Mixed Integer Ink Selection for Spectral Reproduction

NAVID ANSARI, Max Planck Institute for Informatics, Germany OMID ALIZADEH-MOUSAVI, Depsys SA, Switzerland HANS-PETER SEIDEL, Max Planck Institute for Informatics, Germany VAHID BABAEI, Max Planck Institute for Informatics, Germany







(a) Photograph of reproduction with selected inks.

(b) Photograph of the painting.

(c) Photograph of the cyan-yellow reproduction.

Fig. 1. Our mixed integer ink selection for spectral duotone (two-ink) printing. Our selected inks reproduce a limited-palette painting, made of four paints, with remarkable accuracy (a). The best pair from a CMYK ink set (cyan and yellow) generates a poor reproduction (c). Duotone reproduction provides powerful visual evidences of the quality of our ink selection method as the smallest mistake would stand out prominently. While the problem size in this particular example is small, we show that our method gives the optimal result when selecting tens of inks from libraries of thousands in a reasonable time.

We introduce a novel ink selection method for spectral printing. The ink selection algorithm takes a spectral image and a set of inks as input, and selects a subset of those inks that results in optimal spectral reproduction. We put forward an optimization formulation that searches a huge combinatorial space based on mixed integer programming. We show that solving this optimization in the conventional reflectance space is intractable. The main insight of this work is to solve our problem in the spectral absorbance space with a linearized formulation. The proposed ink selection copes with large-size problems for which previous methods are hopeless. We demonstrate the effectiveness of our method in a concrete setting by lifelike reproduction of handmade paintings. For a successful spectral reproduction of high-resolution paintings, we explore their spectral absorbance estimation, efficient coreset representation, and accurate data-driven reproduction.

$\label{eq:CCS} \text{Concepts:} \bullet \textbf{Computing methodologies} \to \textbf{Computer graphics}.$

Additional Key Words and Phrases: Spectral reproduction, mixed-integer programming, computational fabrication

Authors' addresses: Navid Ansari, Max Planck Institute for Informatics, Saarbrücken, Germany, nansari@mpi-inf.mpg.de; Omid Alizadeh-Mousavi, Depsys SA, Puidoux, Switzerland, omid.mousavi@depsys.ch; Hans-Peter Seidel, Max Planck Institute for Informatics, Saarbrücken, Germany, hpseidel@mpi-sb.mpg.de; Vahid Babaei, Max Planck Institute for Informatics, Saarbrücken, Germany, vbabaei@mpi-inf.mpg.de.

ACM Reference Format:

Navid Ansari, Omid Alizadeh-Mousavi, Hans-Peter Seidel, and Vahid Babaei. 2020. Mixed Integer Ink Selection for Spectral Reproduction. *ACM Trans. Graph.* 39, 6, Article 255 (December 2020), 16 pages. https://doi.org/10.1145/3414685.3417761

1 INTRODUCTION

Fabricating objects with desired appearance properties, known as appearance fabrication, is an important but complex, multifaceted research mission. One of the main challenges of appearance fabrication is rooted in *materials*. Real-life objects are made of a variety of materials, most of which cannot be used by digital fabrication. This situation only exacerbates when materials for fabrication come prepacked from factory. For example, inkjet printers (2D and 3D), despite the flexibility originated in their drop-on-demand technology, posses a fixed number of inks with predetermined optical properties. While sufficient for customary applications, freedom in selecting fabrication materials could bring significant benefits to demanding domains, such as fine art reproduction [Morovič et al. 2012].

Aiming at further utilizing the inherent flexibility of digital printers, we propose to adapt the printing materials to the appearance of input. This ensures that the fittest materials make their way to the printer's limited ink channels. In this research we focus on ink selection for spectral painting reproduction, where we measure and print the reflectance¹ of a painting. We select the optimal subset

This is the author's version of the work. It is posted here for your personal use. Not for redistribution. The definitive Version of Record was published in ACM Transactions on Graphics, https://doi.org/10.1145/3414685.3417761

¹Computer graphics refers to "reflectance" as the angular surface reflection (BRDF). In color science, on the other hand, the term reflectance (along with "spectrum/spectra" and "spectral reflectance") denotes spectral, diffuse color. In this manuscript, we adopt the latter convention.

from a large library of inks, such that it results in the most accurate appearance reproduction of the input. There are however significant computational challenges facing the ink selection problem. The main challenge is the huge combinatorial search space that renders traditional methods, such as genetic algorithm, ineffective in realworld scenarios. Further challenges include the spectral modeling of the combination of a large number of inks, and evaluating the performance of different ink subsets on high-resolution spectral images.

We introduce a novel algorithm that solves the spectral ink selection problem reliably and efficiently. The binary nature of the ink selection problem makes integer programming a viable solution. We demonstrate, however, that a problem formulation in terms of reflectance is intractable. A reflectance-based ink selection problem leads to a mixed-integer nonlinear programming (MINLP) which is non-convex with no optimality guarantee and very challenging to solve. Our main insight is to solve the physical dual of the reflectancebased ink selection problem. Instead of searching for inks whose reflectances result in the most faithful reproduction of a painting's reflectance, we explore the ink library to find inks' absorbances that optimally reproduce the painting's spectral absorbance. The insignificant scattering of printing inks allows us to perform color mixing in the absorbance space. The absorbance-based problem formulation results in a mixed-integer linear programming (MILP) whose continuous relaxation is convex and scales gracefully with larger problems.

The spectral absorbance of an ink or a painting can be obtained from its spectral transmittance using a simple analytical formula. But measuring the transmittance of a painting on a non-transparent substrate is challenging. We thus first measure the spectral reflectance of paintings relying on RGB-based spectral reconstruction [Connah et al. 2001]. We then use a data-driven method based on a neural network to map the measured reflectances to their corresponding transmittances. The ink-selection problem requires a set of new optimization variables for each pixel in the candidate image for reproduction. This can lead to huge optimization problem when dealing with high-resolution spectral images. We show that extracting a small, representative subset, or coresets, of all spectra present in the painting can effectively replace the use of all spectral pixels during ink selection. For a thorough evaluation of our ink-selection method, we devise a complete system of physical reproduction of paintings. We demonstrate high-fidelity spectral reproduction through spectral separation using a neural network, a library of 43 inks, and a custom printing setup. Furthermore, we make a simple but important observation about the neural spectral separation [Shi et al. 2018] where we adapt the neural network to different inputs leading to a considerable accuracy gain.

Our main contributions are:

- A versatile formulation based on mixed-integer programming for the problem of ink selection by solving this problem in the absorbance space.
- An analysis of different coreset approximation algorithms that represent high-resolution spectral images with significantly lower number of data points.



Fig. 2. A halftone patch made of three inks gives rise to 8 Neugebauer primaries. Such a halftone is represented by its area coverage vector (here of size 3). The area coverage of Neugebauer primaries (here of size 8) are computed using Demichel equations (Equations 2 and 4).

- A thorough spectral reproduction framework (including capture and print) of paintings.
- An adaptive neural separation strategy for higher reproduction accuracy.

To ensure full reproducibility of our results, both implementation and data will be released. Notably, a spectral dataset of captured paintings, the print datasets for calibrating our neural networks, our ink library, and the set of optimal inks for each painting will be publicly available.

2 BACKGROUND

2.1 Halftone Spectral Prediction

Almost all printers are binary devices: they either print at a given location or leave it unprinted. For printing images with a continuous tone impression, halftoning algorithms take advantage of the low-pass filtering property of the human visual system [Kang and Anderson 1992]. Halftoning methods create spatial binary patterns of dots with different sizes and spacings so that, when viewed from a sufficient distance, give the impression of continuous tone with the least objectionable artifacts. In most color halftoning methods, a halftone layer is created for each ink separately and the final colorhalftone image is formed by the superposition of all layers. The partial superposition of different binary ink layers helps creating a limited number of additional colors, known as *Neugebauer primaries* (Figure 2).

Any spectral reproduction workflow requires a prediction model that computes the reflectance of a printed halftone. *Spectral Neugebauer* model [Yule 1967] is the most widely known spectral prediction model for calculating the spectrum **r** of a halftone

$$\mathbf{r} = w_1 \mathbf{r}_1 + w_2 \mathbf{r}_2 + \dots + w_k \mathbf{r}_k,\tag{1}$$

where \mathbf{r}_i denotes the spectrum of the Neugebauer primaries made of base inks and their possible superpositions each with w_i relative area in the halftone, known as *area coverage*. The Neugebauer model can be extended to the well-known Yule-Nielsen modified spectral Neugebauer model [Hersch and Crété 2005; Chen et al. 2004] in a straightforward manner by mapping the reflectances to 1/n space, *n* being an empirical real number known as the Yule-Nielsen value.

We can easily find the spectra of the Neugebauer primaries, e.g., by first printing and then measuring them using a measurement device, such as a spectrophotometer. To determine the area coverages of Neugebauer primaries, we rely on a set of probabilistic equations known as the Demichel equations [Demichel 1924]. Demichel's insight was that printing k superposed dot layers gives rise to 2^k Neugebauer primaries. Furthermore, an ink area coverage gives its presence/absence probability in the halftone. Assuming ink dots are laid out independently from one another², the computed area coverage (or probability) of a given Neugebauer primary in a halftone is the multiplication of the presence/absence probabilities (or area coverages) of its constituting inks which are given for any halftone.

Let us further clarify by considering the classic CMY print with cyan, magenta and yellow *inks* (Figure 2) with area coverages *a*, *b* and *c*, respectively. The resulting halftone is composed of 8 Neugebauer *primaries*: white (w), cyan (c), magenta (m), yellow (y), blue (b), green (g), red (r) and black (k) whose weights w_x is given by the Demichel equation:

$$w_w : w_{000} = (1-a)(1-b)(1-c) \qquad w_r : w_{011} = (1-a)bc$$

$$w_c : w_{100} = a(1-b)(1-c) \qquad w_g : w_{101} = a(1-b)c$$

$$w_m : w_{010} = (1-a)b(1-c) \qquad w_b : w_{110} = ab(1-c)$$

$$w_t : w_{001} = (1-a)(1-b)c \qquad w_k : w_{111} = abc.$$

(2)

In this notation, w_c (w_{100}), for example, denotes the area coverage of cyan *primary* by finding its probability which is the probability of presence of cyan *ink* (1) and absence of magenta (0) and yellow (0) *inks*. Later, for the first time to the best of our knowledge, we introduce a general and compact Demichel formula in Equation 4.

2.2 Mixed Integer Programming

Mixed-Integer Programming (MIP) is a category of mathematical optimization problems in which all or some of variables are integers [Floudas 1995]. The objective and constraints of the optimization problem can be linear, quadratic or other non-linear functions. MIP problems are inherently non-convex because the feasible region of a problem with integer variables is not convex. There are however important implications if the continuous relaxation of a MIP is convex. That is, if we relax the MIP by ignoring the integrality condition, then we obtain a convex optimization problem. The solution of a convex optimization problem can be guaranteed to be the optimal solution of that problem. Therefore, the solution of the original MIP problem can be compared against the global optimum of its continuous relaxation and give a reliable measure of optimality, called the duality gap. Note that an important class of MIP problems with convex continuous relaxation is the mixed-integer linear programs (MILP).

MIP problems are solved by a variety of algorithms, such as branch-and-bound [Borchers and Mitchell 1994] and cutting planes [Marchand et al. 2002]. There are several open-source and commercial solvers that are based on these methods and can solve MIP problems efficiently [Lofberg 2004]. The commercial solvers [Gurobi Optimization 2018] have better performances, in terms of computation time and types of problems they can solve, especially for larger scale problems. The ink selection problem is inherently a binary decision problem. The proposed formulation, based on the absorbance and with further linearization, allows us to develop a tractable MILP optimization problem with linear constraints and linear objectives. The optimization can be efficiently solved by existing commercial solvers to determine the optimal set of inks with strong optimality evidences.

2.3 Related Work

Ink Selection. In a brute-force ink selection, we can evaluate the reproduction capacity of any N-ink subset of a K-ink library on each of P spectral pixels of the captured painting. This combinatorial search space is too expensive to compute for any problem with a real-life size. A handful of previous works on this topic [Stollnitz et al. 1998; Power et al. 1996; Tzeng and Berns 1999] have therefore resorted to either stochastic algorithms or hand-crafted heuristics. Power et al. [1996] investigate the ink selection problem within their duotone reproduction of color images. Using a simulated annealing approach with an objective function that evaluates the colorimetric performance of any pair of inks, they select the two best inks that result in the color matching of an input image. Stollnitz et al. [1998] try a similar approach for custom-ink color reproduction except that they improve the optimization by using a genetic algorithm instead of simulated annealing. Tzeng and Berns [1999] use heuristics to reduce the size of the ink library first, and then evaluate the reproduction accuracy of all remaining ink subsets. Contrary to these methods, we introduce a novel ink selection formulation based on mixed-integer programming [Belotti et al. 2013] that can solve very large problems efficiently with strong optimality evidences. In Section 6.6, we show comparison to genetic algorithm, the most promising previous method.

Spectral Reproduction. Spectral 2D printing is an established field of research in color reproduction. We refer the readers to Shi et al. [2018] for a nice literature summary. The aims of spectral reproduction is to recreate a spectral image whose pixels represent spectral bands sampled throughout the visible range of the electromagnetic spectrum [Berns et al. 2008]. The main motivation behind spectral reproduction is to avoid metamerism [Wyszecki and Stiles 1982] where the color of a reproduction matches the original under one light source, but shows a mismatch under another light source with a different spectrum. In order to expand the spectral gamut [Rosen and Derhak 2006] of printers, N-ink printers employ custom inks in addition to the process CMYK [Ostromoukhov 1993]. The most important challenge of spectral reproduction with N-ink printing (even with N = 6 or 7) is the forward modeling of the combination of multiple inks. The forward modeling most commonly relies on the Yule-Nielsen but this model is known to scale very poorly with the number of employed inks (N), both in terms of the model complexity and the number of required calibration prints [Babaei and Hersch 2016].

Our ink selection method encounters the same problem, albeit on remarkably larger scales: we need to deal with modeling any *N*-ink combinations of *K* inks in our library with $K \gg N$. Our solution for dealing with this problem is to rely on an approximate

²Amidror and Hersch [2000] extensively discuss the validity of this assumption for different halftones.



Fig. 3. Our pipeline to reproduce paintings including ink selection and neural-network spectral reproduction. The bottom rectangular box shows the steps for reproducing a painting after the ink selection is performed and the inverse neural network, based on selected inks, is computed.

model, based on a linear spectral absorption, in the ink selection optimization where many ink combinations exist. Apart from its lightweight calibration (only one spectral measurement per each ink in the library), it leads to a highly scalable optimization problem. We demonstrate that for obtaining high-accuracy results, more dataintensive prediction models could be later used for final spectral reproduction using selected inks. In particular, we rely on a datadriven method based on neural networks [Shi et al. 2018]. We show that the accuracy of this network can be considerably improved by adapting it to the input content without requiring additional data.

Color Reproduction of 3D Objects. The recent introduction of multichannel inkjet 3D printers has enabled creating physical objects with spatially-varying surface colors. Inkjet 3D printing, the offspring of digital 2D printers, juxtaposes voxels with different colors at high volumetric resolutions resulting in unprecedented quality. The recent works address different challenges from spatial placement of materials, through halftoning [Brunton et al. 2015] or contoning [Babaei et al. 2017] to suppressing subsurface scattering crosstalk [Elek et al. 2017] and reproducing the color of thin geometric features faithfully [Sumin et al. 2019]. All these works can potentially benefit from our ink selection algorithm (see Section 7). Ink selection could also be useful in the fused filament fabrication as hinted by Song et al. [2019].

3 OVERVIEW

Our proposed method takes a painting, finds the optimal inks for its spectral reproduction, and creates the reproduction using a datadriven approach. Figure 3 sketches different steps of our method. Note that the paper does not necessarily follow the order shown in this figure.

Our end-to-end spectral reproduction begins with capturing an RGB image of a painting (as we don't have access to an accurate spectral camera). Since incorporating all pixels of the image gives rise to prohibitively large ink-selection optimizations, we extract a representative set from the painting, called the coreset (Section 5.2), and use this set in the ink selection method. We then reconstruct the coreset's spectral reflectance from the captured RGB images using a nonlinear regression with high accuracy (Section 6.2). As we perform the ink selection in the spectral absorbance space, spectral

reflectance of the coreset is converted into spectral transmittance, handled by a reflectance-to-transmittance (R2T) neural network (Section 5.1). From spectral transmittances, we obtain the spectral absorbances using an analytical formula. The spectral absorbances of our inks in the library are also provided. Given the coreset of a painting and the ink library in the spectral absorbance space, we carry out the ink selection based on a mixed integer linear programming approach (Section 4.2). The MILP-based ink selection outputs, from within the library, a desired number of inks that best reproduce the painting.

Once the inks are selected, we rely on an accurate data-driven model to reproduce input spectral images faithfully. To that end, we print a large dataset of halftones made of selected inks and obtain their spectral reflectances. We then train, in a two-step approach (Section 5.3), an inverse neural network that performs the *neural spectral separation*: mapping any reflectance to its corresponding ink area coverages. Finally, for actual reproduction of the input painting (shown in a rectangular box in Figure 3), we estimate its spectral reflectance and feed it, pixel by pixel, to the inverse network. The resulting ink area coverage *layers* are halftoned independently and printed.

4 MIXED-INTEGER INK SELECTION

In this section, we introduce our optimization formulation for ink selection based on mixed-integer programming. We start with the natural choice of optimizing in the spectral reflectance space (Section 4.1) and show that the reflectance-based problem is intractable. Although we will not use the reflectance-based problem for results shown in this paper, we build on its mathematical formulation to introduce a significantly more scalable ink selection method in the absorbance space (Section 4.2).

4.1 Reflectance-based Ink Selection Formulation

In this section, we put forward our mixed-integer formulation for selecting N inks from an ink library composed of K inks with known spectral *reflectances*. The N selected inks will best reproduce a given spectral image represented via its coreset by P pixels each containing a spectral *reflectance* of size S (due to sampling spectral reflectances)

at S points). We write the optimization problem as

$$\underset{\mathbf{x},\mathbf{W}}{\operatorname{argmin}} \left\| \mathbf{H} \ \mathfrak{D}(\mathbf{W}) - \mathbf{R} \right\|_{F}$$
(3a)

$$\mathbf{x}_k \in \{0, 1\} \qquad \forall \ k \in \gamma \tag{3b}$$

$$\sum_{k=1}^{K} \mathbf{x}_k \le N \tag{3c}$$

$$\mathbf{W}_{k,p} \in \mathbb{R} \quad \forall \ k \in \gamma, \ \forall \ p \in \phi$$
(3d)

$$0 \le \mathbf{W}_{k,p} \le \mathbf{x}_k \qquad \forall \ k \in \gamma, \ \forall \ p \in \phi \tag{3e}$$

$$\gamma = \{1, 2, \cdots, K\}, \ \phi = \{1, 2, \dots, P\},$$
 (3f)

where $\|\cdot\|_F$ denotes the Frobenius norm and we have the following entities:

Variables

- **x** is the *ink-selection variable*, a *binary* vector of size $K \times 1$ whose *k*-th element (pointing at *k*-th ink in the library) is denoted by **x**_k.
- W is the matrix of continuous variables representing the area coverages of each library *ink* corresponding to each pixel in the input image, thus a matrix of dimension $K \times P$ ($W_{k,p}$ denotes the element in the *k*-th row and *p*-th column).

Parameters

- *N* is the maximum number of desired inks to be selected from the library *N* < *K*.
- H is the matrix of *D* Neugebauer primary reflectances with dimensions *S*×*D* which can be measured for small libraries or approximated for the large ones [Phan Van Song et al. 2016].
- **R** is the spectral image, a 2D matrix with dimensions *S* × *P* for which we select optimal inks.

Functions

• $\mathfrak{D}(\cdot)$ is the Demichel *operator*, a nonlinear transformation applied on W (Equations 2 and 4). Collectively, this operator maps the ink area coverage matrix W with dimensions $K \times P$ to the Neugebauer primary area coverage matrix A with dimensions $D \times P$ ($D = 2^K$).

Simply put, this formulation continuously changes the area coverages (**W**) of different inks in the library, predicts their spectral reflectances with the Neugebauer model³ (the term $\mathbf{H} \mathfrak{D}(\mathbf{W})$), minimizes the difference between the predicted reflectances and the target (**R**), and gives the best up to *N* inks for this purpose. Although we are not necessarily interested in optimal ink area coverages **W**, varying them is imperative as it explores the spectral gamut of the considered inks.

A set of *K* independent, continuous variables (corresponding to each column of **W**) is required to *explore* each of *P* spectral pixels of the target spectral image, hence the size of **W** ($K \times P$). This implies that adding a pixel to a spectral image, or more exactly to its coreset, adds a new column in **W**, i.e., an additional *K* continuous variables. Moreover, adding an ink to the library introduces a new row in **W**, i.e., an additional *P* continuous variables. In contrast, the binary ink selection vector only relates to the inks in the library and *not* the

target image, leading to a *K*-vector **x**. Adding an ink to the library introduces only a single binary element in **x**.

The inequality constraint 3c along with the set of $K \times P$ other inequalities (3e) ensure that at most N inks are used for reproducing the considered P spectra. These constraints confine *different* columns of **W** to have *the same* up to N non-zero elements. Indices of these non-zero elements in each **W**'s column match the indices of non-zero elements in vector **x** and pinpoint the location of optimal inks in our library.

 $\mathfrak{D}(\cdot)$ is the Demichel operator for which we gave an example of a three-ink CMY halftone in Equation 2. We introduce here a more compact form of this operator for any *K*-ink halftone. For an image with *P* pixels, it should be applied on *all* ink area coverages of that pixel, hence γ in notation $\mathbf{W}_{\gamma,p}$ in the following equation. We have $\mathbf{A} = \mathfrak{D}(\mathbf{W})$, such that

$$\mathbf{A}_{d,p} = \left(\mathfrak{D}(\mathbf{W}_{\gamma,p})\right)_{d,p} = \prod_{k=1}^{K} 1 - \mathbf{U}_{d,k} - (-1)^{\mathbf{U}_{d,k}} \mathbf{W}_{k,p}, \quad (4)$$

where U is a $D \times K$ matrix containing all $D = 2^K$ binary encoding of a string of length K:

$$\mathbf{U} = \begin{bmatrix} 0 & 0 & \dots & 0 \\ 1 & 0 & \dots & 0 \\ 0 & 1 & \dots & 0 \\ \vdots & \vdots & \vdots & \vdots \\ 1 & 1 & \dots & 1 \end{bmatrix}.$$
 (5)

While a novel, valid ink selection formulation (Equations 3a-3f), optimizing this problem is far from trivial. First, the size of the problem grows exponentially with the size of the library. This includes the size of the Neugebauer primary matrix **H** with a rate of 2^{K} , and the number of terms made of our continuous variables W corresponding to the rows of A in Equation 4 (recall that the number of variables grows linearly with the the library size). Second, due to the Demichel operator $\mathfrak{D}(\cdot)$, the objective function becomes nonlinear and this nonlinearity increases with more inks in the library. The nonlinearity results from the multiplication of continuous variables and there is not any suitable linearization technique to cope with this type of non-linearity [Bisschop 2012]. Thus the optimization problem is a mixed-integer nonlinear programming (MINLP) which generally is non-convex. Even the state-of-the-art commercial MIP solvers face serious difficulties to address this type of problem and there is no guarantee for the solution to be optimal.

4.2 Absorbance-based Ink Selection Formulation

With these observations, we rethink the problem of optimal ink selection at a more conceptual level. Our main insight is to solve the physical dual of the ink selection problem by transferring the problem from the reflectance space to the absorbance (or equivalently transmittance) space. The question then becomes: *which inks' absorbances do reproduce a given spectral absorbance image optimally?* The main implication of using the absorbance space is that we can model the mixing of *inks* themselves rather than their Neugebauer primaries, evading the Demichel equations and its scalability and nonlinearity issues.

 $^{^3}$ We may use the Yule-Nielsen model instead and, working in 1/n space, have the same implications.

Inspired by the literature [Babaei et al. 2017, 2012; Berns 1993], we can use *the thickness* of each ink (instead of primaries) in a linear manner to predict the spectral absorbance of the halftone. While a concrete explanation of why in the absorbance space the primaries can be bypassed requires further studies, intuitively it could be related to the logarithmic conversion. Similar to the use of densities in photographic printing [Evans 1948], the absorbance space emphasizes the (important) absorptive regions of the inks spectra which have small values in the reflectance space. Thus, in the absorbance space, the ambiguities of the contribution of inks in the aggregate spectrum of a halftone is significantly reduced. In an early experiment, we performed ink selection using a linear model of inks (and not primaries) reflectances and obtained completely inaccurate results.

We propose the absorbance-based ink selection formulation as

$$\underset{\mathbf{x},\mathbf{C}}{\operatorname{argmin}} \sum_{p=1}^{P} \sum_{s=1}^{S} \left| (\mathbf{G} \ \mathbf{C} - \mathbf{Q})_{s,p} \right|. \tag{6}$$

We use the sum of absolute differences (equivalent to 1-norm in vectors) and after a *linearization* step, have a completely linear mixed integer problem:

$$\underset{\mathbf{x},\mathbf{C},\mathbf{Z}}{\operatorname{argmin}} \sum_{p=1}^{P} \sum_{s=1}^{S} \mathbf{Z}_{s,p} \qquad (7a)$$

$$\mathbf{Z}_{s,p} \ge (\mathbf{G} \ \mathbf{C} - \mathbf{Q})_{s,p} \quad \forall \ s \in \psi, \ \forall \ p \in \phi \qquad (7b)$$

$$\mathbf{Z}_{s,p} \ge -(\mathbf{G} \ \mathbf{C} - \mathbf{Q})_{s,p} \quad \forall \ s \in \psi, \ \forall \ p \in \phi \qquad (7c)$$

x

$$k \in \{0, 1\} \quad \forall k \in \gamma \quad (7d)$$

$$\sum_{k=1}^{K} \mathbf{x}_k \le N \qquad (7e)$$

$$\mathbf{C}_{k,p} \in \mathbb{R} \quad \forall \ k \in \gamma, \ \forall \ p \in \phi$$
 (7f)

$$0 \le C_{k,p}/t \le \mathbf{x}_k \quad \forall \ k \in \gamma, \ \forall \ p \in \phi \quad (7g)$$

$$\gamma = \{1, 2, \cdots, K\}, \ \phi = \{1, 2, \dots, P\}, \ \psi = \{1, 2, \dots, S\},$$
(7h)

which has similar mathematical and conceptual implications to the reflectance-based formulation. Apart from the repeated entities in Equation 3, we have the following: Variables

• C is the matrix of the library ink thicknesses corresponding to each pixel in the input image, thus of dimension *K* × *P* (it is the counterpart of **W** in reflectance formulation).

 Z is a matrix of size S×P containing a set of auxiliary variables that linearize the absolute value function (Equation 6) through constraints 7b and 7c.

Parameters

- G is the matrix of (measured) spectral absorbances of the library inks (*S* × *K*).
- **Q** is the target, spectral absorbance image, a matrix with dimensions *S* × *P*.
- *t* controls the maximum thickness of library inks.

Note that the term G C gives the *modeled* spectral absorbance image through linear mixing of the measured ink absorbances G in

the library with weights C. We discuss how to obtain spectral absorbance image Q in Section 5.1. As in the case of reflectance-based optimization, we are not interested in the optimal thickness values C but rather the binary decision variables **x**. This frees us from finding any explicit relationship between ink thicknesses and ink area coverages. Moreover, by setting a parameter *t* in Equation 7g we let the optimization explore the thicknesses beyond 1. In other words, the inks can have more than a single layer in a halftone. Note that in practice also our custom printer driver (Section 6.1.1) allows us to print multiple layers of each ink. In all our experiments we set *t* to 4 allowing each ink in the library to have up to four layers.

The optimization formulations in Equations 3 and 7 bear a lot of similarities. The key difference is that in the proposed absorbancebased formulation the Demichel operator is bypassed, leaving us with a mixed-integer programming that scales gracefully with the size of the ink library. The resulting optimization problem is a mixed integer programming with linear objective and linear constraints. This problem can be efficiently solved with existing commercial solvers for a large ink library.

5 SPECTRAL REPRODUCTION

With the optimization in hand, there still remain a few questions to address before being able to select the optimal inks for an input spectral image, and employ the selected inks for a final reproduction. First, how do we obtain spectral absorbances (Section 5.1)? Second, how do we handle high resolution images with potentially millions of spectra during ink selection (Section 5.2)? Third, having selected optimal inks for a particular image, how do we use those inks for high-fidelity spectral reproduction (Section 5.3)?

5.1 Reflectance to Transmittance Network

The spectral absorbance for each visible wavelength (A_{λ}) is most conveniently obtained from its corresponding spectral transmittance (T_{λ}) using Beer-Bouguer law

$$A_{\lambda} = -\ln T_{\lambda}.$$
 (8)

The transmittance of the ink library, when printed on a transparent substrate, is straightforward to measure using a spectrophotometer or a camera (Section 6.2). The challenge is the transmittance measurement of the paintings, which are painted on relatively opaque substrates, such as paper and canvas.

We propose to compute the spectral transmittance of an image from its measured spectral reflectance. Unlike the Fresnel reflectance, which is related to transmittance in a simple way (T = 1 - R), the *bulk* reflectance of a material measured on a diffuse substrate has a non-trivial relationship with its transmittance due to, for example, light scattering and multiple reflections between the substrate and the print-air interface across the inks. We therefore resort to a datadriven method based on a neural network to solve this problem by mapping measured spectral reflectances to their counterpart transmittances. In the absence of fluorescence, there is a one to one correspondence between reflectance and transmittance. This implies that the optimal ink selection in absorbance (transmittance) space gives an optimal answer for the original reflectance space. We train a fully-connected feed-forward neural network, R2T, that takes a spectral reflectance and gives a spectral transmittance. The loss

ACM Trans. Graph., Vol. 39, No. 6, Article 255. Publication date: December 2020.

$$\text{RMSE\%}\left(\mathbf{r}_{\text{gt}}, \mathbf{t}_{\text{gt}}\right) = \frac{1}{\sqrt{S}} \left\|\text{R2T}\left(\mathbf{r}_{\text{gt}}\right) - \mathbf{t}_{\text{gt}}\right\|_{2} \times 100. \tag{9}$$

In this equation, \mathbf{r}_{gt} is a member of the ground truth reflectance data set, \mathbf{t}_{gt} is a member of the ground-truth transmittance data set, R2T is the reflectance to transmittance neural network model, and *S*, as in previous section, is the size of spectral vectors.

5.2 Ink Selection Coresets

in color science

An important challenge in optimal ink selection is to incorporate all spectra in the target spectral image. The number of pixels in a captured painting, denoted by *P* in Equations 3 and 7, can easily amount to millions. Unlike finite-sum minimization problems [Schmidt et al. 2017], typically found in machine learning, the number of optimization variables in our formulation grows with the number of pixels *P*. As we mentioned earlier, this is because the nature of our problem requires to assign a new set of (*K*) continuous variables to each spectral pixel. Therefore, as the number of pixels increases so does the size of our optimization variables **W** and **C** (both of size $K \times P$) in the optimization problems 3 and 7.

This motivates us to investigate different *coreset* estimation methods for ink selection. With origins in computational geometry, a coreset is a subset that can serve as a proxy for an original set [Agarwal et al. 2005]. The result of a particular algorithm applied on a full dataset is *approximated* when the same algorithm is applied on the dataset's coreset⁴. In our case, a coreset is extracted from the full set of painting's pixels.

As the task is similar to color quantization, we evaluate (Section 6.3) some clustering methods, namely *k*-means and *k*-medoids clustering. Furthermore, we investigate the geodesic simplex volume maximization [Heylen et al. 2011] a well-known, nonlinear algorithm in remote sensing for performing *spectral unmixing*. Spectral unmixing is concerned with computing the spectra of distinct materials that are mixed within a single satellite pixel, which are know as *endmembers* (equivalent to coreset) [Keshava and Mustard 2002].

5.3 Neural Spectral Separation

The optimal inks for reproduction of a particular painting, or more exactly its coreset, are selected using an approximate forward model based on linear combination of spectral absorbances of inks. For actual spectral reproduction, we run an accurate, data-driven, neural network prediction model. This network is very similar to the one proposed by Shi et al. [2018] for reproducing spectral oil paintings except we use it for predicting multi-layer halftones while Shi et al. [2018] predict *contone* ink stacks. The goal of the neural spectral separation, composed of a *forward* and an *inverse* network as a decoder and encoder respectively, is to take a spectral reflectance as input and estimates a halftone area coverage that, after printing,



Fig. 4. The neural spectral separation [Shi et al. 2018], composed of a forward and an inverse network, finds a set of ink area coverages that approximates any input reflectance.

reproduces the target spectra. An interesting feature of this architecture, originally proposed by [Tominaga 1996], is to first train the forward network that predicts the reflectance of any halftone. In the second phase of training, the inverse network maps reflectances to halftone area coverage while the forward network is used only for spectral error prediction without being trained (see Figure 4).

The network can be trained using printed patches with known area coverages paired with their corresponding measured reflectances (Section 6.4). Note that the spectral separation network is trained for a given set of inks. That is, a separate training is required for printing an input image with a set of inks specifically selected for its optimal reproduction. The loss function has a spectral term (similar to Equation 9) but not a colorimetric one. Given the ink selection is carried out on a spectral basis, we chose to stay with a purely spectral reproduction. As multiple halftone configurations can produce the same spectra, we introduce a halftone area coverage regularization term to limit the total ink amount printed in order to prevent the paper wetting.

5.4 Adaptive Neural Spectral Separation

Having trained a forward model using halftone-reflectance pairs, the role of the inverse network is to take a reflectance and map it into a set of halftone area coverages. In traditional spectral separation methods, the area-coverage for a single input reflectance is computed by numerically inverting the employed forward model, such as the well-known Yule-Nielsen spectral Neugebauer [Rolleston and Balasubramanian 1993]. Thanks to the expressive power of neural networks [Hornik et al. 1989], the neural inverse model proposed by Shi et al. [2018] is computed for a distribution of reflectances at once. At the test time, each input spectrum is *separated* using the computed weights and biases of the inverse network.

We crucially observe that the inverse network need *not* be trained in a supervised fashion. That is, having trained the forward network using halftone-reflectance pairs, training the inverse network requires only spectral reflectances thanks to the encoder-decoder architecture. Thus, instead of computing a *universal* inverse model using a representative set of spectra (as in [Shi et al. 2018]), we propose to adapt the inverse network to the set of input spectra directly at the test time. This adaptation is possible by both training the network (for a particular input) from the scratch or build on the

 $^{^4\}mathrm{As}$ an example, the convex hull of a point set is a coreset for computing the diameter of the set.

existing universal network similar to transfer learning [Tan et al. 2018b]. While we rely on a universal inverse network for producing the majority of the results, we evaluate the effect of this modification in Section 6.8.

6 RESULTS AND EVALUATION

6.1 Experimental Setup

6.1.1 Printing Hardware. As the reproduction platform, we use Epson Stylus Photo R2880, a 720 dpi, 8-channel 2D printer with Piezo inkjet printhead. The printer is expected to place different layers of halftone patterns, sent to printer as a sequence of binary bitmaps, on the paper as precisely as possible. In practice, the printer's default driver may change these patterns at will. We therefore rely on a custom-written printing driver developed using the printer's command language programming⁵. In addition to precise printing of binary patterns, our driver can print multiple layers of the same ink, an otherwise non-straightforward task.

6.1.2 Ink Library. We develop a library of 43 different inks. We begin with a larger number (around 80) of inks and keep those inks which are jettable and don't show clogging or other inconsistency issues. Our library consists of pigment-based and dye-based inks used in both thermal and Piezo inkjet printhead technologies. The list of inks in our library and our paper substrate along with their spectral plots are included in the supplementary materials.

6.1.3 Paintings and Paint Swatches. We collaborate with an artist [Asadi 2017] who produces different paintings and a manually prepared swatch made of uniform paint patches. For our watercolor paintings, we purchase 14 Schmincke paint tubes. The list of paints and their spectral reflectances are reported in the supplementary materials. We mix these paints manually to prepare our paint swatch made of 423 patches, each with different mixing ratios so that the swatch is a representative of the colors present in our paintings. Note that our ink selection method does not require a swatch made of paints used in each input painting. The main advantage of the swatch is that we can obtain ground-truth spectral measurements, with a spectrophotometer, and use them in the spectral reconstruction stage for a better accuracy (Section 6.2); an accurate spectral camera would render the use of swatch obsolete. Apart from being used in spectral reconstruction, the paint swatch proves useful in some quantitative evaluation of our method notably the evaluation of coreset methods (Section 6.3).

6.2 Spectral Reflectance and Transmittance Acquisition and Reconstruction

Spectrophotometers are the standard tools for accurate measurement of spectral reflectances and transmittances. Unfortunately, they are unable to capture the spatially varying spectra found, e.g., in paintings. *Spectral cameras* with dedicated hardware can capture the visible spectrum at high spatial resolution. Despite the high cost, they might still sacrifice the spectral accuracy [Shi et al. 2018]. We adopt a more cost-effective approach by using a color camera and relying on *spectral reconstruction* to convert RGB signals to spectral reflectances [Maali Amiri and Fairchild 2018]. This well-known inverse problem can be solved with remarkable accuracy since most natural and man-made materials, including our inks and paints, have smooth spectra. The smoothness property implies that the true dimensions of most spectra is significantly lower than 31 (the typical sampling rate from 400 to 700 nm at 10 nm intervals which we also use in all our experiments) [Maloney 1986]. In our case, a principal component analysis (PCA) on the watercolor swatch reveals that the paint spectra can be represented almost perfectly using only 8 dimensions.

We use a nonlinear regression for spectral reconstruction, which learns the relationship between camera signals and ground truth spectral data measured by a spectrophotometer. In order to have a more accurate spectral reconstruction for a given type of surface, our regression learns a separate model with ground truth data for that surface. The watercolor swatch (Section 6.1.3) is used for spectral reconstruction of paintings, printed patches on paper are used for spectral reconstruction of thousands of small patches used in calibrating the spectral separation network (Section 6.4), and printed patches on transparency sheets are used for spectral transmittance reconstruction of small patches used in the R2T network (Section 6.5). The ground truth samples are printed or painted large enough $(2 \times 2 \text{ cm})$ so they can be measured by a spectrophotometer. In this work, all ground-truth spectral transmittance and reflectance measurements are performed by X-Rite Color i7, a sphere spectrophotometer. The measurement geometry, in accordance with observer's experience, provides a diffuse illumination and discards the specular reflection.

In order to evaluate the spectral reconstruction quality of watercolor paints, we learn the regression model using 399 patches in our paint swatch and test it on the remaining, unseen 24 patches. This model results in average RMSE = 2.77% and average CIEDE2000 color difference (ΔE_{00}) [Sharma et al. 2005] 2.41 and 2.74 under D65 and A illuminants, respectively⁶. For all CIELab calculations in this work we use the paper substrate as the white reference.

To ensure the reconstructed spectra are not metameric to the original ones, we plot the 24 test spectra against their ground-truth measurements carried out by the spectrophotometer in Figure 5. We observe that the reconstructed paint spectra follow the spectrophotometric measurements very well without sinusoidal interceptions typical to metameric pairs. In order to see a more detailed assessment of the metamerism potential, we refer the reader to the supplementary materials.

6.3 Coreset Estimation Evaluation

In this section, we evaluate different coreset estimation methods discussed in Section 5.2. We compare the ink selection performance when carried out on an entire set against when done on a subset extracted from that original set. We cannot use an image as original set since computing the ink selection on millions of pixels is intractable. Instead, we rely on our watercolor swatch (Section 6.1.3) as the original set. The swatch includes the underlying paints and their different mixtures, thus a good representative of the colors in

⁵The printer's programming guide is not anymore available online but can be found on the project's website.

⁶Similar experiments evaluating the accuracy of our regression-based spectral reconstruction method used for estimating the reflectance or transmittance spectra of printed surfaces are reported in supplementary materials.

ACM Trans. Graph., Vol. 39, No. 6, Article 255. Publication date: December 2020.



Fig. 5. Spectral reconstruction of 24 test watercolor paint patches compared against ground truth measurements carried out by a spectrophotometer.



Fig. 6. The ink-selection loss on the watercolor swatch (423 spectra) as a function of different coreset methods with different sizes. The dashed line is the loss obtained using all 423 samples of the original set as the coreset.

our paintings. The performance metric of each coreset is the ink selection loss computed using Equation 6. More specifically, we: 1) estimate different coresets of the color swatch using different coreset extraction methods with different sizes, 2) use these subsets to select the optimal inks, and 3) use the selected inks to compute the ink selection loss on *all* spectra of the original set.

Figure 6 shows the ink selection loss on 423 samples of the watercolor swatch as a function of different coreset sizes for different coreset estimation methods. As expected, choosing the original set as the ink selection "coreset" results in the best loss (shown as a horizontal dashed line). More interestingly, all methods show performances near the best possible outcome (dashed line) with relatively small coreset sizes. This has practical implications for us as it suggests a significantly smaller optimization problem can still work reliably for ink selection on large input images.

The spectral unmixing using geodesic simplex volume maximization [Heylen et al. 2011] performs better than other methods by converging to the best answer at smaller coreset sizes. Its downside, compared to other clustering methods, is its slower performance when applied to high-resolution images. Furthermore, we noticed that coreset extraction methods behave similarly when applied on a spectral or an RGB image. Therefore, for all our painting reproductions and analyses we extracted the coreset using a *k*-means clustering (k = 200). Another interesting experiment is to use the 14 underlying paints used in the watercolor swatch as the ink-selection coreset. The underlying paints, shown as a small circle in Figure 6,

	RMSE%		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
Forward	0.63	0.44	1.41	0.84	1.36	0.78	1.34	0.76
	0.61	7.38	1.25	10.45	1.22	9.14	1.21	10.12
Inverse	0.85	0.66	3.38	2.50	3.22	2.33	3.05	2.16
	0.69	8.94	2.72	14.73	2.54	14.80	2.44	13.40

Table 1. Spectral and colorimetric accuracy of the forward and inverse networks on a dataset made with 8 original Epson inks.



Fig. 7. The performance of our linear absorbance model used in our absorbance-based optimization. For 300 test patches printed with 8 Epson inks, after converting the modeled absorbances to transmittances using Equation 8 we obtain mean RMSE = 2.83%, 90 percentile RMSE = 8.57%, mean ΔE_{00} (D65) = 2.73 and mean ΔE_{00} (A) = 3.10. From the test set, 24 randomly picked transmittance spectra are shown here.

perform the best when the coreset size is 14, equal to the number of underlying paints.

6.4 Evaluation of Prediction Models

In this section, we evaluate the performance of our spectral prediction models, i.e., the linear absorbance model used in our ink selection optimization along with our neural network based forward and inverse models.

The role of the absorbance model is to predict the absorbance spectra of *digital* mixtures of arbitrary inks in the library. This analytically simple model works with inks (and not primaries) absorbances, resulting in tractable optimization problems. It is also light-weight to calibrate requiring only a single spectral transmittance measurement per ink. We test the absorbance model using 300 printed patches on transparency made with different combinations of the 8 original inks of our Epson printer. Using only the measured transmittance of the 8 base inks, converted to absorbance with Equation 8, we obtain relatively high accuracy as shown in Figure 7 where 24 diverse spectral transmittances from the set of 300 patches are plotted. Note that for each prediction, the ink thickness (*t* in Equation 7g) is fitted, as it also occurs during optimization.

For the utmost reproduction accuracy, we use a fully data-driven neural spectral separation [Shi et al. 2018] whose function is to

255:10 • Ansari, Alizadeh-Mousavi, Seidel, Babaei

	RMSE%		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
R2T	4.46	2.73	6.87	3.84	6.76	3.73	6.57	3.65
	3.90	19.36	6.21	33.70	6.16	33.74	5.99	33.53

Table 2. Spectral and colorimetric accuracy of the R2T network. This network converts reflectances into transmittances to provide data for ink selection in the absorbance space. The training dataset is spectral reflectance and transmittance pairs of the same ink combinations printed on paper and transparency sheet, respectively.

map a spectral reflectance vector to a vector of halftone area coverages through its inverse network. As explained in Section 5.3, for computing a meaningful loss, the separation network relies on a pre-trained forward model that maps the area coverages to spectral reflectances. The evaluations in this section (and Section 6.7) are based on a universal inverse model. In Section 6.8, we evaluate the role of adaptive separation approach.

Both forward and inverse networks of the neural spectral separation are multi-layer perceptrons with 4 hidden layers and ReLU activation functions. The separation network possesses 160 nodes per layer while the forward network has 300, 300, 200, and 100 nodes per each layer, respectively. The training protocol is identical for the two networks where we use the Adam optimizer [Kingma and Ba 2014] with initial learning rate 5×10^{-3} and learning decay 0.95, randomly split the training data into 80% training and 20% validation, and allow 20 epochs with batch size 31. Both models are implemented using PyTorch, and trained and tested on an NVIDIA Titan X GPU very efficiently: training each network takes about 1 minute. At the test time, predicting 1 million spectra takes roughly 10 seconds.

Table 1 summarizes the performance of both forward and separation networks. The two networks show very high accuracy in terms of both colorimetric (under several illuminations) and spectral metrics. The inverse network shows a slightly lower accuracy as its task is more challenging than the forward network. Recall that these networks are ink specific. That is, they should be completely recalibrated for each new set of inks. The ink set used to generate the results shown in Table 1 are the 8 original inks of our Epson printer. We use 33,945 printed patches for training and testing these specific networks. As measuring these patches with a spectrophotometer is extremely time and material consuming, we print all patches in a small size $(1.5 \times 1.5 \text{ mm})$, fit them to 4 A4-size paper substrates, and obtain their spectra using our camera and spectral reconstruction procedure.

6.5 Evaluation of the R2T Network

The R2T network (Section 5.1) converts spectral reflectances to spectral transmittances allowing us to carry out the ink selection in the absorbance space. Similar to the forward and inverse networks, R2T is a fully-connected feed-forward neural network with the same configurations except it has 5 hidden layers with 350, 760, 160, 100, and 50 nodes, respectively. It is trained using 33,945 pairs of transmittance and reflectance spectra printed with 8 original Epson



Fig. 8. (a) The effect of coreset size on MILP optimization time. (b) The effect of increasing the ink-library size on MILP optimization time. (c) The spectral absorbance loss of GA at multiple runs each with 50 generations where the loss of best generation is reported. The dashed horizontal line shows the loss obtained by MILP for an identical problem. (d) The effect of increasing the ink-library size on GA optimization time.

inks on transparency sheets and paper substrates, respectively. Both types of spectra are measured using our spectral reconstruction procedure. The results of evaluating the R2T network, using 20% of the data for validation, are shown in Table 2. In practice, the R2T network, trained on a set of inks, is used to compute the transmittance of paints. As uniform application of paints on transparency is not possible, we don't have a direct way of evaluating the network's performance on paints.

The R2T network shows an acceptable performance yet its error is larger than previous networks. One reason for the inaccuracy is the ink *blotting* occurring during printing training patches on transparency sheets. Unlike prints on the paper, inks tend to agglomerate on the transparency sheets resulting in nonuniform patches. In future, using a more professional printing setup that avoids blotting, can enhance the results of this module. Also, analytical flux transfer approaches, such as the four-flux model [Phan Van Song et al. 2016] remain a viable alternative for separating the transmittance from reflectance. A main advantage of such models would be that they can operate for inks and paints separately. Nevertheless, our current approach leads to high-quality ink selection and spectral reproduction as we will see in next sections.

6.6 Ink-selection Optimization Performance

In Figure 8 we summarize the performance and scalability of the MILP-based ink selection. We solve all mixed-integer problems using Gurobi version 9.0.1, a commercial solver with free academic license [Gurobi Optimization 2018]. We compare MILP with genetic

algorithm (GA), a method already used for ink selection [Stollnitz et al. 1998]. Figure 8a shows how MILP scales with the size of the ink-selection coreset size. In this experiment we use our default ink library of size 43 and test coresets of size 5 to 420 (at irregular intervals). From Equation 7, we know that the number of continuous variables increases linearly with the size of the coreset. The experimental complexity shown in Figure 8a suggests that MILP scales well with the number of continuous variables enabling incorporation of large coresets.

In a second experiment shown in Figure 8b, we fix the coreset size to 5 (extracted from *Cat*) and study the effect of the library size on the MILP performance. As ink library, we work with *Munsell* [Tyler and Hardy 1940] spectral samples (around 1200 spectra). We perform the MILP ink-selection (number of inks N = 5) multiple times where each time we increase the library size with 50-ink steps. Once again, Equation 7 suggests that increasing the library size increases the the number of binary variables in a linear fashion. This time, however, the optimization time increases more drastically. Yet we obtain reasonable solution times for even very large library sizes. Note that the fluctuations in solution time are typical to the underlying methods used in MILP.

Figure 8c shows the performance of the GA method for selecting 5 inks from a library of size 43 for a coreset of size 5. When performing GA-based ink selection, due to its stochastic nature, we perform multiple optimizations (each called a *run*) and allow 50 generations within each run. In Figure 8c, we report the loss of the best generation of each independent run. For comparison, we solve the same problem using MILP and show its loss alongside the GA loss in the same figure. We observe that after 28 runs, GA sometimes approaches but never outperforms the MILP performance.

Finally, in Figure 8d, we repeat the experiment in Figure 8b but use GA (instead of MILP) for optimization in order to study the effect of library size on the GA solution time. The problem setting is identical in both experiments but as running 50 times the GA even with 50-ink intervals is slow, we perform it once and multiply the solution time by 50. We also stop proceeding with GA after the library size approaches 900 as the method becomes prohibitively slow. Comparing Figures 8b and 8d, regardless of the the raw values, reveals that GA is significantly disadvantaged against MILP, in terms of scaling to problems with large sizes of ink library.

On a more general note, all solutions of different MILP-based ink selections in this paper reached a user specified duality gap of 10^{-4} . This means the objective value of the selected inks using MILP is at most 10^{-4} loss unit far from the global optimal of the relaxed, convex problem. Given the magnitude of our absorbance loss, we can ensure our solution is practically optimal. This result has significant implications compared to stochastic methods: we can solve ink selection problems with high confidence using MILP.

6.7 Spectral Painting Reproduction

Thus far, we have evaluated all different components of our proposed method separately: spectral acquisition and reconstruction, reflectance to transmittance transformation, MILP ink selection and linear absorbance model within it, and neural-network based spectral separation. In this section, we evaluate these components jointly



Fig. 9. Photographs of a painting and its printed reproduction under three different physical illuminations. The selected inks by our MILP algorithm are Epson Cyan (ID 41 in the ink library in supplementary materials), Epson Magenta (ID 42), Waterproof Yellow (ID 16), and Gray Sublimation (ID 21)).



Fig. 10. Simulated reproduction of 24 patches from the watercolor swatch. The average RMSE for 423 patches is 3.40%. The average ΔE_{00} under D65, TL85, A and D50 is 3.09, 3.22, 3.07 and 3.11, respectively.

through physical reproduction of a number of paintings. All our paintings exhibit challenging features, such as diverse and saturated colors, high spatial frequencies, and color gradients. Gradient images are challenging to reproduce because they are prone to banding artifact, however in the supplementary materials we show that our spectral separation does not suffer from such a problem. In Figure 9,

255:12 • Ansari, Alizadeh-Mousavi, Seidel, Babaei



Fig. 11. Our mixed integer ink selection tested on a painting produced by our artist collaborator (e). The photograph of the reproduction using selected inks (d) and the photograph of reproduction using a set of CMYK Canon inks (f) are shown. A separate spectral separation network has been calibrated per each ink set. The colorimetric error map of reproduction with selected inks (a) and Canon CMYK (c) (over all pixels) have average values of $\Delta E_{00} = 3.41$ and $\Delta E_{00} = 4.15$ under D65, respectively. The spectral error maps of reproduction with selected inks (g) and Canon CMYK (i) (over all pixels) have average RMSE values of 2.84% and 3.18%, respectively. The area-coverage maps of the two unusual selected inks, i.e., Waterproof Yellow (b) and Sublimation Red (h) are also shown.

we show a watercolor painting (*Flower*) along with its printed reproduction captured under three physical illuminations with different color temperatures⁷. We select 4 inks from our ink library using our proposed approach and deploy a separation network calibrated using these inks and their combinations. Subjectively, the reproduction quality, observed under a variety of illuminations, is high such that discerning a difference between the painting and print from a normal viewing distance is difficult. Also, with changing the illumination, corresponding pairs of colors change consistently. When capturing these photographs, we disabled the white balance function of the camera in order to pronounce the effect of the illumination and spot possible mismatches due to metamerism. Note that, in this experiment, we are not concerned with how a human observer would judge these colors under different illuminations. This evaluation would require *color appearance models* [Fairchild 2013], which we consider out of scope of the current work.

For a quantitative evaluation of our reproduction, we use a calibrated separation network to simulate the reproduction of all 423 patches in our watercolor paint swatch. In Figure 10, we show a subset of the swatch in form of both spectral plots and visualizations under four different light sources. Note that the simulation is highly reliable as it relies on a dense, data-driven neural network calibrated

 $^{^7\}mathrm{High}$ resolution photographs of all reproductions are included in supplementary materials.

ACM Trans. Graph., Vol. 39, No. 6, Article 255. Publication date: December 2020.

with four selected inks. The average spectral RMSE (3.40%) is higher than the error encountered when evaluating the (inverse) network (0.91%). This is because many patches in the highly diverse water-color swatch lie outside the gamut of the four inks used for the reproduction.

6.7.1 Comparison with CMYK Reproduction. An interesting analysis of the ink selection is to compare it with the gold standard CMYK. In Figure 11 we show the painting *Sunset* reproduced by 4 selected inks using our MILP method. The selected inks consist of Epson Cyan (ID 41) and Magenta (ID 42), Waterproof Yellow (ID 16) and Sublimation Red (ID 20). We choose a set of standard Canon inks as the CMYK reference set whose reflectance spectra are reported in supplementary materials. We create two separation networks calibrated using these two sets of inks. From Figure 11, we can see an overall better reproduction when using the selected inks. This is backed by both average spectral and colorimetric errors of the two reproductions evaluated over all image pixels (see the caption of Figure 11).

A closer look at the selected inks leads to some interesting insights. In addition to cyan and magenta, the method selects two atypical inks, i.e., Waterproof Yellow, a more brownish than usual yellow ink, and Sublimation Red instead of a highly expected black. The area coverage maps of these two inks in Figure 11 reveal that these inks (especially the yellow) are behind the better reproduction of the selected ink set at regions where the CMYK set struggles.

Although both the ink selection algorithm and the separation network work on purely spectral basis, the colorimetric performance of the selected inks is higher than the CMYK set. There are however areas in the colorimetric error maps where the CMYK set performs better than the selected inks, most notably the violet area in the painting's bottom right. But, as the spectral error maps in Figure 11 suggest, even in the regions with poorer colorimetric accuracy, the selected inks outperform the CMYK in terms of spectral accuracy. It is also noteworthy to mention that the ink selection (Equation 7) optimizes for a set of inks that reproduce the spectra of an input painting *on average*. That is, despite the guaranteed superiority in overall performance, in some regions the selected inks might be outperformed by another set even in a spectral sense.

6.7.2 Optimal Inks for Duotone Reproduction. When reproducing our paintings using our library inks, we observe that the visible differences diminish when selecting three and more inks. In contrast, duotone reproduction [Power et al. 1996], i.e., color reproduction with only two inks, typically leads to significant visual differences. This behavior is summarized in Figure 12a where the ink selection loss as a function of number of selected inks is plotted for a few paintings. Additionally, in this figure, the result of ink selection loss is shown for a spectral gray ramp. The gray ramp, visualized in Figure 14a, is procedurally generated using a range of *perfect* grays that have identical reflectance values at all wavelengths. Due to its significantly lower color diversity, the gray ramp shows different loss behavior in Figure 12a where even a two-ink reproduction produces relatively small loss.

The effect of number of inks is also shown in Figure 13 where a painting (*Cat*) has been reproduced using 2 to 4 selected inks. Thus, duotone ink selection is an attractive showcase as, first, the effect of



Fig. 12. The ink-selection loss for different painting coresets shwon in Figures 12b, 9, 13, and a spectral gray ramp's coreset as a function of increasing number of selected inks. Photographs of our paintings used for different experiments: "Lady", "Couple", and "Sunrise"



Fig. 13. Effect of the number of selected inks on the reproduction quality for *Cat.* Photographs of reproduction with (a) 2 inks: Epson Yellow (ID 36), Waterproof Blue (ID 24), (b) 3 inks: Epson Cyan (ID 41), Magenta (ID 42), and waterproof Yellow (ID 16), (c) 4 inks: Epson Cyan (ID 41), Magenta (ID 42), Waterproof Yellow (ID 16), and Gray Sublimation (ID 21). (d) Photograph of the original painting.

ink selection is highly visible and, second, smallest mistake during ink selection would stand out prominently.

In Figure 13a, we show the result of duotone ink selection for *Cat*, the most challenging painting for our ink library according to Figure 12a where our method selects Epson Yellow (ID 36) and Waterproof Blue (ID 24). For a less challenging painting *Music* shown in Figure 1 with a limited palette of 4 paints, our duotone ink selection performs much better resulting in a near perfect reproduction. In

comparison, the best two-ink combination of the CMYK Canon inks (cyan and yellow) are significantly far from the original painting.

6.8 Evaluation of Adaptive Neural Separation

In this part, we compare the adaptive and universal approaches to the network-based spectral reproduction when using an identical set of inks. We *simulate* the reproduction of several spectral images, using both approaches, and show their reproduction accuracy in Table 3. For simulations, we take the area-coverages computed using different separation strategies and, instead of sending them to printer, compute their corresponding spectra using the forward network. As the network has been trained on printed data (in this case Canon CMYK inks), the simulations are highly reliable.

First, we compute a universal inverse network trained with Canon CMYK printed patches. In a second approach using an adaptive strategy, we take advantage of the unsupervised nature of the inverse network and train it using only the spectra found in the input images. Note that we still rely on the forward network trained on the Canon CMYK dataset but build the inverse neural network on the spectra of a given input image.

For training the adaptive network, inspired by transfer learning methods, we take weights computed for the universal network as the initial guess and further optimize them using the Adam optimizer. We also experimented with a random initial guess and obtained comparable results.

The spectral reproduction accuracy, listed in Table 3, shows a considerable improvement when using the adaptive separation for all images. The colorimetric accuracy is also higher when taking the adaptive separation approach. In Figure 14 we show the simulated spectral reproduction of the gray ramp (under D65 illumination) using both adaptive and universal approaches. We observe that the reproduction using the universal neural network leads to visible imbalance in the grayness at several locations on the ramp. On the other hand, the result of the adaptive network exhibits significant improvement of the visual results.

A universal network trained to work well on a diverse set of spectra, does not guarantee a high level of performance for images at a tiny corner of the spectral space. The spectral gray ramp is exactly such an example. A second advantage of adaptive separation, beside performance, is the robustness. Neural network weights and biases, computed using stochastic gradient descent, are known to fall in different local minima [Ruder 2016]. Therefore, with identical training data, we may compute different networks, with slightly different losses, depending on the local minimum found by the optimization. A slight difference in the overall loss on a diverse training data, however, can cause significant changes for a single image, especially if it has large, single-color regions. In contrast, in adaptive separation, we *learn* to perform well on a more specific set of spectra.

7 LIMITATIONS AND FUTURE WORK

Scattering Inks. We assume the employed inks don't exhibit significant volume scattering. This is a valid assumption as the world of graphic printing is, to a great extent, scattering free. One notable exception is the recent 3D printing inks with considerable volume





Fig. 14. Simulated spectral reproduction of a spectral gray ramp using universal and adaptive inverse networks.

	Un	iversal separa	tion	Adaptive separation			
	RMSE%	ΔE_{00} (D65)	ΔE_{00} (A)	RMSE%	ΔE_{00} (D65)	ΔE_{00} (A)	
Couple	4.43	2.77	2.82	3.17	2.58	2.79	
Cat	2.86	3.60	3.56	2.20	2.05	2.01	
Flower	2.74	2.20	1.96	2.23	1.44	1.42	
Lady	4.60	3.58	3.36	3.06	2.97	2.89	
Sunrise	2.21	3.28	3.05	1.85	2.70	2.33	
Sunset	3.18	4.15	3.48	2.60	2.42	2.65	
Gray ramp	4.36	6.61	5.19	3.60	3.05	3.66	

Table 3. Reproduction accuracy of universal and adaptive neural separation approaches evaluated on all pixels of several images.

scattering [Sumin et al. 2019; Elek et al. 2017]. For these types of inks, one can use the well-known Kubelka-Munk (KM) theory [Kubelka and Munk 1931] as the ink mixing model in the optimization framework. Thanks to the KM model's linearity, the resulting problem would still be a linear integer programming.

Ink Library. Our library is formed on an ad hoc basis where the inks are added if the printer can jet them reliably. An expanded ink library with diverse dyes/pigments, e.g., daylight fluorescent pigments, would significantly enhance the ink selection capability. The example of duotone reproduction mimics such a scenario as our current library is fairly diverse relative to only two printing channels.

The question of which inks to include in the library triggers an opportunity for task-specific *ink design* where, instead of selecting inks, we design inks with desired reflectance properties. Ink design can also be combined with ink selection by allowing the bulk mixing of inks in the library. This means that we need to design an appearance model for bulk ink mixing (similar to Papas et al. [2013]) on top of a model for digital mixing mentioned in this paper.

Reflectance to Transmittance Conversion. An important limitation of our method is the evaluation of the R2T network. Although we could not spot inaccuracies due to this network, it has been evaluated

on the printer inks while being used to predict paint transmittances. As mentioned in Section 6.5, since applying a uniform film of paints on transparency was not possible, we could not obtain ground truth paint transmittance measurements to compare against R2T predictions. While, in future, we can find a technique to produce these samples, we believe a reflectance-transmittance separation based on physical models is feasible and would be a more scalable solution. Another interesting possibility is to use a physically-based path tracer to render the training samples for our network.

Multi-objective Ink Selection. In the current work, spectral accuracy is the sole performance criterion for ink selection. It is highly desirable to include further optimization criteria, such as colorimetric accuracy and cost of inks [Hunt 2006]. As new objectives are likely in conflict with each other, mixed-integer, multi-objective programming schemes are required [Alves and Clímaco 2007]. Furthermore, with the colorimetric term and its non-linearity, keeping the problem tractable by new linearization schemes is an interesting challenge.

Coreset Estimation. Our current problem setup relies on minimizing the spectral error on a coreset that is an *overall* representation of the input image. Through alternative coreset estimation, we can prioritize important regions in the image. This *importance map*, analogous to region of interest in image processing, can come from the painting semantics and statistics or even be user specified. We believe that some approaches in computer graphics, such as soft color segmentation [Aksoy et al. 2017], image-to-layer decomposition [Tan et al. 2017] or pigment-based analysis of images [Tan et al. 2018a; Aharoni-Mack et al. 2017] can inspire more methods in this direction.

Network Scalability. Our forward network learns to connect ink labels and their ratios in a halftone to the corresponding spectrum. With a larger number of inks, the number of combinatorial training examples increases significantly. This drawback motivates developing neural networks that can generalize to arbitrary inks by learning the underlying color mixing from physical ink properties.

8 CONCLUSION

We presented an algorithm for selecting a desired number of inks for optimal spectral reproduction. Our technique handles the huge combinatorial search space and selects inks from within an ink library of potentially thousand and more. In contrast to previous approaches, our ink selection method is scalable and comes with measurable optimality. We also use the selected inks in an endto-end workflow for accurate, spectral reproduction of watercolor paintings.

We foresee a future where materials are selected and designed to fulfill multiple objectives across multiple modalities, e.g., selecting a limited number of inks for both appearance and mechanical purposes.

ACKNOWLEDGMENTS

The authors are grateful to Azadeh Asadi for creating all paintings in this work and Niloofar Maani for her kind help with experiments. We thank Roger Hersch for providing the base of, and Serdar Durdyyev for building on, the custom printing driver. This work benefited greatly from early discussions with Wojciech Matusik and Szymon Rusinkiewicz. We also thank the anonymous reviewers, Mathieu Hébert, Petar Pjanic, Sebastian Cucerca and Mojtaba Bemana for useful feedback.

REFERENCES

- Pankaj K Agarwal, Sariel Har-Peled, and Kasturi R Varadarajan. 2005. Geometric approximation via coresets. Combinatorial and computational geometry 52 (2005), 1–30.
- Elad Aharoni-Mack, Yakov Shambik, and Dani Lischinski. 2017. Pigment-based recoloring of watercolor paintings. In Proceedings of the Symposium on Non-Photorealistic Animation and Rendering. 1–11.
- Yağiz Aksoy, Tunç Ozan Aydin, Aljoša Smolić, and Marc Pollefeys. 2017. Unmixingbased soft color segmentation for image manipulation. ACM Transactions on Graphics (TOG) 36, 2 (2017), 19.
- Maria Joao Alves and João Clímaco. 2007. A review of interactive methods for multiobjective integer and mixed-integer programming. European Journal of Operational Research 180, 1 (2007), 99–115.
- Isaac Amidror and Roger D Hersch. 2000. Neugebauer and Demichel: Dependence and independence in n-screen superpositions for colour printing. *Color Research & Application* 25, 4 (2000), 267–277.
- Azadeh Asadi. 2017. Freelance artist. http://www.azadehasadi.com/index.html. Accessed: 2018-08-15.
- 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, Romain Rossier, and Roger D Hersch. 2012. Reducing the number of calibration patterns for the two-by-two dot centering model. In IS&T/SPIE Electronic Imaging. International Society for Optics and Photonics, 829208–829208.
- Vahid Babaei, Kiril Vidimče, Michael Foshey, Alexandre Kaspar, Piotr Didyk, and Wojciech Matusik. 2017. Color contoning for 3D printing. ACM Trans. Graph. (SIGGRAPH) 36 (2017).
- Pietro Belotti, Christian Kirches, Sven Leyffer, Jeff Linderoth, James Luedtke, and Ashutosh Mahajan. 2013. Mixed-integer nonlinear optimization. Acta Numerica 22 (2013), 1–131.
- Roy S Berns. 1993. Spectral modeling of a dye diffusion thermal transfer printer. Journal of Electronic Imaging 2, 4 (1993), 359–370.
- 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.
- Johannes Bisschop. 2012. AIMMS Optimization Modeling. Paragon Decision Technology. Brian Borchers and John E Mitchell. 1994. An improved branch and bound algorithm for mixed integer nonlinear programs. Computers & Operations Research 21, 4 (1994), 359–367.
- Alan Brunton, Can Ates Arikan, and Philipp Urban. 2015. Pushing the limits of 3d color printing: Error diffusion with translucent materials. ACM Trans. Graph. (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.
- D Connah, S Westland, and MGA Thomson. 2001. Recovering spectral information using digital camera systems. *Coloration technology* 117, 6 (2001), 309–312.
- E Demichel. 1924. Le procédé. Société Française de Photographie 26 (1924), 17-21.
- Oskar Elek, Denis Sumin, Ran Zhang, Tim Weyrich, Karol Myszkowski, Bernd Bickel, Alexander Wilkie, and Jaroslav Křivánek. 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

Ralph M. Evans. 1948. An introduction to color. John Wiley Sons, Inc. 262–267 pages. Mark D Fairchild. 2013. Color appearance models. John Wiley & Sons.

Christodoulos A Floudas. 1995. Nonlinear and mixed-integer optimization: fundamentals and applications. Oxford University Press.

- LLC Gurobi Optimization. 2018. Gurobi Optimizer Reference Manual. http://www. gurobi.com
- Roger David Hersch and Frédérique Crété. 2005. Improving the Yule-Nielsen modified Neugebauer model by dot surface coverages depending on the ink superposition conditions. In *Electronic Imaging 2005*. International Society for Optics and Photonics, 434–447.
- Rob Heylen, Dževdet Burazerovic, and Paul Scheunders. 2011. Non-linear spectral unmixing by geodesic simplex volume maximization. *IEEE Journal of Selected Topics* in Signal Processing 5, 3 (2011), 534–542.
- Kurt Hornik, Maxwell Stinchcombe, Halbert White, et al. 1989. Multilayer feedforward networks are universal approximators. *Neural networks* 2, 5 (1989), 359–366.
- R. W. G. Hunt. 2006. The reproduction of colour. John Wiley Sons. 163-179 pages.

255:16 • Ansari, Alizadeh-Mousavi, Seidel, Babaei

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.

Nirmal Keshava and John F Mustard. 2002. Spectral unmixing. *IEEE signal processing magazine* 19, 1 (2002), 44–57.

- Diederik Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. arXiv preprint arXiv:1412.6980 (2014).
- P. Kubelka and F. Munk. 1931. Ein Beitrag zur Optik der Farbanstriche. Zeitschrift f
 ür technische Physik 12 (1931), 593–601.
- Johan Lofberg. 2004. YALMIP: A toolbox for modeling and optimization in MATLAB. In 2004 IEEE international conference on robotics and automation (IEEE Cat. No. 04CH37508). IEEE, 284-289.
- Morteza Maali Amiri and Mark D Fairchild. 2018. A strategy toward spectral and colorimetric color reproduction using ordinary digital cameras. *Color* (2018).
- Laurence T Maloney. 1986. Evaluation of linear models of surface spectral reflectance with small numbers of parameters. *JOSA A* 3, 10 (1986), 1673–1683.
- Hugues Marchand, Alexander Martin, Robert Weismantel, and Laurence Wolsey. 2002. Cutting planes in integer and mixed integer programming. *Discrete Applied Mathematics* 123, 1-3 (2002), 397–446.
- Peter Morovič, Ján Morovič, 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.
- Victor Ostromoukhov. 1993. Chromaticity gamut enhancement by heptatone multicolor 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.
- Marios Papas, Christian Regg, Wojciech Jarosz, Bernd Bickel, Philip Jackson, Wojciech Matusik, Steve Marschner, and Markus Gross. 2013. Fabricating Translucent Materials Using Continuous Pigment Mixtures. ACM Trans. Graph. 32, 4 (July 2013), 146:1–146:12.
- Théo Phan Van Song, Christine Andraud, and Maria V Ortiz-Segovia. 2016. 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.
- Joanna L Power, Brad S West, Eric J Stollnitz, and David H Salesin. 1996. Reproducing color images as duotones. In Proceedings of the 23rd annual conference on Computer graphics and interactive techniques. ACM, 237–248.
- Robert Rolleston and Raja Balasubramanian. 1993. Accuracy of various types of Neugebauer model. In *Color and Imaging Conference*, Vol. 1993. Society for Imaging Science and Technology, 32–37.
- 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.
- Sebastian Ruder. 2016. An overview of gradient descent optimization algorithms. arXiv preprint arXiv:1609.04747 (2016).
- Mark Schmidt, Nicolas Le Roux, and Francis Bach. 2017. Minimizing finite sums with the stochastic average gradient. *Mathematical Programming* 162, 1-2 (2017), 83–112.
- Gaurav Sharma, Wencheng Wu, and Edul N Dalal. 2005. The CIEDE2000 color-difference formula: Implementation notes, supplementary test data, and mathematical observations. Color research and application 30, 1 (2005), 21–30.
- 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 (Dec. 2018), 271:1–271:15.
- Haichuan Song, Jonàs Martínez, Pierre Bedell, Noemie Vennin, and Sylvain Lefebvre. 2019. Colored fused filament fabrication. ACM Transactions on Graphics (TOG) 38, 5 (2019), 1–11.
- 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.
- Denis Šumin, Tobias Rittig, Vahid Babaei, Thomas Nindel, Alexander Wilkie, Piotr Didyk, Bernd Bickel, Jaroslav Křivánek, Karol Myszkowski, and Tim Weyrich. 2019. Geometry-Aware Scattering Compensation for 3D Printing. ACM Transactions on Graphics (Proc. SIGGRAPH) 38, 4 (July 2019), 111:1–111:14. https://doi.org/10.1145/ 3306346.3322992
- Chuanqi Tan, Fuchun Sun, Tao Kong, Wenchang Zhang, Chao Yang, and Chunfang Liu. 2018b. A survey on deep transfer learning. In *International conference on artificial neural networks*. Springer, 270–279.
- Jianchao Tan, Stephen DiVerdi, Jingwan Lu, and Yotam Gingold. 2018a. Pigmento: Pigment-based image analysis and editing. *IEEE transactions on visualization and* computer graphics 25, 9 (2018), 2791–2803.
- Jianchao Tan, Jyh-Ming Lien, and Yotam Gingold. 2017. Decomposing images into layers via RGB-space geometry. ACM Transactions on Graphics (TOG) 36, 1 (2017), 7.
- 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.

- John E Tyler and Arthur C Hardy. 1940. An analysis of the original Munsell color system. JOSA 30, 12 (1940), 587–590.
- Di-Yuan Tzeng and Roy S Berns. 1999. Spectral-based ink selection for multiple-ink printing II. Optimal ink selection. In *Color and Imaging Conference*, Vol. 1999. Society for Imaging Science and Technology, 182–187.
- Gunter Wyszecki and Walter Stanley Stiles. 1982. Color Science. Vol. 8. Wiley New York.
- John Arthur Carslake. Yule. 1967. Principles of color reproduction: applied to photomechanical reproduction, color photography, and the ink, paper, and other related industries. Wiley.