we will make it publicly available, of 20,878 contone ink stack spectra and layouts, spectrally captured oil paintings, together with their optimized layouts using our ink library, and photographs of our printed reproductions under multiple illuminations.

2 PREVIOUS WORK

Our work builds on the rich history of research on color reproduction while embracing new printing technologies and techniques. Below we briefly review the work most relevant to ours.

Custom-Ink Color Printing. The 2D printing literature is abundant with custom-ink printing (sometimes known as *n*-ink printing), in which the printer employs inks different from the traditional CMYK, mostly to expand the gamut of the printer. There are multiple challenges in custom-ink 2D printing, however, including color prediction of halftones, color separation, increased halftoning complexity, and the total ink limit [Babaei and Hersch 2016]. Stollnitz et al. [1998] propose a color reproduction workflow using custom inks that addresses these challenges. In addition, their framework is capable of choosing the best inks for reproducing a certain image. Ostromoukhov [1993] introduces a heptatone (7-ink) printing system that includes cyan, yellow, black, red, green, blue, and purple. Rossier [2013] expands the color gamut significantly by combining daylight fluorescent inks with the CMYK palette. Custom-ink printing may also be used for reducing the visibility of halftones or ink consumption [Son et al. 2011]. Today there are consumer-grade desktop 2D printers available that use 10 inks or more.

Color 3D Printing. While the use of multiple, custom inks is prevalent in 2D printing, 3D printing has recently begun to accommodate color printing. Although the effort to introduce multiple colors in 3D printing has been started with the fused deposition modeling (FDM) technology [Reiner et al. 2014; Hergel and Lefebvre 2014], the best quality for fine art reproduction is currently achieved by resin-based inkjet 3D printers. Brunton et al. [2015] introduced an error diffusion algorithm for the surfaces of 3D printed geometry. In order to eliminate the halftoning artifacts, Babaei et al. [2017] proposed a *contoning* scheme in which different inks are layered to form the print; our printing workflow builds on this idea. Recently, Elek et al. [2017] proposed a color reproduction workflow that preserves the texture by simulating the crosstalk between neighboring voxels using Monte Carlo path tracing. It is also possible to incorporate translucency in color printing pipelines, owing to clear resins, and thereby creating spatially-varying translucency and color using a 3D printer [Brunton et al. 2018]. Yet another approach to decorating a surface is to transfer a previously printed texture to the 3D object, using methods such as hydrographic printing [Zhang et al. 2015; Panozzo et al. 2015] and thermoforming [Schüller et al. 2016; Zhang et al. 2016]. All these methods, however, focus on colorimetric reproduction using CMYK inks. Compared to them, our work aims at spectral reproduction and uses a significantly larger number of inks to achieve a larger color gamut.

Spectral Printing. Spectral printing has been an active research area in 2D printing. Most of these works, however, focus on spectral modeling of 2D printers, tested on in-gamut colors, and not a complete workflow for generating printouts. The prediction models are often based on the Yule-Nielsen spectral Neugebauer model [Taplin and Berns 2001], or its cellular version [Chen et al. 2004]. Berns et al. [2008] designed a spectral acquisition and modeling system, specifically for works of art. The most promising approaches to spectral reproduction are based on higher-dimensional interim connection spaces [Derhak and Rosen 2006; Tsutsumi et al. 2008]. However, as pointed out by Morović et al. [2012], the major limit on the quality of spectral reproduction is the spectral gamut of a 2D printer, which is significantly smaller than the naturally occurring spectral gamut. It is noteworthy that spectral gamut mapping [Rosen and Derhak 2006; Urban and Berns 2011] can improve the quality of spectral reproduction when working with the limited spectral gamut of 2D printers. With the new flexibility offered by 3D printing, especially our freedom to choose different inks and increase their thickness, and our focus on painting reproduction, we significantly expand the printer gamut and for the first time, to the best of our knowledge, target full spectral reproduction.

Spectral Vector Error Diffusion. Spectral vector error diffusion (sVED) [Kawaguchi et al. 1999] is an extension of classical RGB or CMY color vector error diffusion [Klassen et al. 1994]: instead of performing a classical error diffusion algorithm on different layers of inks, the spectrum at every pixel is compared against the spectra of available primary inks and the spectral error is diffused to the pixel's neighborhood, wavelength by wavelength. sVED was introduced as a promising spectral reproduction technique that encapsulates color separation and halftoning in a single step. It was soon realized, however, that it suffers from low reproduction accuracy.

Moreover, due to the often-large spectrally diffused error, it tends to change the texture of images [Gerhardt and Hardeberg 2007]. As hinted in Norberg and Nyström [2013], increasing the number of primaries resolves these issues. This, however, renders the algorithm very inefficient. In our work, we introduce a highly efficient variant of sVED (§ 5.4) that uses almost one billion primaries made of contone stacks.

Painting Reproduction. Our hardware is very similar to the 2.5D printing introduced recently by Océ. This technology can be used in a straightforward manner for reproduction of oil paintings. In the study of Elkhuizen et al. [2014], experts are asked about the quality of such reproductions in a side-by-side comparison with the originals. According to this evaluation, while the recreation of geometry brings the reproduction to life, the color reproduction quality is low, the visibility of halftones is objectionable, the reproduced gloss is too uniform, and the translucency is missing. This system has also been used to control the gloss of a 2.5D print, in the context of fine art reproduction [Elkhuizen et al. 2015; Baar et al. 2015, 2016]. Recently, Elkhuizen et al. [2017] devised a system capable of measuring and printing the color, gloss and relief of paintings. Given the sensitivity of fine art reproduction to fine color differences, our work concentrates on high-quality spectral color reproduction; we leave the integration of color with other appearance attributes (gloss, relief and translucency) to future work. Regarding the color prediction of 2.5D prints, the work by Phan Van Song et al. [2016b; 2016a] has shown promising results. They rely on a four-flux estimation of the radiative transport equation [Rozé et al. 2001] in order to predict the spectral reflectance and transmittance of a multi-layer print, and the model is further extended to account for surface roughness [Phan Van Song et al. 2017]. In this work, we opt for a neural network to perform spectral reflectance prediction. As we show in § 5.2, our model outperforms the four-flux model, while not requiring explicit knowledge of the physical properties of materials.

Neural-Network-based Color Management. Neural networks have a long history of modeling non-linear ink-mixing behavior and learning the mapping between device control values and device-independent responses. Kang et al. [1992] apply a cascade correlation network [Fahlman and Lebiere 1990] to characterize 2D printer ink intensity (CMYK) and the output color in CIELAB space. Colorimetric accuracy is later improved by using multi-layer perceptrons, also known as fully-connected feedforward networks [Marcu and Iwata 1993; Abet and Marcu 1994; Drakopoulos and Subbarayan 2002]. Tominaga [1996; 1998] applies an encoder-decoder model to learn the same mapping bidirectionally within a single network. Xu et al. [2007] introduce a neural-network-based physical model that learns the mapping from inks' physical attributes to the output color. Littlewood et al. [2002] use neural networks for a color prediction model and formulate a Pareto-optimal problem to optimize the ink intensities jointly for colorimetric accuracy and additional user-defined objectives (i.e., usage of inks). Our proposed network builds on these models and extends them to work with spectral, not necessarily in-gamut input through a multi-objective loss function.




Fig. 2. *Left*: A photograph of one of our printed sample patches. Each color square is 1 mm \times 1 mm. *Right*: Our spectral acquisition setup.

3 OVERVIEW

We introduce a workflow for painting reproduction with unprecedented spectral accuracy, using multi-layer composition of different inks implemented via 3D printing. Given the difficulty of high-fidelity physically-based modeling of the spectral properties of ink stacks as well as the 3D printing process, we are motivated to address the problem using a data-driven approach. To this end, we design a spectral acquisition setup for accurate high-dynamic-range scanning of fine art paintings, as well as a printing setup with a selection of inks used for reproduction (§ 4). Equipped with the experimental setup, we develop our data-driven approach for modeling the spectral behavior of the ink stacks as realized by the 3D printing process (§ 5). We design a bidirectional prediction model between ink stack layouts and their corresponding spectra, and exploit our ink stack dataset to learn both directions of the model to enable faithful spectral reproduction. Our novel spectral vector error diffusion complements the “vertical” multi-layer ink stacks by “horizontally” modulating them with billions of possible primaries obtained through ink stacks. This combined approach of contouring and halftoning leads to high-quality reproduction of paintings, which we validate with an extensive evaluation and a variety of results (§ 6).

4 HARDWARE SETUP

In this section, we describe in detail our hardware setups for the spectral acquisition of original paintings and the printing process to realize the reproductions.

4.1 Spectral Acquisition

We use a spectral imaging setup as shown in Figure 2 to capture both calibration samples and paintings to reproduce. At top, the Nuance FX multispectral imaging system is coupled with a Coastal Optical 60 mm 1:4 UV-VIS-IR APO macro lens for multispectral image capture. Within the camera, a tunable liquid crystal filter can be dynamically adjusted to transmit one narrow range of wavelengths at a time, while the sensor captures a 1392 \times 1040-pixel monochromatic image of that spectral band. We perform high-dynamic-range (HDR) multi-spectral capture [Gkioulekas et al. 2013] within the spectral range of 420–720 nm, at 10 nm steps, yielding a 31-dimensional spectrum per pixel. Note that an alternative to the tuned LCD filter would be to use advanced hyperspectral recovery




Fig. 3. Our ink library. *Left*: The pigments and concentrations used in our inks, arranged in the order from top layer to bottom. *Right*: The spectra of saturated color of our inks printed with 30 layers. The color of each line corresponds to the color of an ink. The higher white line represents the Transparent White and the lower represents the Opaque White.

algorithms that can reconstruct the spectra using wide-band RGB channels [Baek et al. 2017; Choi et al. 2017].

At the bottom of the setup, we illuminate the sample with two 12-watt ROSCO Daylight (5800 K) 12 \times 12 inch (30 \times 30 cm) LED LitePads at 45 degrees to horizontal, centered symmetrically around the multi-spectral camera. The light emitted by each panel is diffused by a ROSCO #3029 diffusing sheet to maximize the lighting uniformity over the camera’s field of view. A black acrylic slab under the camera serves as a backing board to minimize reflection and indirect illumination. The entire setup is fully sealed in an enclosure to eliminate environmental illumination.

We correct radial distortion by photographing checkerboards and solving for radial distortion parameters using the MATLAB Camera Calibration Toolbox [Bouquet 2008]. We image X-Rite white balance reference and color-checkers to calibrate and compensate the light fall-off (vignetting), lighting non-uniformity, and spatial non-uniformity of pixel gain [Berns et al. 2015]. We normalize the spectrum of the illumination by dividing every multi-spectral measurement by the multi-spectral measurement of a calibrated X-Rite white reference target and multiplying by the reference’s spectrum.

4.2 Printing and Inks

We use MultiFab [Sitthi-Amorn et al. 2015], a laboratory-scale, multi-material inkjet 3D printer with a photopolymer printing process similar to commercial printers developed by Stratasys [2016] and 3D Systems [2013]. Our printer provides greater flexibility on the selection of printing materials. Its spatial resolution (xy -resolution) and vertical resolution (z -resolution) are 35 μ m and 11.25 μ m respectively. Our printer has five channels, enabling us to print with up to 15 inks using three passes of the printer. Our inks comprise a UV curable clear photopolymer carrier mixed with commercially available color pigments. We prepare the inks by first mixing the carrier and pigment and then milling the inks using a bead mill.

We have developed a set of 11 inks, including 9 color inks, a low-concentration white ink, and a high-concentration white ink (Figure 3). The color inks include the *process* cyan, magenta, yellow, and black (CMYK) inks. Following the n -ink 2D printers, we also add red, green, blue, orange, and violet inks to our printer’s palette. We emphasize that, though the orange at our disposal shares a

similar response with the red, it is more reflective in long wavelengths and has a narrower ($\sim 30\text{--}40\text{ nm}$) absorption band at short wavelengths. Given the importance and dominant use of orange pigment in art history and its reputation as difficult to be reproduced [Ostromoukhov 1993], we decided to include it. Inspired by the *veiling* technique used by painters, we add a low-concentration white ink, which we call transparent white, on top of the ink stack. The transparent white “veils” the highly saturated color inks beneath it and enables low-saturation colors. The high-concentration white ink, which we call opaque white, is the background of our prints and plays the same role as the paper in 2D printing. The transparent and opaque white use different concentrations of the same pigment.

The color ink concentrations are optimized for saturation, and are bounded by the hardware constraints of the printing process. In principle, we desire highly saturated colors with minimal ink-stack thickness. We start by finding, for each individual ink, the maximum concentration that the printer can properly print. Exceeding these concentrations, causes unstable jetting, printhead clogging and incomplete UV curing. Then, we experimentally determine the maximum number of layers required to obtain saturated colors at their highest concentration. This number depends on the inks and for a majority of the more transparent inks is around 30 layers. We adjust the concentration of the remaining inks such that they all saturate at 30 layers. For example, for some inks, such as violet and black, the color of the resulting stack saturates more quickly. Therefore, we decrease the concentration such that they saturate at 30 layers.

We opt for the highest concentration because when ink concentration increases, the number of layers required to obtain saturated colors decreases. A lower number of layers is preferred because the mechanical and optical dot gain, and therefore blurring, increase with the number of layers [Babaei et al. 2017]. On the other hand, highly concentrated inks cause pronounced quantization artifacts. We partially address this challenge by incorporating veiling top layers using the transparent white, and therefore the concentration of that ink is tuned to allow the colors beneath to remain visible, although it is handled in a more principled way by spectral vector error diffusion (§ 5.4). Finally, we set the concentration of the opaque white to be as high as possible. This prevents subsurface scattering and results in a brighter white substrate.

5 MODELING SPECTRAL REPRODUCTION

In this section, we describe our approach to computing a layout of inks that faithfully reproduces a given spectral reflectance. The bidirectional relation between the spectral reflectance and the ink-stack layout is modeled using a pair of neural networks, one for each direction, which are learned from a dataset consisting of printed contone-stack layouts. The spatial modulation is implemented using a spectral vector error diffusion, which also serves a means of layout discretization. We first introduce our contoning dataset (§ 5.1), which we use to learn our neural-network-based *forward* model (§ 5.2), capable of predicting the spectrum of a stack of inks. The forward model, in turn, is used to train our *backward* model

(§ 5.3) that predicts an optimal ink-stack layout for a given spectrum. Finally, we describe our spectral vector error diffusion (§ 5.4), which is followed by practical considerations and details (§ 5.5).

5.1 Ink-Stack Dataset

We prepare a dataset of contone ink-stacks to train our model. Throughout this work, the order of inks in each stack is fixed: a section of 30 layers consisting of any inks, arranged from the most transparent to the least transparent starting from the top (in the order shown in the table of Figure 3), followed by additional 20 opaque white layers at the bottom. Although we fix the order of inks in the ink stacks, the number of layers and that of inks lead to about 0.8 billion possible layer layouts. To efficiently sample from all layer-layouts to build our training dataset, we introduce the following heuristic sampling rules to maximize the coverage of the spectral gamut while keeping the number of samples manageable. In the following, the “color” inks include the transparent white as well, i.e., denote the first ten inks in the table of Figure 3, while still excluding the opaque white ink.

- Due to the smooth change of spectral properties of traditional inks, combinations of a large number of different inks does not result in large spectral variety [Rosen et al. 2004]. We thus limit the maximum number of different color inks in a stack to 5.
- We limit the number of color ink layers depending on the number of different inks used in the stack. Specifically, when using 3, 4, or 5 color inks, we set the maximum number of color layers to 10, 15, or 20, respectively. The remaining layers are complemented with additional opaque white layers at the bottom.
- If an ink stack contains 2 or 3 “dark” colors, we further limit the number of color layers to 8 or 4, respectively. Dark colors are the colors that have high absorption and/or scattering, which in our case are blue, violet, and black inks.
- For any layer-layout consisting of more than 2 color inks, we enforce the layer increment step of an ink to be 2.

We obtain 20,878 valid layer layouts following these rules. Each ink stack occupies a $1\text{ mm} \times 1\text{ mm}$ square in the calibration print. We aggregate 13×16 ink stacks into one calibration patch, resulting in the total of 101 patches (Figure 2, left). Within each patch, the neighboring ink stacks are separated by a 0.3 mm wall made of transparent material to prevent pixel cross-talk. We favor transparent wall over black pigment wall for both better curability and free of color contamination, with the black backing further minimizes reflection, resulting negligible sub-surface scattering. To ensure measurement accuracy, the spectral reflectance of each ink stack is calculated by averaging the measurement over the $0.5\text{ mm} \times 0.5\text{ mm}$ central region of each square.

Based on the maximum operation area of our printer, we print 4×4 patches at each printing iteration. Printing proceeds in 3 passes, and is done upside down—i.e., the topmost layers are printed first, in order to achieve a diffuse surface finish. In the first pass, the top white layers and the transparent wall are printed and cured. In the next two passes, the first 5 color layers and the next 5 color layers are sequentially printed. Completion of all three passes takes




Fig. 4. Illustration of our network structure. F denotes our spectral prediction model (§ 5.2), which is used to estimate the spectrum of the predicted layout by our layout prediction model B (§ 5.3) and provides the means to measure the quality of it. Q represents our soft quantization layer (§ 5.4). ψ denotes the spectrum and ξ the layer-layout vector. Both F and B are implemented as fully-connected feed-forward neural networks.

approximately 3.5 hours. The HDR multi-spectral acquisition takes 5 minutes to complete measurement for one calibration patch.

5.2 Spectral Prediction Model

Spectral reflectance is represented as a 31-dimensional real vector, each element denoting the reflectance at a discretized wavelength. A layout is represented as a 11-dimensional non-negative integer vector, encoding the numbers of layers used for each ink. We learn the spectral prediction function $F: \mathbb{Z}_+^{11} \rightarrow \mathbb{R}_+^{31}$, which maps a layout ξ to its spectral reflectance ψ , from our ink-stack dataset (§ 5.1). Given the low dimensionality of our problem, we model the function using a straightforward fully-connected feed-forward neural network (also known as a multi-layer perceptron) [Cybenko 1989]. Our model includes 4 hidden layers, each with 300 neurons (hidden units), forming a network that is wider than it is deep. Both the hidden layers and the final output layer use the rectified linear unit (ReLU) as a nonlinear activation function.

5.2.1 Spectral Loss. We optimize the network for the minimal spectral error, which is defined as a Euclidean distance between the prediction and the measurement (scaled by the square root of the number of wavelength bands, in accordance with the practice in colorimetry literature [Imai et al. 2002]):

$$E_{\text{spec}}(\xi, \psi^*) = \frac{1}{\sqrt{31}} \|F(\xi) - \psi^*\|_2. \quad (1)$$

In our implementation, we normalized the input ink layout vector by the total number of layers (i.e., 30) such that each layout sums to one.

5.3 Layout Prediction Model

Ultimately, we are interested in the inverse function of F , through which we retrieve the optimal layout for a given spectral reflectance. Given that F is a differentiable function that maps layouts to spectra, a straightforward approach for the inverse problem would be fixing F 's parameters and running gradient descent on layouts by minimizing the spectral error. However, the distribution of spectra over layouts will likely be highly non-uniform and multi-modal. Further, as the layout is inherently discrete, its relationship with the spectrum forms a non-convex space. In practice, we find that

running gradient descent directly on a randomly initialized layout usually leads to a local minimum and results in a poor prediction.

To pick a good initial guess, reducing the risk of being trapped in a local minimum, we could train a separate network that would learn the inverse mapping of F and provide a reasonable initial layout estimate from a target spectrum, from which we could initiate the gradient descent through F . However, it is difficult to measure how good one predicted layout is compared to another, since the L^2 difference between layouts does not translate directly to the spectral difference, in particular when different layouts could lead to an identical spectrum. A better alternative is to construct an encoder-decoder model that uses F to “decode” the predicted layout to the spectrum and evaluate the error with respect to the target input spectrum; i.e., to use the forward network F (decoder) to set up the loss and train the reverse mapping (encoder) through back-propagation directly [Tominaga 1996] (Figure 4).

We define a layout prediction function $B: \mathbb{R}_+^{31} \rightarrow \mathbb{Z}_+^{11}$ that predicts an optimal layout given a spectrum and is implemented by another fully-connected neural network. B consists of 8 hidden layers, deeper than F , as it attempts to learn a more complicated distribution. Each hidden layer consists of 160 neurons with the ReLU activation function. As for F , the layout vector is normalized. A softmax layer is applied to the output layer to guarantee the validity of the layout estimate—nonnegative and summing to one.

5.3.1 Perceptual Color Loss. We first experimented with training B with the spectral loss (1) only. However, we observed that the reproduced spectral reflectance (from the predicted layout) of an out-of-gamut target spectrum varied noticeably across trained models, although all shared a similar L^2 error. Since we do not explicitly perform spectral gamut mapping nor apply any gamut constraints, the trained model was allowed to freely select any in-gamut spectrum minimizing the L^2 error. We, however, wish the network to predict the one among all solutions that also minimizes the perceptual color difference under a variety of standard illuminants and common light sources [Morović et al. 2012]. To implement this, we define a perceptual color loss over chosen illuminants:

$$E_{\text{LAB}}(\xi, \psi^*) = \sum_i \Delta E_{ab}^* (\text{LAB}_{I_i}(F(\xi)), \text{LAB}_{I_i}(\psi^*)) , \quad (2)$$

where $\text{LAB}_{I_i}(\cdot)$ represents a series of transformations from a spectrum to CIEXYZ coordinates under the illuminant I_i with the CIE 1931 2° standard observer, and further to CIELAB color. Each I_i denotes the spectrum of a particular illuminant, and for ΔE_{ab}^* we use the CIE 1976 color difference (ΔE_{76}) in CIELAB space [Wyszecki and Stiles 1982]. For this perceptual loss, we used 8 different light sources: halogen, incandescent, fluorescent (2700 K), Philips Hue LED (5000 K), ROSCO Daylight LED (5800 K) as well as the CIE standard illuminants D65, D50 and A. The ΔE_{76} is chosen for simplicity; ΔE_{94} , also differentiable, could as well be an accurate perceptual loss metric [Kauvar et al. 2015]. Although more accurate, ΔE_{00} is not differentiable.

We observe more stable and consistent training and preferred results with the addition of the perceptual loss (2). We emphasize that, for in-gamut input spectra, adding the perceptual loss will *not*

misguide the network to learn metameric pairs since only the desired spectra drive both the spectral loss and perceptual loss to zero. However, training with only the perceptual loss also yields suboptimal results, since the mapping from a low dimensional CIELAB space to a high dimensional spectral space is multi-modal. Without the measurement of spectral fidelity and under a set of limited illuminants only, the network can easily converge to a metameric spectrum that gives an equivalently low perceptual difference.

5.3.2 Layer Thickness Loss. In addition to minimizing the perceptual difference, we also minimize the number of color layers (color inks plus transparent white) in a layout estimate (excluding the twenty or more *opaque* bottom white layers). Babaei et al. [2017] show that fewer color layers result in smaller minimum feature size and reduced optical and physical dot gain in color contouring. In both 2D- and contouring-based colorimetric reproduction, the mapping from a CIELAB color to CMYK ink intensity is multi-modal, which is a result of the trade-off between the black ink (K) and three chromatic inks (CMY). Our ink set also shares this problem, and the layout with the thickest black layer often results in the minimal total layer thickness (i.e., the number of color layers). In addition, given the significantly increased number of primaries, a target spectrum may also be reproduced by combination of different inks. Therefore, to encourage the use of a minimal number of color inks, we define the layer thickness loss:

$$E_{\text{thick}}(\xi) = \sum_{i=1}^{10} |\xi_i| = \|\xi\|_1 - |\xi_{11}|, \quad (3)$$

where ξ_{11} is the number of opaque white ink layers in the predicted layout.

The backward function **B** is trained to minimize the following loss over our training dataset:

$$E = E_{\text{spec}} + \alpha E_{\text{LAB}} + \gamma E_{\text{thick}}, \quad (4)$$

with $\alpha = 10^{-3}$ and $\gamma = 10^{-3}$.

5.4 Spectral Vector Error Diffusion

When using the backward model in order to print an ink layout best reproducing a given spectrum, we witness two problems which, at first glance, seem independent. First, a denormalized layout prediction of **B** has continuous values and thus has to be discretized to integer layers for actual printing. Second, because of the high concentration of our inks, a mere contour-based model would suffer from color quantization [Babaei et al. 2017]. Although the use of transparent white mitigates this problem, quantization artifacts are still visible at regions with smooth color gradients. This is in essence a halftoning problem: we have a spatially-varying continuous input (from **B**) that needs to be represented with a discrete ink layout at each location. Halftoning techniques simulate continuous tone images through spatial modulation of a limited number of primaries [Baqai et al. 2005].

5.4.1 Layer Layout Discretization. A straightforward approach would be to round the continuous output to the nearest integer. Alternatively, we can enumerate all possible combinations of rounding up and down for the prediction on the thickness of each ink,

which would amount to 2^{11} possibilities, and search for the best combination, in terms of spectral accuracy, in a brute-force manner.

Since the complexity of enumeration grows exponentially, it would be preferable if the network could directly generate integer-valued layouts. However, rounding is not differentiable, and thus cannot be directly used during training. To address that, we apply a soft quantization layer to encourage the network to predict close-to-integer layouts, mitigating the influence of a posterior rounding operation on the prediction. Our soft quantization layer is placed between the spectrum prediction network **F** and the layout prediction network **B**, such that the loss is calculated on (soft) rounded layout predictions. The soft quantization layer is constructed as a sum of cascaded soft unit-step functions:

$$Q(\xi_i) = \frac{1}{N_s} \sum_{k=1}^{N_s} s(N_s \cdot \xi_i - (k-1)), \quad (5)$$

where the number of steps equals to the number layers $N_s = 30$, and the unit-step function $s(\cdot)$ could be any sigmoid function that has a steep transition around a half and has vanishing gradients outside the unit interval. We use a shifted logistic function:

$$s(\xi) = \sigma(\kappa \cdot (\xi - 0.5)), \quad (6)$$

where $\sigma(x) = (1 + \exp^{-x})^{-1}$ and κ controls the steepness of the transition, which we set $\kappa = 15$ for training our network.

5.4.2 Error Diffusion. Since in contouring, the only way of modulating color is to add or remove a layer, when highly-concentrated inks are used, the color space is not covered as densely as required, especially at low-saturation areas of the gamut. Our solution to the color quantization problem is to perform an error diffusion with contour stacks as primaries. This bears resemblance to spectral vector error diffusion (sVED) algorithms. However, our proposed method is significantly more efficient. Unlike sVED algorithms that for each pixel look for the nearest neighbor among all primaries, which in our case spans 0.8 billion 31-dimensional points, we simply round the continuous layout predicted by **B** and feed the rounded layout to **F**. The resulting spectrum is subtracted from the target spectrum to obtain the spectral error vector. The spectral error is then carried over to the neighborhood pixels, wavelength by wavelength, according to the weights specified by the Floyd-Steinberg [1976] error diffusion kernel. The subsequent prediction for the neighborhood pixels then takes into account the carried-over spectral error. To further improve the computational efficiency, we parallelized the error diffusion following the optimal scheduling proposed by Metaxas [1998].

The advantage of this approach is twofold. First, the discretization of the continuous layout from the backward model can be a simple rounding, which is orders of magnitude faster than a search among the huge space of possible ink stacks. More importantly, the color quantization artifacts are removed through spatial mixing of a large number of contour-stack primaries. Unlike with RGB or CMYK dithering, the extremely large pool of primaries makes halftoning artifacts virtually invisible, since for any input spectrum there exists a very similar “primary.”

Table 1. The accuracy of different spectral reflectance prediction models, in terms of the spectral error E_{spec} (Eq. 1) and the CIEDE2000 (ΔE_{00}) [Luo et al. 2001] under three illuminants. We report the mean, standard deviation, median, and maximum. The lowest errors are denoted in boldface.

Models	E_{spec}		ΔE_{00} (D65)		ΔE_{00} (TL84)		ΔE_{00} (A)	
	Mean	SD	Mean	SD	Mean	SD	Mean	SD
	Median	Max	Median	Max	Median	Max	Median	Max
Contoning	6.91	2.82	9.08	5.67	9.13	5.53	9.11	5.20
	5.48	61.73	8.56	41.49	8.56	41.73	8.65	40.77
1C-KM	4.90	2.82	9.89	5.67	9.02	5.53	8.26	5.20
	4.22	25.73	9.12	43.37	8.14	34.37	7.31	35.33
2C-KM	3.83	3.09	6.64	4.03	6.13	3.96	5.72	3.93
	2.99	34.29	5.85	39.79	5.24	40.10	4.80	40.80
4-Flux	1.99	1.60	3.57	2.84	3.70	2.85	3.54	2.70
	1.60	27.40	2.92	30.45	3.03	30.87	2.93	27.46
Ours	1.44	1.06	2.50	1.58	2.38	1.58	2.20	1.69
	1.16	10.95	2.19	11.79	1.97	11.36	1.72	12.24

5.5 Training Protocol

We randomly split (and fixed) our spectral ink-stack dataset into a training set of 18,878 samples and a test set of 2,000 samples. We verified that the number of layers for each color ink in the test set follows the distribution in the entire dataset.

We train our spectral prediction network F using the spectral loss (1) only. After the training of F is finished, we train the layout prediction network B using the combined loss (4), with the weights of the trained F fixed. During the reproduction, a target multi-spectral image is fed to B to yield an optimal layout prediction, which subsequently is quantized by the spectral error diffusion. The resulting (denormalized) layout map consists of 11-dimensional vectors, each of which dictates how many layers of each ink have to be deposited in order.

Both networks are trained using Adam [Kingma and Ba 2014], a stochastic optimization algorithm. We use an initial learning rate of 10^{-3} for F and 10^{-4} for B , with $\beta_1 = 0.9$ and $\beta_2 = 0.999$ for both models. Training samples are drawn randomly from the training dataset. Both models are trained for 500,000 iterations with batch size of 64. A learning rate decay of 0.1 is employed every 50,000 iterations. Weights are regularized by penalizing their L^2 norm weighted by 10^{-5} .

6 RESULTS AND EVALUATIONS

We evaluate the performance of our method both quantitatively and qualitatively. In addition to individual spectra that have been historically considered important to reproduce, we validate our full reproduction pipeline with a number of real paintings, which are intentionally painted to maximize the diversity in color, and include challenging textures and smooth gradients.

6.1 Model Evaluation

In this section, we evaluate our spectral and layout prediction models, and spectral vector error diffusion. We also provide the reproduction quality of our model tested on historically important colors.

6.1.1 Spectral Reflectance Prediction Model. We compare our spectral reflectance prediction model with several physically-based and data-driven models: the one-constant Kubelka–Munk [Kubelka

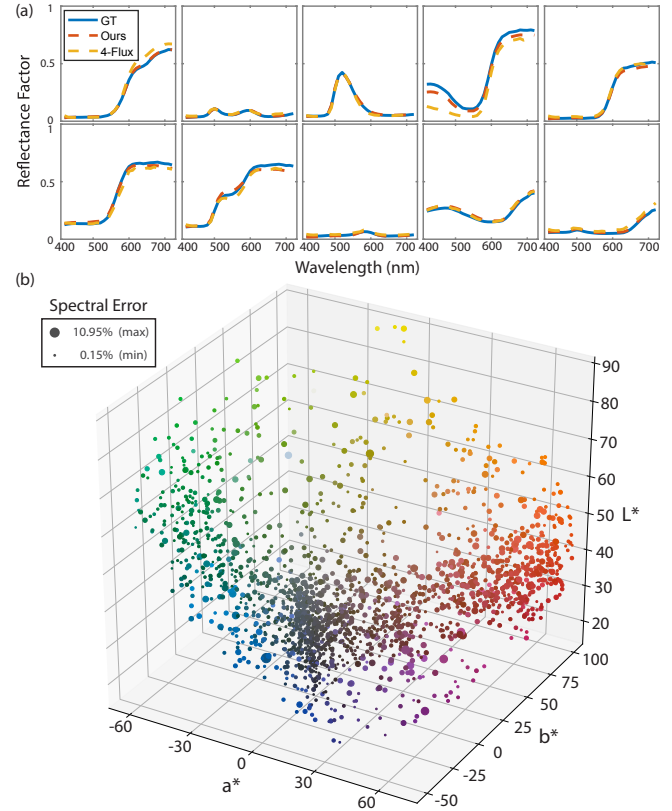


Fig. 5. (a) Comparison of test spectra from our contoning ink stack dataset predicted by our spectral prediction model F and by the 4-flux method. (b) Projection of predicted spectra onto CIELAB space (under D65) with the dot volume proportional to the spectral error E_{spec} .

and Munk 1931] and two-constant Kubelka–Munk models [Davidson and Hemmendinger 1966] (1C-KM and 2C-KM, respectively); the four-flux model [Rozé et al. 2001]; and the color contoning model [Babaei et al. 2017]. We follow different strategies for building these models. For 1C-KM model, where the model parameters (absorption to scattering ratio at each wavelength) are straightforward to compute, we use measurements. For more advanced models, i.e., 2C-KM and four-flux, since measuring their parameters is prone to error, we fit their parameters using our training set and evaluate the model accuracy on the test set. The fitting procedure ensures the maximum capacity of these models as it does not rely on physical measurements. For training data-driven models (all except 1C-KM), including ours, we use the training samples of our dataset (§ 5.1).

In Table 1, we show the spectral and colorimetric accuracies of all models. We report the statistics of the spectral error and ΔE_{00} under three illuminants, D65, TL84, and A. Under all metrics, our model achieves a significantly lower prediction error, and is on average close to the measurement noise (1%), which was measured through repeated scanning. In Figure 5(a), we show the quality of spectral prediction with a number of randomly-drawn test samples of varying spectra, where our prediction closely follows the

Table 2. The accuracy of our layer prediction model trained with different objectives, measured in the spectrum error E_{spec} (Eq. 1) and the CIEDE2000 ΔE_{00} under three illuminants. We report the mean and standard deviation. The lowest errors are shown boldface and the second lowest underlined.

Losses	E_{spec}		ΔE_{00} (D65)		ΔE_{00} (TL84)		ΔE_{00} (A)		thickness	
	Mean	SD	Mean	SD	Mean	SD	Mean	SD	Mean	SD
(a) E_{spec}										
–	0.81	0.49	1.71	1.28	1.84	1.54	1.52	1.53	16.20	3.32
rounded	<u>2.28</u>	2.39	3.75	2.41	3.64	2.59	3.17	2.39	15.73	3.36
(b) $E_{\text{spec}} + E_{\text{LAB}}$										
–	0.81	0.50	1.20	<u>0.93</u>	1.37	<u>1.17</u>	1.17	1.18	16.57	3.16
rounded	2.32	<u>2.31</u>	<u>3.44</u>	<u>2.32</u>	<u>3.25</u>	<u>2.24</u>	<u>2.94</u>	2.10	16.13	3.28
(c) $E_{\text{spec}} + E_{\text{LAB}} + E_{\text{thick}}$										
–	0.82	0.49	<u>1.34</u>	0.91	<u>1.50</u>	1.15	<u>1.31</u>	<u>1.20</u>	14.43	3.17
rounded	2.52	2.38	3.71	2.38	3.57	2.34	3.13	<u>2.07</u>	14.03	3.26
(d) $E_{\text{spec}} + E_{\text{LAB}} + E_{\text{thick}}$ w/ soft quantization layer										
–	1.20	0.71	2.02	1.27	2.16	1.44	1.82	1.36	<u>14.72</u>	3.20
rounded	1.60	1.09	2.70	1.62	2.65	1.61	2.31	1.53	<u>14.80</u>	3.23

measured spectra. In Figure 5(b), we plot the predicted test set spectra (2000 samples) in CIELAB space with dot volume proportional to the spectral error. In the supplementary material, we provide additional visualization of equivalent 2D plots with 9 different luminance bins. The standard deviation of mean spectral error at each quadrants is 0.10%, showing no obvious bias towards towards a particular tone. The samples at high luminance region are more sparsely populated, likely due to the use of high concentration inks. Using more layers with lower concentration shall provide denser coverage at the cost of more severe blurriness.

The merit of physically-based models is that they rely on only a limited physical measurements. Therefore, they are scalable to a large number of inks and do not require a large dataset and lengthy training process required by data-driven models. While they usually work well within a limited operational range, these models have difficulties to incorporate the added complexity of the printing process. The contouring model, which is also a data-driven method, assumes only absorbing materials and is unable to predict the spectral properties of our complex stack that includes different inks with a wide range of absorption and scattering.

6.1.2 Layer Prediction Model. In order to evaluate our layout prediction model, we compare the performance of the model trained with varying losses (Table 2). We report the errors using the same metrics as in the spectral prediction model, measured from the undiscretized network output and the discretized output, which has been rounded to the nearest integer.

With a negligible difference on the achieved spectral error, adding perceptual loss and layer thickness loss results in roughly 30% reduction on the perceptual difference, and one layer reduction on the total color layers. However, we choose to apply small weights to the perceptual loss and layer thickness loss, since a high perceptual loss weight results in a metameric reproduction. A high weight for layer-thickness loss induces high quantization errors. Depending on the application and number of inks, weights for perceptual loss and ink thickness loss may be further tuned.

The layer prediction model could be directly trained without the pre-trained spectral prediction model as a loss, where the loss itself

would also be learned through, e.g., a conditional generative adversarial network [Isola et al. 2017]. We experimented with a model consisting of our **B** as a generator and a fully-connected feedforward network with a 2-way softmax at the end, resembling our **F** but performing binary classification, as a discriminator, trained with our loss terms. We found, however, that a such-trained model was prone to overfitting with poor generalization.

6.1.3 Spectral Vector Error Diffusion. We evaluate the performance of our proposed soft quantization layer and sVED. Figure 7 presents the simulated reproduction of a painting under D65 illuminant using varying configurations. Our soft quantization layer helps reduce the errors introduced by rounding (e vs. f). While the local (per-pixel) brute-force enumeration (d) helps alleviate artifacts, the use of our spectral error diffusion greatly improves overall reproduction quality (c). Combined with the brute-force search (b), the sVED shows the best results, with less noticeable graininess, which can be best seen in the electronic version of the paper when zoomed in. Yet the sVED on both the brute-force enumeration (b) and nearest-integer rounding (c) exhibit comparable quality, but with rounding exhibiting better run-time performance (see § 6.3). We note that, at high luminance and region with smooth gradient (last row), the halftone pattern becomes more visible, likely due to the relatively sparse coverage over the particular gamut area and the high color contrast.

The effectiveness of the soft quantization layer is validated quantitatively in Table 2. The reproduction error of soft-quantized results before rounding (d; first row) is slightly higher than that of the results without the soft quantization layer (a–c; first rows each). However, while the error increases sharply after rounding without the soft quantization layer (a–c; second rows), there is a much smaller increase for the soft-quantized results (d; second row).

6.1.4 Color Gamut Evaluation. We compute the color gamut volume of our proposed 10-ink setup in CIELAB color space under D65 illumination. We densely sample 16.8 million layouts, including all possible layouts for stacks with up to 5 inks as well as layouts for stacks with more than 5 inks at a minimal increment step of 2. The layouts are fed to the our spectral prediction model **F** to obtain the resulting spectra. We reconstruct the non-convex gamut surface of our CIELAB point cloud using the ball-pivoting algorithm [Bernardini et al. 1999]. We then compute the volume enclosed by this surface, which for our 10-ink setup is 79% of the sRGB gamut volume. This may be compared to 42% for a standard CMYK 2D printer based on halftoning, 45% for a typical CMYK contouring system, and 65% for a fluorescent-ink printer (CMYK plus two fluorescent inks). We observe that our gamut is significantly larger than the fluorescent-ink gamut, the state-of-the-art method for wide-gamut printing [Rossier 2013]. In Figure 6, we visualize the color gamut comparisons over several iso- L^* slices in CIELAB. As shown in the figure, our 10-ink setup provides additional gamut coverage over color contouring most prominently at green, red, orange and pink tone regions, particularly when the luminance is high. To account for the sRGB gamut, further incorporation of blue and purple tone primaries would be necessary.

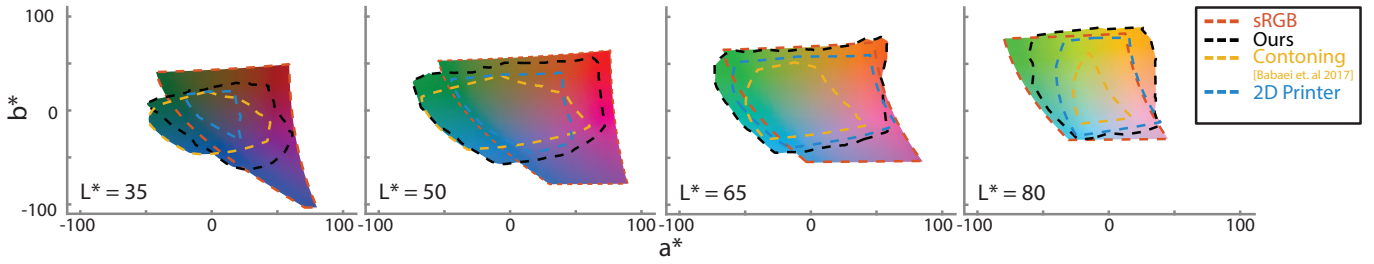


Fig. 6. Comparison of our 10-ink gamut with two other printing methods and sRGB (under D65). The figure shows a^*b^* slices at different L^* values.

6.1.5 Reproducing Historical Pigments. We evaluate the ability of our proposed model to reproduce real pigments that have been widely used for artwork. We use the FORS spectral database, which contains 54 pigments of historical interest [Cosentino 2014]. This dataset contains the spectra of the pigments in pure powder form, as well as mixed with different binders (6 sets in total). Among these, we reproduce the measured spectra of pigments mixed with gum arabic, which gives the most saturated colors.

In Figure 8, we compare the spectra and the resulting color under the D65 illuminant predicted by color contoning [Babaei et al. 2017] and our model (with rounding enumeration), to the exhaustive search over 16.8 millions of densely sampled spectra which we used to compute the contone color gamut (§ 6.1.4). The color contoning achieves $E_{\text{spec}} = 4.02\%$ and $\Delta E_{00}(D65) = 5.64$ on average, while our network achieves $E_{\text{spec}} = 2.60\%$ and $\Delta E_{00}(D65) = 4.77$, compared to $E_{\text{spec}} = 2.31\%$ and $\Delta E_{00}(D65) = 2.39$ of the exhaustive search. The spectral error and LAB error of individual pigments are reported in the supplementary material Table 1. For most in-gamut spectra, our network prediction yields results very close to (or indeed identical to) exhaustive search over 16.8 million samples.

We identify 12 challenging (out-of-gamut) spectra that are not well reproduced by our ink set ($E_{\text{spec}} > 5\%$ or $\Delta E_{00}(D65) > 7$ when predicted by our model). For example, the cobalt blue (row 4 at column 8 in Figure 8; R4C8 hereafter) and smalt (R5C4) share a steep tail after 700 nm, and similarly, the cobalt violet (R5C8) possesses a steep head before 450 nm, all of which cannot be reproduced by our violet without causing a fat tail from 630-700 nm or a fat head from 450-500 nm. The fat head also appears in the carmine lake (R2C3), whose ramp starts from 600 nm and is 50 nm behind the spectra of our red, orange and magenta. The cadmium yellow (R3C7) starts the ramp at 450 nm, which is 30 nm earlier than our saturated yellow in and causing a lifted yellow spectra (with white layers) to match the curve. The downhill of lithopone (R6C9), which starts at 630 nm, also cannot be reproduced, given that none of our inks' spectra has a descending trend after 600 nm. Overall, our current ink set has difficulties in reproducing yellow- and gray-tone pigments. Ultimately, we believe a further expanded and more carefully selected ink set is necessary to reproduce all pigments' spectra well.

6.2 Painting Reproduction Evaluation

We scanned small oil paintings with varying scales, color and texture characteristics, using the same spectral acquisition setup we

used for measuring our dataset. During capture, the distance between the camera and painting is adjusted to match the measurement resolution to the printer's spatial resolution (35 μm). For paintings larger than the camera's field of view, we scan them part-by-part and stitch the measurements into a single spectral image [Brown and Lowe 2007].

We captured the photographs of all results reported in the paper using a Canon 5D Mark-III DSLR camera with a Canon 100 mm f/2.8 macro lens. The photos of the paintings and their reproductions were taken simultaneously in a single shot under the same light source, then cropped appropriately. We used a fluorescent light and a halogen light to represent cold and warm color temperatures. Additionally, we used a Philips Hue light, which includes different colors of LEDs inside to simulate different lighting.

We include the high-resolution images of both simulations and physical printings in our paper, and refer the readers to the electronic version of it, where more details of our reproduction can be seen and examined when zoomed in.

In Figure 9, we show the side-by-side comparisons of the selected paintings patches and our printed reproductions under different physical light sources. The physical dimension of each painting patch is 3.2cm \times 2.4cm. Our results faithfully reproduce a variety of color used in real paintings under a wide range of color. In particular, the paintings exhibits different characteristics: the "water lily" is characterized by its desaturated colors and smooth gradients, while the others use more saturated colors.

We compare our reproduction with the color contoning method [Babaei et al. 2017] in Figure 10, using simulation. Our sVED based on 10-ink stacks outperforms the CMYK contoning method with more faithful reproduction of various colors. This is expected as our gamut volume is significantly larger than the color contoning gamut. In contoning, the ink concentrations are kept low to prevent color quantization. Our use of more saturated inks help achieve a larger gamut at the cost of more prominent color quantization, which is handled by our spectral vector error diffusion method (the halftoning pattern can be examined in the electronic version when zoomed in).

In Figure 11 we show the printed reproductions of two paintings generated using our 10-ink pipeline, and the same pipeline but with our CMYK inks only. Thanks to our sVED method, which allows for high-concentration inks, we can achieve a significantly large color gamut even with CMYK. However, although our highly-concentrated CMYK inks yield a large gamut close to a fluorescent-ink printer (65% of sRGB), the 10-ink method shows consistently

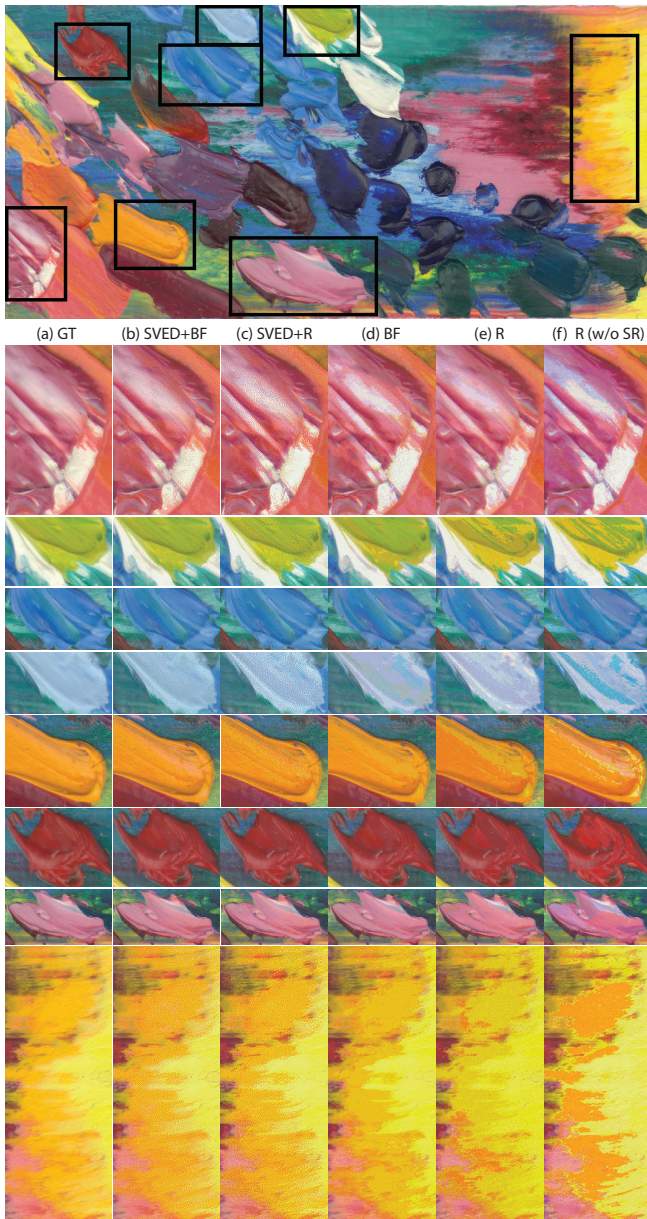


Fig. 7. Comparisons of different quantization methods. (a) ground truth; (b) spectral vector error diffusion with brute-force enumeration; (c) spectral vector error diffusion with nearest-integer rounding; (d) per-pixel brute-force enumeration; (e) per-pixel nearest-integer rounding; (f) per-pixel nearest-integer rounding *without* using our soft quantization layer. The results in (b–e) were created using the network *with* the soft quantization layer. Each row at bottom corresponds to a rectangular region in the painting at top. Paintings ©Azadeh Asadi.

better reproduction quality under different lights. This justifies the use of a larger number of inks for high-fidelity spectral reproduction. For a quantitative evaluation, we also show the spectra of some sampled points from the original painting, the reproduction by our 10-ink method, and the reproduction by our CMYK inks, along

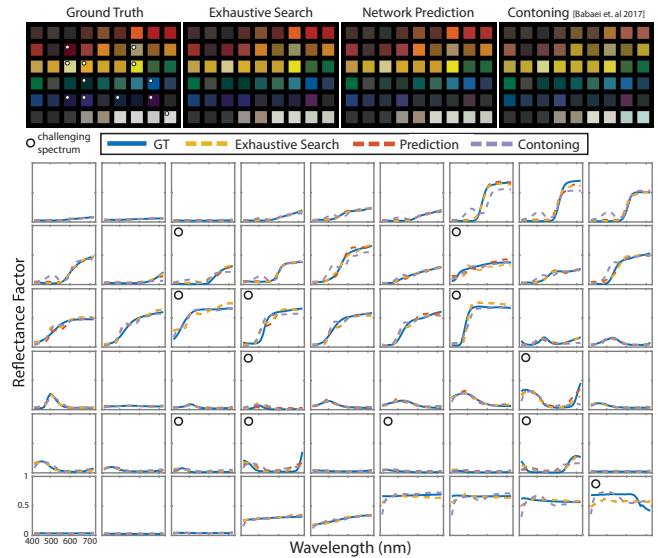


Fig. 8. Comparison of historical pigment spectra reproduced by the proposed neural network and the nearest neighbour search. The white dots mark the challenging spectra that are difficult for our ink set.

with spectral and colorimetric errors. In general, the CMYK set, even though at high concentration, is not able to reproduce some of spectra faithfully (on average $E_{\text{spec}} = 6.22\%$ and $\Delta_{00} = 7.97$). Our 10-ink printer, on the other hand, gives excellent spectral match between the paintings and prints (on average $E_{\text{spec}} = 2.78\%$ and $\Delta_{00} = 3.80$). Note the spectral difference in the inset of Figure 11, where the 10-ink set consistently gives less spectral error. Concerning colorimetric reproduction, the 10-ink set outperforms or does equally well compared to the CMYK. There are a few exceptions, especially out-of-gamut spectra, such as spectrum number 3 in the second painting, a brilliant red color. Although the color reproduction accuracy of the CMYK set is relatively good, it has unacceptable spectral accuracy, indicating the risk of a metameric reproduction. On the other hand, the 10-ink reproduction has much better spectral reproduction accuracy, while the colorimetric error is high, typical for out-of-gamut colors.

The reproduced paintings have to undergo a series of color transformations (e.g., by the camera, publishing software, and the printer or display) to be presented in this paper, whether it is printed or viewed electronically. Thus, the difference between our reproductions and those by other methods may have been washed out or become subtle. To further provide evidence of the significant differences between them when seen in person, we carried out a small perceptual study, in which we compared our 10-ink with our CMYK prints. We asked participants to evaluate the color fidelity of the two candidate methods as compared to the original, under three different lighting conditions (fluorescent at 3500 K, LED at 5800 K, and cloudy daylight). We asked participants to evaluate printed patches on a scale between -3 , indicating a strong preference *against* ours, and $+3$, *for* ours, with 0 being indifferent. Based on a study involving 8 participants, the mean preference score was $+1.69$ with




Fig. 9. Reproduced paintings using our method. We show 4x2 sets of our printed reproductions. For each set, we show our printed reproduction with CMYK inks (top row), the original painting (middle row), and our printed reproductions with 10 inks (bottom row). Each painting is lit by four different light sources (halogen (2750K), Phillips Hue LED (Warm), fluorescent (5000K), and Phillips Hue LED (Cool)), which are shown so ordered. The spectrum of each light is illustrated in the supplementary material Figure A.1. Paintings ©Azadeh Asadi.




Fig. 10. Simulated comparisons under D65 illumination of our method to the CMYK contoning. Our reproduction (a) of three paintings are compared against the original painting (b) and the CMYK contoning method (c) of Babaei et al. [2017]. Our method provides consistently more faithful reproductions to the original. Paintings ©Azadeh Asadi.

standard deviation of 1.14, which shows a consistent and statistically significant ($p < 0.001$) preference for our reproduction.

6.3 Run-time Performance

Our model is implemented using TensorFlow [Abadi et al. 2016], and trained and tested on an NVIDIA Titan X (Pascal) GPU. It takes about 15 minutes to train the spectral prediction network F and about 45 minutes for the layout generation network B . Note that these are one-time computations and once the model is built, it can be used for different input paintings. To predict the layout of a painting consisting of one million spectral pixels, it takes on average 0.3 seconds.

Our spectral vector error diffusion is implemented in Python (with NumPy) and TensorFlow. We follow the optimal scheduling pattern proposed by Metaxas [1998] to batch the evaluations of uncorrelated pixels at each time step, which significantly reduces the number of network inference calls. During error propagation, a damping factor between 1 and 0.6 (for highly saturated paintings) is applied to the error diffusion kernel to prevent the algorithm from diverging. Error diffusion on a painting of one million spectral pixels takes on average 23 seconds using simple rounding and 590 seconds using brute-force rounding enumeration.

For physical spectral prediction models, although the forward model can be very efficient, the backward model needs a model inversion with iterative optimization algorithms, which renders the print-data generation into a slow process. We tested the color contoning run-time performance using our training dataset. The forward model takes more than 40 hours in MATLAB on a MacBook Pro with a 2.8 GHz i7 processor and 16 GB of RAM. This is because the model inverts a large weight matrix ($n \times n$, n being the number of training data samples) for predicting every layout. We did not test the contoning backward model, as a nearest neighbor search among 0.8 billion points in a 31-dimensional spectral space is intractable.

6.4 Limitations and Future Work

Our experiment with the historical pigments revealed that our ink library is suboptimal, unable to faithfully reproduce certain spectral curve shapes, such as cobalt blue. An exciting extension of our

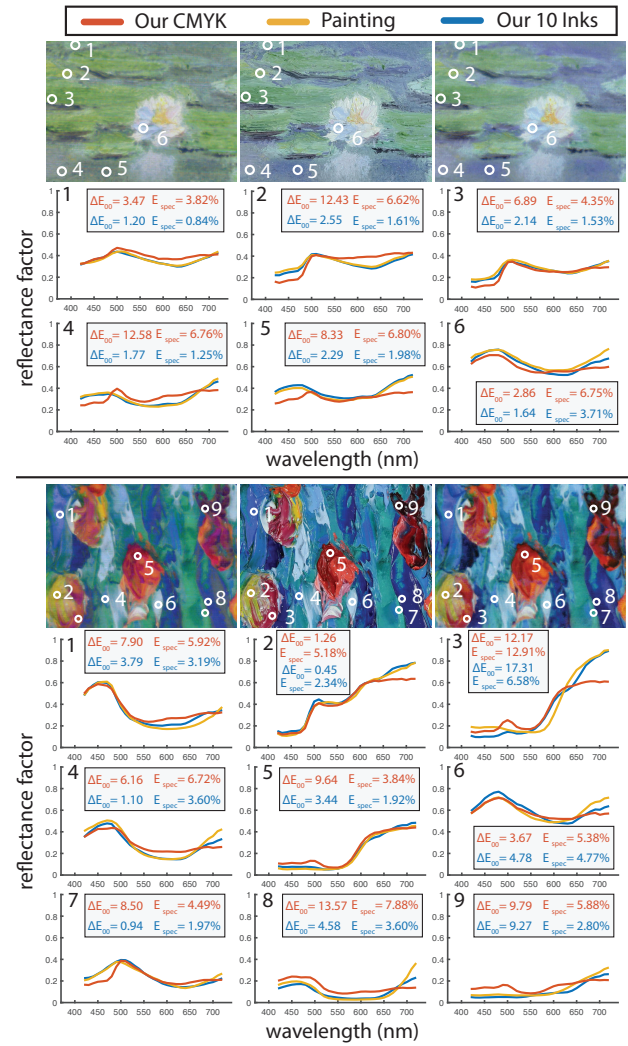


Fig. 11. Comparisons of the printed results and measured spectral reflectance between *our* CMYK reproduction, 10-ink reproduction and the original painting. In the lily example (top), the reconstruction of green and purple hues are extensively examined, which are known to be difficult for CMYK. In the bottom example, we examine a wide variety of both saturated and desaturated colors. In both results, 10 inks result consistently outperforms the CMYK both perceptually and in terms of spectral RMSE. Paintings ©Azadeh Asadi.

system is to further expand the ink library. With an extended ink library, an efficient painting-specific ink selection algorithm would be required to allocate inks from the library to the limited number of printer channels.

A data-driven model fits well with our current problem size. However, for maximum scalability, striving for more advanced physical prediction models is an important research direction for the future. For this, one can take inspiration from physically-based models for rendering of layered materials [Jakob et al. 2014]. Another appealing direction is to construct a physically-aware neural network that, instead of working with ink labels, is built on proper spectral

measurements. By learning from the physics of the problem, we expect the resulting network to be generalizable to arbitrary inks not present in the training set. However, we expect a significant larger training set to be necessary given the increase in the number of degrees of freedom.

Compared to the previous contoning method [Babaei et al. 2017], the dot gain in our approach is significantly lower due to the thin (high-concentration) ink stacks. Consequently, the blur is much less significant without any preprocessing [Babaei et al. 2017] or optimizing the material arrangement for the least crosstalk [Elek et al. 2017]. Having said that, we believe that improving the blur problem through careful tuning of material thickness and concentration, optimizing halftoning and contoning parameters, and finding optimal material arrangements [Elek et al. 2017] is necessary for high-quality reproduction.

Finally, a painting reproduction without considering the rich spatially-varying gloss and translucency found in many paintings as well as the 3D trail of the brush-stroke is far from complete. Incorporating gloss and microgeometry will also improve the spectral reproduction by recognizing the unavoidable highlights measured during the capture, and compensating for them by surface reflection (and not by diffuse color). Although there have been recent efforts in combined fabrication of these appearance attributes Elkhuisen et al. [2017], there is still a long way to go for archival-quality fabricated fine art.

7 CONCLUSIONS

We propose a complete pipeline capable of reproducing spectral reflectance using a 3D printer. To this end, we present a data-driven approach for predicting the optimal stack of different inks. Our framework provides accurate and efficient forward and backward predictions. Our proposed spectral model outperforms state-of-the-art physical and data-driven prediction models, mostly by large margins. We demonstrate the effectiveness of our workflow by reproducing a number of challenging oil paintings, painted by our artist collaborator [Asadi 2017]. We further propose a novel spectral vector error diffusion that combines both halftoning and contoning techniques in a complementary manner to leverage the potential of both methods. Our sVED algorithm uses error diffusion on a very large number of potential contone stacks, very efficiently, thereby resolving both layout discretization and color quantization problems. Our framework takes the first step in the exciting direction of fine art reproduction, and we hope it triggers more works, leading to the exploitation of rapid advances in fabrication technologies and computational techniques to protect our cultural heritage.

ACKNOWLEDGMENTS

We highly appreciate the help by David Kim for formulating our inks. We would like to thank Azadeh Asadi for painting all presented paintings, Todd Zickler and Ioannis Gkioulekas for spectral camera hardware and software, and Hossein Amirshahi for pointing to suitable pigments. We also thank our pigment providers: BASF, Lansco Colors, Penn Color, Sun Chemical and Toyo Ink Group. Vahid Babaei and Changil Kim are supported by the Swiss National Science Foundation (SNSF) fellowships P300P2 171212

and P2EZP2 168785, respectively. This work is partially supported by the U. S. National Science Foundation (NSF) grants IIS-1421435, CHS-1617236, IIS-1815070, and IIS-1815585.

REFERENCES

- Martín Abadi, Paul Barham, Jianmin Chen, Zhifeng Chen, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Geoffrey Irving, Michael Isard, Manjunath Kudlur, Josh Levenberg, Rajat Monga, Sherry Moore, Derek G. Murray, Benoit Steiner, Paul Tucker, Vijay Vasudevan, Pete Warden, Martin Wicke, Yuan Yu, and Xiaoqiang Zheng. 2016. TensorFlow: A System for Large-scale Machine Learning. In *Proceedings of the 12th USENIX Conference on Operating Systems Design and Implementation (OSDI'16)*. USENIX Association, Berkeley, CA, USA, 265–283. <http://dl.acm.org/citation.cfm?id=3026877.3026899>
- Satoshi Abet and Gabriel Marcu. 1994. A neural network approach for RGB to YMCK color conversion. In *TENCON'94. IEEE Region 10's Ninth Annual International Conference. Theme: Frontiers of Computer Technology. Proceedings of 1994*. IEEE, 6–9.
- Azadeh Asadi. 2017. Freelance artist. <http://www.azadehasadi.com/index.html>. (2017). Accessed: 2018-08-15.
- Teun Baar, Hans Brettel, and Maria V Ortiz Segovia. 2015. Towards gloss control in fine art reproduction. In *Measuring, Modeling, and Reproducing Material Appearance 2015*, Vol. 9398. International Society for Optics and Photonics, 93980T.
- Teun Baar, Sepideh Samadzadegan, Philipp Urban, and Maria V Ortiz Segovia. 2016. Interrelation between gloss and texture perception of 2.5 D-printed surfaces. *Electronic Imaging* 2016, 9 (2016), 1–6.
- Vahid Babaei and Roger D Hersch. 2016. N-Ink printer characterization with barycentric subdivision. *IEEE Transactions on Image Processing* 25, 7 (2016), 3023–3031.
- Vahid Babaei, Kiril Vidimčec, Michael Foshey, Alexandre Kaspar, Piotr Didyk, and Wojciech Matusik. 2017. Color contoning for 3D printing. *ACM Trans. Graph. (SIGGRAPH)* 36 (2017).
- Seung-Hwan Baek, Incheol Kim, Diego Gutierrez, and Min H Kim. 2017. Compact single-shot hyperspectral imaging using a prism. *ACM Transactions on Graphics (TOG)* 36, 6 (2017), 217.
- Farhan Baqai, Je-Ho Lee, Jan P Allebach, et al. 2005. Digital color halftoning. *Signal Processing Magazine, IEEE* 22, 1 (2005), 87–96.
- Fausto Bernardini, Joshua Mittleman, Holly Rushmeier, Cláudio Silva, and Gabriel Taubin. 1999. The ball-pivoting algorithm for surface reconstruction. *Visualization and Computer Graphics, IEEE Transactions on* 5, 4 (1999), 349–359.
- Roy S. Berns, Brittany D. Cox, and Farhad Moghareh Abed. 2015. Wavelength-dependent spatial correction and spectral calibration of a liquid crystal tunable filter imaging system. *Appl. Opt.* 54, 12 (Apr 2015), 3687–3693. <https://doi.org/10.1364/AO.54.003687>
- Roy S Berns, Lawrence A Taplin, Philipp Urban, and Yonghui Zhao. 2008. Spectral color reproduction of paintings. In *Conference on Colour in Graphics, Imaging, and Vision*, Vol. 2008. Society for Imaging Science and Technology, 484–488.
- Jean-Yves Bouguet. 2008. A Release of a Camera Calibration Toolbox for Matlab. (2008).
- Matthew Brown and David G Lowe. 2007. Automatic panoramic image stitching using invariant features. *International journal of computer vision* 74, 1 (2007), 59–73.
- Alan Brunton, Can Ates Arikan, Tejas Madan Tanksale, and Philipp Urban. 2018. 3D Printing Spatially Varying Color and Translucency. *ACM Trans. Graph.* 37, 4, Article 157 (July 2018), 13 pages. <https://doi.org/10.1145/3197517.3201349>
- Alan Brunton, Can Ates Arikan, and Philipp Urban. 2015. Pushing the limits of 3d color printing: Error diffusion with translucent materials. *ACM Transactions on Graphics (TOG)* 35, 1 (2015), 4.
- Yongda Chen, Roy S Berns, and Lawrence A Taplin. 2004. Six color printer characterization using an optimized cellular Yule-Nielsen spectral Neugebauer model. *Journal of Imaging Science and Technology* 48, 6 (2004), 519–528.
- Inchang Choi, Daniel S Jeon, Giljoo Nam, Diego Gutierrez, and Min H Kim. 2017. High-quality hyperspectral reconstruction using a spectral prior. *ACM Transactions on Graphics (TOG)* 36, 6 (2017), 218.
- Antonino Cosentino. 2014. FORS Spectral Database of Historical Pigments in Different Binders. 2 (09 2014), 57–68.
- George Cybenko. 1989. Approximation by superpositions of a sigmoidal function. *Mathematics of Control, Signals and Systems* 2, 4 (01 Dec 1989), 303–314. <https://doi.org/10.1007/BF02551274>
- Hugh R. Davidson and Henry Hemmendinger. 1966. Color Prediction Using the Two-Constant Turbid-Media Theory*. *J. Opt. Soc. Am.* 56, 8 (Aug 1966), 1102–1109. <https://doi.org/10.1364/JOSA.56.001102>
- Maxim Derhak and Mitchell Rosen. 2006. Spectral colorimetry using LabPQR: an interim connection space. *Journal of Imaging Science and Technology* 50, 1 (2006), 53–63.
- Paul Apostolos Drakopoulos and Ganesh Subbarayan. 2002. Color printer characterization using optimization theory and neural networks. (Nov. 12 2002). US Patent 6,480,299.

- Oskar Elek, Denis Sumin, Ran Zhang, Tim Weyrich, Karol Myszkowski, Bernd Bickel, Alexander Wilkie, and Jaroslav Krivánek. 2017. Scattering-aware Texture Reproduction for 3D Printing. *ACM Trans. Graph.* 36, 6, Article 241 (Nov. 2017), 15 pages. <https://doi.org/10.1145/3130800.3130890>
- Willemijn S Elkhuizen, Tessa TW Essers, Boris Lenseigne, Clemens Weijkamp, Yu Song, Sylvia C Pont, Jo M-P Geraedts, and Joris Dik. 2017. Reproduction of Gloss, Color and Relief of Paintings using 3D Scanning and 3D Printing. (2017).
- Willemijn S Elkhuizen, Boris AJ Lenseigne, Teun Baar, Wim Verhofstad, Erik Tempelman, Jo MP Geraedts, and Joris Dik. 2015. Reproducing oil paint gloss in print for the purpose of creating reproductions of Old Masters. In *Measuring, Modeling, and Reproducing Material Appearance 2015*, Vol. 9398. International Society for Optics and Photonics, 93980W.
- Willemijn S Elkhuizen, Tim Zaman, Wim Verhofstad, Pieter P Jonker, Joris Dik, and Jo MP Geraedts. 2014. Topographical scanning and reproduction of near-planar surfaces of paintings. In *Measuring, Modeling, and Reproducing Material Appearance*, Vol. 9018. International Society for Optics and Photonics, 901809.
- Scott E Fahlman and Christian Lebiere. 1990. The cascade-correlation learning architecture. In *Advances in neural information processing systems*. 524–532.
- Robert W. Floyd and Louis Steinberg. 1976. An Adaptive Algorithm for Spatial Greyscale. In *Proceedings of Society for Information Displays*, Vol. 17. 75–77.
- Jérémie Gerhardt and Jon Y Hardeberg. 2007. Controlling the error in spectral vector error diffusion. In *Color Imaging XII: Processing, Hardcopy, and Applications*, Vol. 6493. International Society for Optics and Photonics, 649316.
- Ioannis Gkioulekas, Shuang Zhao, Kavita Bala, Todd Zickler, and Anat Levin. 2013. Inverse Volume Rendering with Material Dictionaries. *ACM Trans. Graph.* 32, 6, Article 162 (Nov. 2013), 13 pages. <https://doi.org/10.1145/2508363.2508377>
- Jean Hergel and Sylvain Lefebvre. 2014. Clean color: Improving multi-filament 3D prints. In *Computer Graphics Forum*, Vol. 33. Wiley Online Library, 469–478.
- Francisco H. Imai, Mitchell R. Rosen, and Roy S. Berns. 2002. Comparative study of metrics for spectral match quality. In *Conference on Colour in Graphics, Imaging, and Vision*, Vol. 2002. Society for Imaging Science and Technology, 492–496.
- Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, and Alexei A. Efros. 2017. Image-to-Image Translation with Conditional Adversarial Networks. In *2017 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2017, Honolulu, HI, USA, July 21-26, 2017*. 5967–5976. <https://doi.org/10.1109/CVPR.2017.632>
- Wenzel Jakob, Eugene d'Eon, Otto Jakob, and Steve Marschner. 2014. A comprehensive framework for rendering layered materials. *ACM Transactions on Graphics (TOG)* 33, 4 (2014), 118.
- Henry R Kang and Peter G Anderson. 1992. Neural network applications to the color scanner and printer calibrations. *Journal of Electronic Imaging* 1, 2 (1992), 125–136.
- Isaac Kauvar, Samuel J Yang, Liang Shi, Ian McDowall, and Gordon Wetzstein. 2015. Adaptive color display via perceptually-driven factored spectral projection. *ACM Transactions on Graphics (TOG)* 34, 6 (2015), 165.
- Takayuki Kawaguchi, Norimichi Tsumura, Hideaki Haneishi, Yoichi Miyake, and M Kouzaki. 1999. Vector error diffusion method for spectral color reproduction. In *PICS*. 394–397.
- Diederik Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980* (2014).
- R. Victor Klassen, Reiner Eschbach, and Krishna Bharat. 1994. Vector error diffusion in a distorted colour space. In *Proc. of IS&T 47th Annual Conference*. 63–65.
- P. Kubelka and F. Munk. 1931. Ein Beitrag zur Optik der Farbanstriche. *Zeitschrift für technische Physik* 12 (1931), 593–601.
- D. J. Littlewood, P. A. Drakopoulos, and G. Subbarayan. 2002. Pareto-optimal Formulations for Cost Versus Colorimetric Accuracy Trade-offs in Printer Color Management. *ACM Trans. Graph.* 21, 2 (April 2002), 132–175. <https://doi.org/10.1145/508357.508361>
- M Ronnier Luo, Guihua Cui, and B Rigg. 2001. The development of the CIE 2000 colour-difference formula: CIEDE2000. *Color Research & Application* 26, 5 (2001), 340–350.
- Gabriel Marcu and Kansei Iwata. 1993. RGB-YMCK color conversion by application of the neural networks. In *Color and Imaging Conference*, Vol. 1993. Society for Imaging Science and Technology, 27–32.
- Panagiotis Takis Metaxas. 1998. Optimal parallel error diffusion dithering. In *Color Imaging: Device-Independent Color, Color Hardcopy, and Graphic Arts IV*, Vol. 3648. International Society for Optics and Photonics, 485–495.
- Peter Morović, Ján Morović, Jordi Arnabat, and Juan Manuel García-Reyero. 2012. Revisiting spectral printing: A data driven approach. In *Color and Imaging Conference*, Vol. 2012. Society for Imaging Science and Technology, 335–340.
- Ole Norberg and Daniel Nyström. 2013. Extending color primary set in spectral vector error diffusion by multilevel halftoning. In *Color Imaging XVIII: Displaying, Processing, Hardcopy, and Applications*, Vol. 8652. International Society for Optics and Photonics, 86520M.
- Victor Ostromoukhov. 1993. Chromaticity gamut enhancement by heptatone multi-color printing. In *IST/SPIE 1993 Symposium of Electronic Imaging: Science and Technology, Conf. on Device Independent Color Imaging and Imaging Systems Integration*, Vol. 1905. 139–151.
- Daniele Panozzo, Olga Diamanti, Sylvain Paris, Marco Tarini, Evgeni Sorkine, and Olga Sorkine-Hornung. 2015. Texture Mapping Real-World Objects with Hydrographics. *Computer Graphics Forum (proceedings of EUROGRAPHICS Symposium on Geometry Processing)* 34, 5 (2015), 65–75.
- Théo Phan Van Song, Christine Andraud, and Maria V Ortiz Segovia. 2016a. Implementation of the four-flux model for spectral and color prediction of 2.5 D prints. In *NIP & Digital Fabrication Conference*, Vol. 2016. Society for Imaging Science and Technology, 26–30.
- Théo Phan Van Song, Christine Andraud, and Maria V Ortiz-Segovia. 2016b. Towards spectral prediction of 2.5 D prints for soft-proofing applications. In *Image Processing Theory Tools and Applications (IPTA), 2016 6th International Conference on*. IEEE, 1–6.
- Théo Phan Van Song, Christine Andraud, and Maria V Ortiz-Segovia. 2017. Spectral predictions of rough ink layers using a four-flux model. In *Color and Imaging Conference*, Vol. 2017. Society for Imaging Science and Technology, 251–257.
- Tim Reiner, Nathan Carr, Radomir Měch, Ondřej Štáva, Carsten Dachsbacher, and Gavin Miller. 2014. Dual-color mixing for fused deposition modeling printers. In *Computer Graphics Forum*, Vol. 33. Wiley Online Library, 479–486.
- Mitchell R Rosen and Maxim W Derhak. 2006. Spectral gamuts and spectral gamut mapping. In *Spectral Imaging: Eighth International Symposium on Multispectral Color Science*, Vol. 6062. International Society for Optics and Photonics, 60620K.
- Mitchell R Rosen, Edward F Hattenberger, and Noboru Ohta. 2004. Spectral redundancy in a six-ink ink jet printer. *Journal of Imaging Science and Technology* 48, 3 (2004), 194–202.
- Romain Rossier. 2013. *Framework for Printing with Daylight Fluorescent Inks*. Ph.D. Dissertation. Ecole Polytechnique Fédérale de Lausanne (EPFL), Lausanne.
- Claude Rozé, Thierry Girasole, and Anne-Gaëlle Tafforin. 2001. Multilayer four-flux model of scattering, emitting and absorbing media. *Atmospheric environment* 35, 30 (2001), 5125–5130.
- Christian Schüller, Daniele Panozzo, Anselm Grundhöfer, Henning Zimmer, Evgeni Sorkine, and Olga Sorkine-Hornung. 2016. Computational thermoforming. *ACM Transactions on Graphics (TOG)* 35, 4 (2016), 43.
- Pitchaya Sitthi-Amorn, Javier E. Ramos, Yuwang Wang, Joyce Kwan, Justin Lan, Wen-shou Wang, and Wojciech Matusik. 2015. MultiFab: A Machine Vision Assisted Platform for Multi-material 3D Printing. *ACM Trans. Graph.* 34, 4, Article 129 (July 2015), 11 pages. <https://doi.org/10.1145/2766962>
- Chang-Hwan Son, Hyung-Min Park, and Yeong-Ho Ha. 2011. Improved color separation based on dot-visibility modeling and color mixing rule for six-color printers. *Journal of Imaging Science and Technology* 55, 1 (2011), 10505–1.
- Eric J Stollnitz, Victor Ostromoukhov, and David H Salesin. 1998. Reproducing color images using custom inks. In *Proceedings of the 25th annual conference on Computer graphics and interactive techniques*. ACM, 267–274.
- Stratasys. 2016. Stratasys J750 the ultimate full-color multi-material 3D printer. <http://www.stratasys.com/3d-printers/production-series/stratasys-j750>. (2016). [Online; accessed 15-October-2016].
- 3D Systems. 2013. ProJet 5500X Multi-Material 3D Printer. (2013). <http://www.3dsystems.com/files/projet-5500x-1113-usen-web.pdf>
- Lawrence A Taplin and Roy S Berns. 2001. Spectral color reproduction based on a six-color inkjet output system. In *Color and Imaging Conference*, Vol. 2001. Society for Imaging Science and Technology, 209–213.
- Shoji Tominaga. 1996. Color control using neural networks and its application. In *Color Imaging: Device-Independent Color, Color Hard Copy, and Graphic Arts*, Vol. 2658. International Society for Optics and Photonics, 253–261.
- Shoji Tominaga. 1998. Control scheme for printers using more than three color inks. In *Electronic Imaging: Processing, Printing, and Publishing in Color*, Vol. 3409. International Society for Optics and Photonics, 286–294.
- Shohei Tsutsumi, Mitchell Rosen, and Roy Berns. 2008. Spectral color reproduction using an interim connection space-based lookup table. *Journal of Imaging Science and Technology* 52, 4 (2008), 40201–1.
- Philipp Urban and Roy S Berns. 2011. Parametric mismatch-based spectral gamut mapping. *IEEE transactions on image processing* 20, 6 (2011), 1599–1610.
- Gunter Wyszecki and Walter Stanley Stiles. 1982. *Color Science*. Vol. 8. Wiley New York.
- Songhua Xu, Haisheng Tan, Xiantao Jiao, Francis Lau, and Yunhe Pan. 2007. A generic pigment model for digital painting. In *Computer Graphics Forum*, Vol. 26. Wiley Online Library, 609–618.
- Yizhong Zhang, Yiying Tong, and Kun Zhou. 2016. Coloring 3D Printed Surfaces by Thermoforming. *IEEE Transactions on Visualization and Computer Graphics* (2016).
- Yizhong Zhang, Chunji Yin, Changxi Zheng, and Kun Zhou. 2015. Computational hydrographic printing. *ACM Transactions on Graphics (TOG)* 34, 4 (2015), 131.