Predicting Stacked-Filament Color from Independently Measured Filament Properties

Garnet Liu

September 1, 2026

Abstract

Consumer multi-material fused deposition modeling (FDM) printers reproduce color images by stacking translucent filament layers; stack color must be predicted before printing. Existing workflows predict with a scalar transmission distance (TD) plus manual tuning, or with dense lookup tables (LUTs) per palette. We ask whether per-filament measurements alone predict arbitrary stack combinations. We present a forward model that composes nominal RGB, per-channel TD, an effective attenuation coefficient per filament, a shared monochromatic scattering prefactor, and stack geometry into stack color through a loss-allocation stacking rule; when only a scalar TD is available, the remaining coefficients are estimated on one printed plate, a LUT-free calibration of the same analytic model. Staircase-measured per-channel TDs predict 256-cell clear-filament permutation plates with no parameters fitted on the plates: 18.6 versus 47.2 CIEDE2000 color difference (ΔE_{00}) at best, 42–60% reduction on all six plate/backing combinations (Table 4). Eight scalars fitted on a plate of five arbitrary non-CMYWK filaments, RGB and per-channel TDs measured off-plate, explain 780 cells at 6.9 ΔE_{00} (white 5.3, black 8.5); a RAW replication reaches 5.5. A staircase, dual-backing plate, and phone camera supply these properties; a single-point optical filament sensor implies a 75× absorption scale across CMYW filaments and saturates on transparent ones, so it cannot deliver them (Section 5.4). On ColorChecker targets across three prints, the model reaches 16.69 versus 30.92 ΔE_{00} for the community TD-table formula (Table 7); coefficients are process-conditioned effective values in the sRGB-encoded domain, within which one shared model predicts arbitrary stacks from one measurement per filament.

1 Introduction

Consumer multi-material fused deposition modeling (FDM) printers reproduce color images by stacking translucent filament layers in Z. The community has converged on a data format: a nominal RGB color plus a scalar transmission distance (TD) per filament [8, 34]. Tools differ in how they blend the two values.

Three families of practice answer the prediction problem. Interactive consumer software treats each filament as a scalar Beer–Lambert (BL) transmission and relies on manual tuning [2, 8, 21, 34]. Palette-calibrated workflows photograph a board per palette and predict from a dense measured lookup table (LUT) [17] or fitted palette models [18]. Physically based optical characterization recovers accurate material parameters under laboratory capture [5, 11, 20]. None composes independently measured consumer-grade filament properties into arbitrary stacks through one shared formula.

This paper separates the measurement of a filament from the composition of a stack. A single-material staircase yields per-channel TDs, a dual-backing structured plate supplies the remaining coefficients, and a phone camera captures both (Fig. 1). One shared forward model then predicts any stack: any layer count, order, or combination.

We answer three research questions.

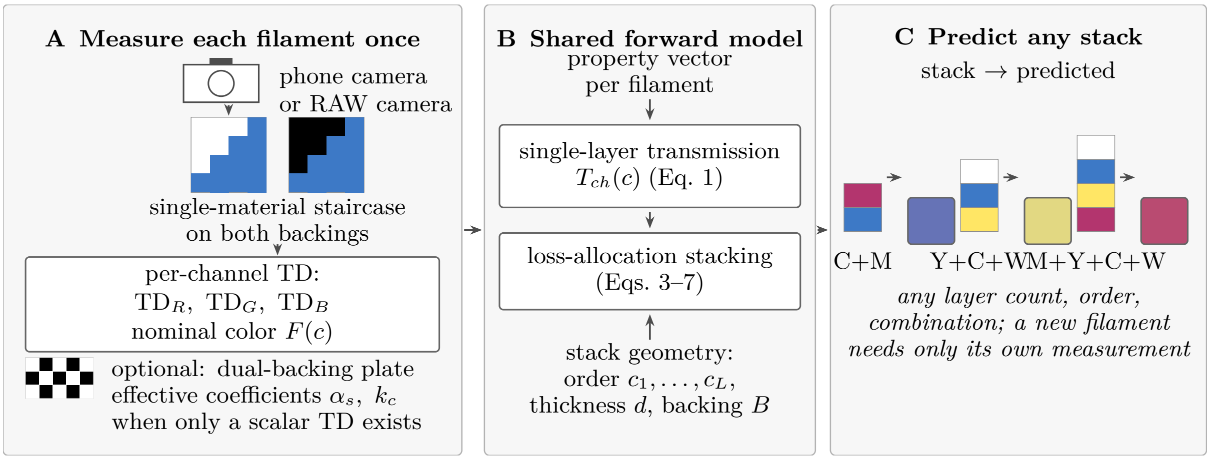

Figure 1: Measure each filament once, then predict any stack. Left: single-filament artifacts (staircase; optional dual-backing plate) photographed with a phone or RAW camera supply nominal color, per-channel transmission distance and, when only a scalar TD exists, effective coefficients. Middle: one shared forward model (single-layer transmission and loss-allocation stacking). Right: predictions for arbitrary layer counts, orders and combinations; a new filament requires only its own measurement.

1. *RQ1.* Do independently measured filament properties predict unseen stack combinations? Staircase-measured per-channel TDs predict 256-cell transparent permutation plates at 18.6 versus 47.2 ΔE_{00} at best, 42–60% lower error on all six plate/backing combinations, with no fitting on the plates (Table 4).

2. *RQ2.* Does the representation transfer to new filament identities? Eight scalars fitted on a plate of five arbitrary non-CMYWK filaments, with RGB and per-channel TDs measured off-plate, explain 780 cells at 6.9 ΔE_{00}; a RAW replication reaches 5.5 (Fig. 4).

3. *RQ3.* How well can the required properties be obtained with consumer tools? A single-point optical filament sensor implies a 75× absorption scale across CMYW filaments and saturates on transparent ones, while the phone-camera procedure recovers the black-filament coefficient k_K at 23.66 ± 0.51 across four prints (Section 5.4, Table 8).

Our contributions are threefold.

1. **A forward model** that predicts stacked-filament color from independently measurable per-filament properties (nominal RGB, per-channel transmission distance, effective attenuation coefficient) and stack geometry through a loss-allocation stacking rule. Calibration, where needed, is LUT-free calibration of an analytic forward model.

2. **Transfer validation.** Staircase measurements predict unseen combinations with no fitting on the target plates (Section 5.1); the same representation carries five arbitrary non-CMYWK filaments with eight fitted scalars (Section 5.2); physical prints lead the community TD-table formula (Section 5.3).

3. **A consumer-accessible characterization procedure** (single-material staircase, dual-backing structured plate, phone camera) that supplies the properties direct sensors cannot, with phone-JPEG and RAW captures reaching the same accuracy.

Section 2 reviews related work; Section 3 presents the forward model, Section 4 the characterization procedure, Section 5 the results by research question, Section 6 the discussion.

2 Related Work

2.1 Consumer stacked-filament tools

HueForge popularized stacked-filament color images on consumer FDM printers through the transmission-distance workflow: users assign a TD value per filament and tune layer transitions until the preview matches intent [8]. Kromacut implements the HueForge-compatible base-10 BL transmission $T = 10^{-d/\mathrm{TD}}$ for single-layer thickness d as an open library [34]. Slicer integrations carry layered color into mainstream tools: Bambu Studio Color Mixing [2] predicts layer-alternation color following the OrcaSlicer-FullSpectrum fork, whose blending is powered by the Kubelka–Munk-based FilamentMixer library [12, 23]; Prusa ColorMix previews layer-interleaved color with a Yule–Nielsen halftone model fitted to measured swatches [21, 22]. OpenColor couples a four-flux adding-doubling radiative-transfer model with Gaussian-process residuals [19]. Across these tools the input is nominal RGB plus a scalar TD; prediction rests on manual tuning and preview features, not on a calibrated stacked-transmission predictor.

2.2 Palette-calibrated prediction

Lumina prints a multi-color calibration board, photographs it, and reproduces target colors by nearest-neighbor lookup in the measured table [17]; the table is the model, and new filaments or thicknesses call for a new board. ChromaPrint3D photographs a calibration board to populate a color database, fits a photo-calibrated forward model on it, and matches print recipes [18]. Both systems measure mixed palettes, and both bind prediction to one palette.

2.3 Optical characterization of printing materials

Color 3D printing research spans contoning [1], high-fidelity surface color [3], dual-extrusion color mixing [24], line-based halftoning [15], scattering-aware texture reproduction [4], geometry-aware compensation [31], neural acceleration [25], multispectral ink reproduction [28], thickness-controlled layered color on filament printers [29], and automated filament inking [16]; thin white cores measurably shift surface color on material-jetting systems [7]. Layered-media color prediction builds on the Kubelka–Munk two-flux lineage [13, 14, 26, 36]. Four-flux models predict 2.5D layered prints [32]; FDM transmittance varies strongly with material and wavelength [6]; spectrophotometer measurements of printed slabs at several thicknesses recover scattering and absorption coefficients [20]; inverse rendering recovers volume-transport parameters of photopolymers [5]; photographing a few thicknesses yields spectral coefficients at low cost [10]. The closest academic counterpart is the scattering-aware calibration of Iser et al. [11]: one printable calibration target, captured with a camera or spectrometer, recovers per-resin albedo and extinction on a full-color photopolymer printer. We target FDM filaments, phone-camera capture, and a property interface, nominal RGB plus per-channel TD, that the consumer ecosystem already exchanges.

2.4 Position

Prior work either measures materials with laboratory instrumentation or predicts stacks from palette-level data; this work composes the first into the second. The per-filament property interface accepts better instrumentation, a spectrometer or a RAW pipeline, without changes to the forward model. Table 1 contrasts the data dependencies of the five systems.

Table 1: Data dependency of stacked-color prediction systems. Ours measures each filament independently (nominal RGB, per-channel TD) and predicts new combinations without further measurement; when a scalar TD is the available input, eight scalars are fitted on one dual-backing plate.

Method	Single-filament input	Mixed-palette measurement	New-palette workflow	New-combination measurement	Forward model
HueForge / Kromacut	RGB + scalar TD	no	supply and tune values	no	scalar BL, manual tuning
Lumina (dense LUT)	optional	calibration board	print new board	often yes	lookup
ChromaPrint-3D	calibrated material data	board + color database	recalibrate palette	training data required	fitted palette model
Iser et al. 2025	per-resin albedo + extinction from one printed target (camera or spectrometer)	single calibration target	re-characterize resins	no	radiative-transfer fit
Ours	independently measured RGB + per-channel TD; with a scalar TD, 8 scalars fitted on one plate	optional (dual-backing plate)	measure new filaments	no	shared analytic model

3 Forward Model

3.1 Material-property interface

The predictor consumes a fixed per-filament property vector and the stack geometry. For each filament c the vector holds the nominal color $F_{ch}(c)$, the transmission distance td_c (per-channel $td_{ch,c} = TD_R$, TD_G, TD_B), and the effective attenuation coefficient k_c; the monochromatic scattering prefactor α_s and the remap pair (s_{td}, γ_{td}), Eq. (2), are global. A stack is an ordered sequence $c_1, \ldots, c_L$, bottom layer first, at single color-layer thickness d, over N primaries.

The nominal color is the filament printed opaque: the eight-layer cell of a single-filament ramp plate (Section 4), photographed on white backing with white-balance correction. It enters as an appearance proxy through the absorbance $A_{ch}(c) = (255 - F_{ch}(c))/255 \in [0, 1]$, not as a reflectance albedo. We keep k_c rather than the radiative-transfer coefficients μ or σ because the model is not radiative transfer. Throughout, $d = 0.08$ mm (0.04–0.20 mm supported), $B_{ch} = 1$ (white) or ≈ 0.05 (black cardstock), $N = 5$ (cyan, magenta, yellow, white, black; 4–16 supported); Table 2 collects the notation.

3.2 Single-layer transmission

The layer transmission has two forms, selected by what is measured. (i) With staircase-measured per-channel distances, $T_{ch}(c) = 10^{-d/td_{ch,c}}$; nothing is fitted on the target plates. (ii) With a scalar distance, from a sensor or a published table,

$$T_{ch}(c) = \exp\left(- \left(\frac{\alpha_s}{td_c} + k_c \cdot A_{ch}(c) \right) \cdot d \right), \tag{1}$$

with α_s and k_c estimated from a printed plate (Section 3.5). The monochromatic term α_s/td_c is identical across R, G, B, keeping white filament semi-opaque; the wavelength-dependent absorption term $k_c \cdot A_{ch}(c)$ varies per channel, separating cyan, magenta, and yellow (Fig. 2(a)). The monochromatic term is exact for a broadband scatterer (TiO_2 white pigmentation) and approximate for tinted scatterers; Section 5 tests it on colored scatterers. A scalar reading

Symbol	Meaning	Unit
c	filament (material) index	
$ch \in \{R, G, B\}$	color channel	
$F_{ch}(c)$	nominal per-channel color (opaque ramp-cell measurement)	8-bit sRGB
$A_{ch}(c)$	per-channel absorbance proxy, $(255 - F_{ch}(c))/255$	dimensionless
d	single color-layer thickness	mm
td_c	transmission distance of filament c	mm
$td_{ch,c}$	per-channel transmission distance (TD_R, TD_G, TD_B)	mm
$TD_{sensor,c}$	single-point optical-sensor reading (TD1S)	mm
α_s	monochromatic scattering prefactor (global)	dimensionless
k_c	effective attenuation coefficient of filament c	mm^{-1}
$T_{ch}(c)$	single-layer per-channel transmission, Eq. (1)	dimensionless
s_{td}, γ_{td}	TD power-law remap scale and exponent, Eq. (2)	dimensionless
L	number of layers in the stack	count
$c_1, \ldots, c_L$	stack composition, bottom layer first	
N	number of filament primaries	count
$R_{ch}(i)$	per-channel light remaining after the first i layers	fraction
$\ell_{ch}(i)$	per-channel loss allocated at layer i	fraction
$\ell_{ch}(bg)$	per-channel loss reaching the backing ($= R_{ch}(L)$)	fraction
$L_{ch}(i)$, $L_{ch}(bg)$	normalized loss-allocation weights	fraction
B_{ch}	per-channel backing reflectance (1 white; ≈ 0.05 black cardstock)	dimensionless
$P_{ij}(\theta)$	predicted color of grid cell (i, j)	8-bit sRGB
G_{ij}, G	measured color of grid cell (i, j); the plate measurement grid	8-bit sRGB
θ	vector of free parameters	
$J(\theta)$	fitting objective, Eq. (8)	ΔE_{76}
ΔE_{00}	CIEDE2000 color difference (reporting metric)	ΔE
ΔE_{76}	CIE 1976 color difference (optimization metric)	ΔE

Table 2: Notation of the forward model.

$TD_{sensor,c}$ enters through a power-law remap,

$$td_c = s_{td} \cdot (TD_{sensor,c})^{\gamma_{td}}, \tag{2}$$

fitted jointly with the optical coefficients ($s_{td} \approx 1.48$, $\gamma_{td} \approx 0.20$ on CMYWK). Consumer tools exchange the scalar form $T = 10^{-d/TD}$, one channel-independent distance per filament [34]; we call it the TD-table formula and use it as the prediction baseline in Section 5.

3.3 Stack composition

The stacking rule allocates the per-channel light budget of a single downward pass across layers and backing (Fig. 2(b)). It is a heuristic single-pass loss-allocation rule, justified empirically in Section 5. Let $R_{ch}(i)$ denote the per-channel light remaining after the first i layers:

$$R_{ch}(0) = 1, \qquad R_{ch}(i) = R_{ch}(i-1) \cdot T_{ch}(i) \quad \text{for } i = 1, \ldots, L. \tag{3}$$

The loss allocated at layer i is

$$\ell_{ch}(i) = R_{ch}(i-1) \cdot (1 - T_{ch}(i)), \tag{4}$$

and the residual reaching the backing is

$$\ell_{ch}(bg) = R_{ch}(L). \tag{5}$$

Normalized by their per-channel sum, the weights are

$$L_{ch}(i) = \frac{\ell_{ch}(i)}{\sum_{j=1}^{L} \ell_{ch}(j) + \ell_{ch}(bg)}, \qquad L_{ch}(bg) = \frac{\ell_{ch}(bg)}{\sum_{j=1}^{L} \ell_{ch}(j) + \ell_{ch}(bg)}. \tag{6}$$

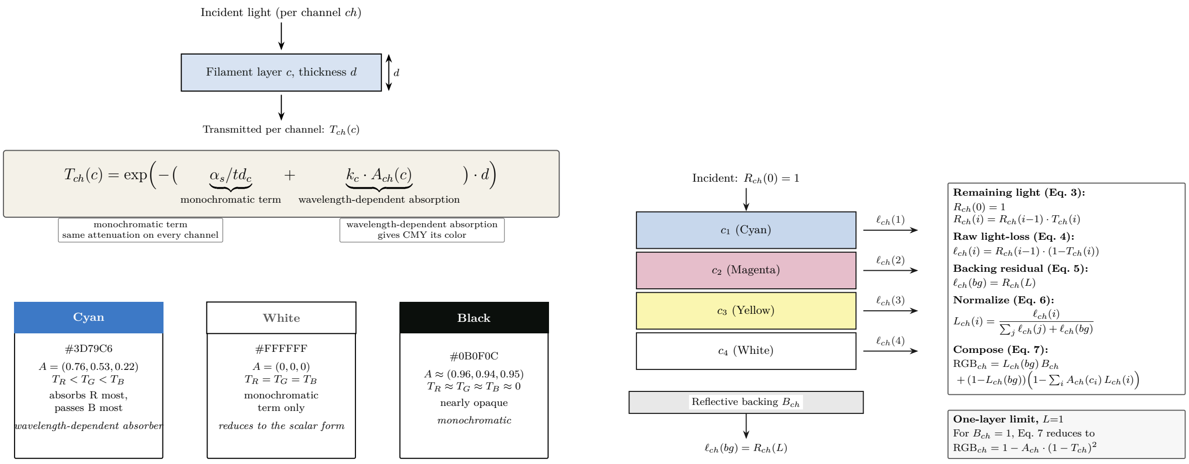

(a) (b)

Figure 2: Forward model. (a) The single-layer transmission of Eq. (1): the monochromatic term α_s/td_c and the wavelength-dependent absorption term $k_c \cdot A_{ch}(c)$. (b) The stacking rule of Eqs. (3)–(7): a single downward pass allocates the per-channel light budget across layers and backing.

The output mixes the backing with the layer-absorbance-weighted stack color:

$$\mathrm{RGB}_{ch} = L_{ch}(bg) \cdot B_{ch} \;+\; \big(1 - L_{ch}(bg)\big) \cdot \left(1 - \sum_{i=1}^{L} A_{ch}(c_i) \cdot L_{ch}(i)\right), \tag{7}$$

the 8-bit output being $\mathrm{RGB}_{ch} \times 255$ clipped to $[0, 255]$.

3.4 Color-space convention

All model quantities (nominal RGB, A_{ch}, per-layer transmission, output) and all measurements and ΔE comparisons live in the 8-bit sRGB-encoded domain; this is the domain in which filament colors and TD values are published and exchanged in the consumer ecosystem, and in which camera images are delivered. Equations (1)–(7) are therefore an empirical model in encoded space; the fitted coefficients (α_s, k_c, td_c) absorb the encoding nonlinearity and are process-conditioned. Staircase transmission is measured in linear light, with the sRGB electro-optical transfer function (EOTF) and EXIF exposure normalization, and converted to the encoded convention with $TD_{enc} = 2.2 \cdot TD_{lin}$. All reported evaluation color differences are CIEDE2000 [9, 27], written ΔE_{00}. To maintain numerical smoothness during fitting and avoid the hue-angle boundary discontinuity of ΔE_{00}, the optimization objective evaluates the continuous Euclidean CIELAB distance (ΔE_{76}), separating the optimization metric from the standard perceptual reporting metric.

3.5 Fitting objective

When direct sensors cannot supply a property, it is estimated from a printed plate. Let G be the plate's measured per-cell color grid, G_{ij} the color of cell (i, j), and $P_{ij}(\theta)$ the prediction for its stack code. The objective minimizes the mean Euclidean CIELAB color difference (ΔE_{76}) plus a weak ordering regularizer:

$$J(\theta) = \frac{1}{|G|} \sum_{(i,j)\in G} \Delta E_{76}\big(P_{ij}(\theta), G_{ij}\big) \;+\; \lambda_k \cdot \Pi_k(\{k_c\}), \tag{8}$$

Material	TD1S (mm)	TD_R	TD_G	TD_B	Fit R^2
Kingroon transparent PLA	100 (sat.)	17.9	18.9	17.2	0.988
Ziro light cyan clear	48.9	1.35	2.54	4.09	0.998
Sunlu transparent yellow	100 (sat.)	22.1	15.4	5.38	0.963
iSANMATE light pink	100 (sat.)	11.2	2.03	2.51	0.990

Table 3: Per-channel transmission distances (mm) recovered from single-material staircases for four transparent filaments, with the raw TD1S sensor reading for comparison; R^2 is the mean log-linear fit quality per material.

where Π_k favors the ordering $k_K > k_W > k_M > \{k_C, k_Y\}$ with $\lambda_k = 0.5$, and sensor-seeded fits carry a weak TD1S ordering prior on the td_c (Section 4). All fits use differential evolution [30] through SciPy [33]. Free parameters vary by experiment: none on the transparent plates, eight scalars on the five-arbitrary-filament plate, eight per PLATE-06 print (Section 5). Bound, prior, and single-backing ablations are tabulated in the supplementary material.

4 Material Characterization Used in This Work

4.1 Staircase per-channel transmission distances

Two single-filament artifacts characterize each material. A ramp plate prints the filament at 1, 2, 4, 6, and 8 layers on white backing; its eight-layer cell gives the nominal color $F_{ch}(c)$ and its attenuation over layer count gives the rate against which sensor readings are interpreted (Section 5.4). A single-material staircase recovers per-channel transmission distances from one white/black photograph pair. The clear track prints eight steps of 0.32–10 mm design thickness at 0.28 mm layers on the Anycubic Kobra X; the opaque track prints twelve steps of 0.08–4.8 mm at 0.08 mm layers on the Bambu Lab H2C. The same step across backings gives the round-trip transmission

$$T^2 = \frac{S_w - S_b}{BG_w - BG_b},$$

with S and BG the step and local-background readings under white (w) and black (b) backing; a log-linear fit over the thin range ($d \leq 3$ mm) yields TD_R, TD_G, TD_B at R^2 0.963–0.998 (Table 3). Where illumination geometries differ, an offset-augmented estimator $s(d) = C + A \cdot 10^{-d/TD}$ uses all steps. Channel orderings match pigment physics unconstrained: cyan $TD_R < TD_G < TD_B$; yellow $TD_B \ll TD_G < TD_R$; pink $TD_G \approx TD_B \ll TD_R$; neutral near-equal. Rounding to the layer grid shifts well-sampled channels by at most 1.2% with every ordering preserved; per-step fits and confidence intervals are in the supplementary material. The TD1S (BIQU AJAX optical filament sensor) saturates at its 100 mm ceiling on three of four materials (Table 3).

4.2 Dual-backing structured plate

One structured plate under two backings supplies enough constraints to estimate the remaining coefficients. PLATE-06 lays out $30 \times 26 = 780$ cells covering the exhaustive 1–4-layer CMYWK codes with corner markers; each print is photographed once under white and once under black out-of-plate backing. Border-anchored photometric normalization removes the ≈1.8-stop exposure difference (1/215 s vs. 1/60 s) with a measured border-leak discount. The objective is $J_{white} + J_{black}$ over eight parameters per print: α_s, s_{td}, γ_{td}, k_C through k_K.

Plate (profile)	Backing	Per-channel TD	Scalar TD	Reduction (%)
CMYG (default)	white	21.5	45.5	52.8
CMYG (adjusted)	white	**18.6**	47.2	60.5
CMYW (crossed)	white	28.3	50.8	44.3
CMYG (default)	black	23.1	40.0	42.2
CMYG (adjusted)	black	22.1	38.7	43.0
CMYW (crossed)	black	31.8	61.8	48.6

Table 4: Clear-plate cross-validation: mean camera-pipeline ΔE_{00} (CIEDE2000) per 256-cell four-layer permutation plate. "Per-channel TD" uses staircase-measured $TD_R/TD_G/TD_B$ with no parameters fitted on the plate; "Scalar TD" replaces them with one scalar TD per filament. Reduction is the relative error decrease of the per-channel route.

4.3 Capture

The same interface accepts a consumer phone and a RAW pipeline. Primary capture uses an iPhone 16 Plus (JPEG, auto exposure and white balance), build-plate side up in a consumer LED light box; glare masking excludes pixels with $V > 245$ and cells are sampled over their central 40–60%. A Sony A7M4 replication decodes RAW deterministically to sRGB. An in-frame ColorChecker fits a per-shot camera-to-CIELAB transform at leave-one-out residual ΔE_{00} 2.64, so reported ΔE values are camera-pipeline-relative. CMYWK validation prints use Bambu PLA Basic filament on Bambu Lab printers (H2C primary, Bambu P1S second-printer replication); clear-track prints use the Kobra X.

5 Results

The results answer RQ1 (Section 5.1), RQ2 (Section 5.2), and RQ3 (Section 5.4), then the stacking-rule comparison (Section 5.3) and an application print (Section 5.5). All plates follow the print and capture protocol of Section 4.

5.1 RQ1: Prediction of Unseen Stacks from Single-Filament Measurements

Profile-matched staircases supply per-channel TDs, and the stacking rule composes them with no parameters fitted on the target plates: 256-cell clear permutation plates (four-layer codes, 0.84 mm color layers) under two slicer profiles and a crossed-infill variant, on both backings (Table 4, Fig. 3).

Measured per-channel TDs cut the scalar-TD error by 42–60% on all six plate/backing combinations; the best case reaches 18.6 versus 47.2 ΔE_{00} (CIEDE2000). Black-backing photographs are EXIF-normalized; the plates carry no white reference and are not white-balanced. Per-cell diagnostics are in the supplementary material.

5.2 RQ2: Transfer to New Filament Identities

Five arbitrary opaque filaments, none used in model development, fill the five roles of the 30×26 five-arbitrary-filament plate (RANDMIX5): translucent grey, matte magenta, matte ice-blue, beige, purple (Table 5). Nominal colors and per-channel TDs come from independent staircases; eight scalars are fitted on the plate (α_s, s_{td}, γ_{td}, one k_c per filament); the claim is eight fitted numbers against a 780-entry measured table.

The model reaches all 780 cells at combined ΔE_{00} 6.9 (white 5.3, black 8.5; median 6.9, maximum 15.4), less than half the default-coefficient error of 15.9 (Fig. 4). The CMYWK primaries reach 9.23–9.76 under the identical protocol (Section 5.4).

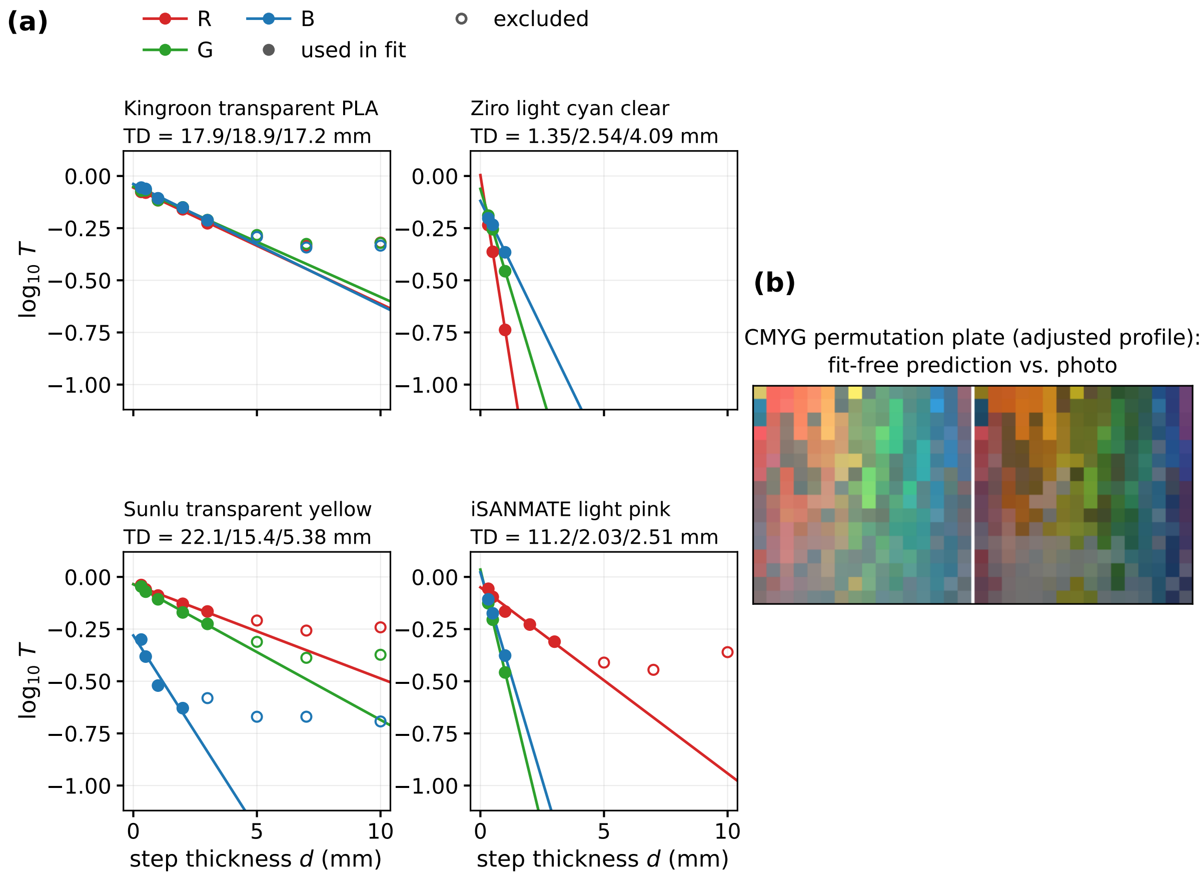

Figure 3: Transparent track. (a) Staircase per-channel TD fits for four transparent filaments; the log-linear fit uses the thin-range steps. (b) Fit-free prediction versus photograph of a 256-cell clear permutation plate; cell structure and hue families align.

Slot	Filament	Nominal color	td_{stair} (mm)	td_{eff} (mm)	k_c
C	grey (Panchroma translucent)	#7B7872	3.44	9.45	0.50
M	magenta (Sunlu matte)	#B43C74	0.37	1.59	7.24
Y	ice-blue (matte)	#A9DDF3	1.09	3.75	7.90
W	beige (Bambu)	#F6E5CC	1.97	6.04	6.07
K	purple (Bambu)	#5950B4	0.59	2.30	1.06

Table 5: The five-arbitrary-filament plate (RANDMIX5). td_{stair} is the staircase-measured TD, locked during the fit; $td_{\text{eff}} = s_{td} \cdot td_{\text{stair}}^{\gamma_{td}}$ with the fitted remap. Fitted on the plate: $\alpha_s = 6.09$, $s_{td} = 3.50$, $\gamma_{td} = 0.80$, and the five k_c; nominal colors and td_{stair} are measured off-plate.

The staircases resolve wavelength dependence. Magenta transmits red while absorbing green at td_G 0.17 mm (R^2 0.9999); its per-channel TDs span 0.17–0.63 mm against a single broadband reading of 1.8 mm, which no scalar TD represents. For near-neutral beige the single TD1S reading, 3.5 mm, matches its fitted green-channel TD of 3.3 mm.

A Sony A7M4 RAW re-capture of the same plate, refit under the identical protocol, reaches ΔE_{00} 5.5 (white 4.9, black 6.1) from a 15.4 baseline (details in the supplementary material): the accuracy belongs to the model and the staircases, not to one camera pipeline. The eight coefficients are what a laboratory photometric instrument would measure directly; under such capture the model runs in pure prediction.

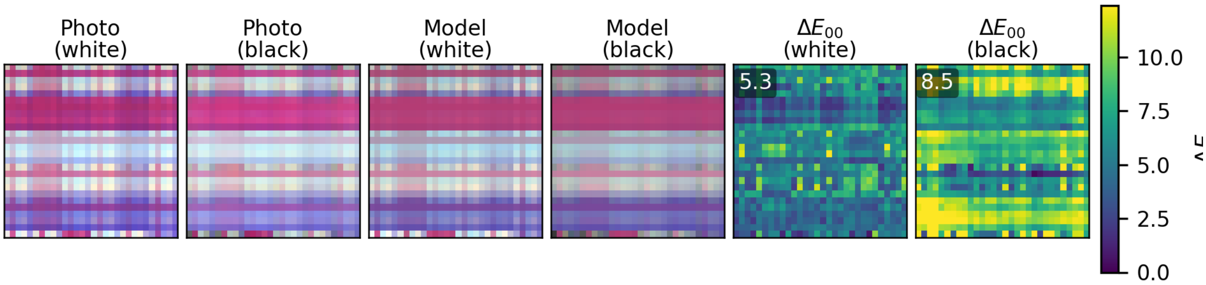

Figure 4: RANDMIX5: photographs, model reconstructions, and per-cell ΔE_{00} maps of the 780-cell arbitrary-filament plate under both backings; combined mean 6.9, white 5.3, black 8.5. All five filaments are unseen non-CMYWK materials; nominal RGB and per-channel TDs come from independent staircases.

Stacking form	Held-out mean ΔE_{00}
Loss allocation (ours)	**15.96**
Sequential compositing	25.96
Simplified Kubelka–Munk	27.41
Calibrated scalar Beer–Lambert	15.91

Table 6: Stacking-rule ablation on the COMPARE-01 plate: mean held-out ΔE_{00} (CIEDE2000) over 72 ColorChecker patches whose codes are not used for fitting. The first three rows share the same transmission model and fitted coefficients; the compositing rule alone differs. The calibrated scalar Beer–Lambert row is a one-coefficient scalar model fitted on the same plate.

5.3 Stacking Rule and Comparison with the Community Formula

With the transmission model and fitted coefficients fixed, Table 6 swaps the compositing rule alone; evaluation is held out on 72 ColorChecker patches not used for fitting. Loss allocation reaches 15.96 ΔE_{00}, sequential compositing 25.96, a simplified Kubelka–Munk form 27.41; the rule alone moves held-out error by 10 ΔE_{00}. On a fixed CMYWK palette with on-plate calibration, a one-coefficient scalar model matches this error; the per-channel representation is what the transparent and arbitrary-filament plates need, where no palette-level fit exists or the filaments are arbitrary.

The head-to-head prints 24 ColorChecker targets per method on one plate, three independent sessions captured together (Table 7). The model cuts the community TD-table formula's error [34] by 45–47% on every print, a mean margin of +14.23 ΔE_{00} (16.69 versus 30.92). The uncalibrated scalar Beer–Lambert formula with the same nominal colors trails by 1.1–1.6 ΔE_{00} on each print. A P1S reprint preserves the ordering: 11.31, 14.25, 21.20 ΔE_{00}.

5.4 RQ3: Characterization Accuracy with Consumer Tools

The properties the model needs are obtainable with a phone camera; a single-point transmission sensor does not supply them. The TD1S optical sensor reads one broadband TD: 6.1, 2.9, 5.0, and 2.0 mm for white, magenta, yellow, cyan. Combined with the attenuation rates measured on the single-filament ramp plate (Section 4), these readings imply a global α spanning 3.0 to 226.5, a 75× range; black saturates the sensor's 0.1 mm floor, and three of the four transparent filaments saturate its 100 mm ceiling.

The dual-backing plate supplies them (Table 8, Fig. 5): four independent prints recover the

Print	Ours	Uncalibrated scalar BL	TD-table formula
w-1	16.89	17.97	31.46
w-2	16.59	18.18	30.18
w-3	16.59	17.71	31.13
Mean	**16.69**	17.95	30.92

Table 7: COMPARE-01 physical head-to-head: mean camera-pipeline reproduction ΔE_{00} (CIEDE2000) over 24 ColorChecker targets per print; three independent print sessions, one capture session. The TD-table formula is the community base-10 Beer–Lambert with sensor TDs [34].

Print	α_s	k_C	k_M	k_Y	k_W	k_K	Mean ΔE_{00}	Default
PLATE-06-H2C-A	2.23	3.50	4.31	3.66	6.32	23.19	9.23	18.91
PLATE-06-H2C-B	2.21	2.48	3.58	2.79	10.93	24.43	9.56	18.30
PLATE-06-H2C-C	2.54	2.57	3.92	2.93	14.00	23.79	9.73	18.38
PLATE-06-H2C-D	1.83	2.61	3.98	2.92	9.81	23.23	9.76	18.67

Table 8: Dual-backing fitted coefficients for four independent PLATE-06-H2C prints (780 cells each); errors are mean camera-pipeline ΔE_{00} (CIEDE2000), combined over both backings. "Default" is the same metric with default coefficients and no fitting. The ordering $k_K > k_W > k_M > \{k_C, k_Y\}$ is recovered on every print, k_K averages 23.66 (sd 0.51), and the combined error spans 9.23–9.76 (mean 9.57, sd 0.21), 47–51% below the default coefficients.

ordering $k_K > k_W > k_M > \{k_C, k_Y\}$ every time and cut the combined error 47–51% below the default coefficients, to 9.23–9.76 ΔE_{00} (mean 9.57, sd 0.21). Re-photographing print D with the Sony RAW pipeline and refitting gives 9.62 versus 9.76 ΔE_{00} with materially different coefficients (s_{td} 4.53 versus 1.18): the transmission constrains the product $\alpha\, d/\mathrm{TD}$, so accuracy transfers while the decomposition is capture-conditioned. Bound-widening and prior ablations leave k_K and the error unchanged (supplementary material).

5.5 Application

The forward model drives an end-to-end image-to-print pipeline: K-means in CIELAB extracts the palette; each pixel takes the stack code minimizing ΔE_{00} to the available filaments; greedy meshing merges same-code regions (70–80% box-count reduction); output is STL or 3MF. The printed translucent CMYW panel is muted front-lit and saturated stained-glass backlit (Fig. 6). The same model supplies both the code search and the predicted appearance.

6 Discussion

Rather than requiring high-dimensional radiative-transfer parameters, the model targets a low-parameter interface: nominal RGB plus transmission distance, the pair the consumer ecosystem already exchanges and a phone can measure. The staircases follow the few-sample, multi-thickness measurement spirit of [10, 11]; the permutation plates validate stacking generalization, not the coefficients. The same interface accepts physically based characterization as its measurement backend.

Four-flux, inverse-rendering, and spectrophotometric characterization [5, 11, 20, 32] deliver exactly the quantities the interface consumes. As such instrumentation reaches the consumer, the dual-backing plate gives way to direct measurement; the forward model is unchanged.

The absorption ordering is not an artifact of one vendor's CMYWK set: five arbitrary non-

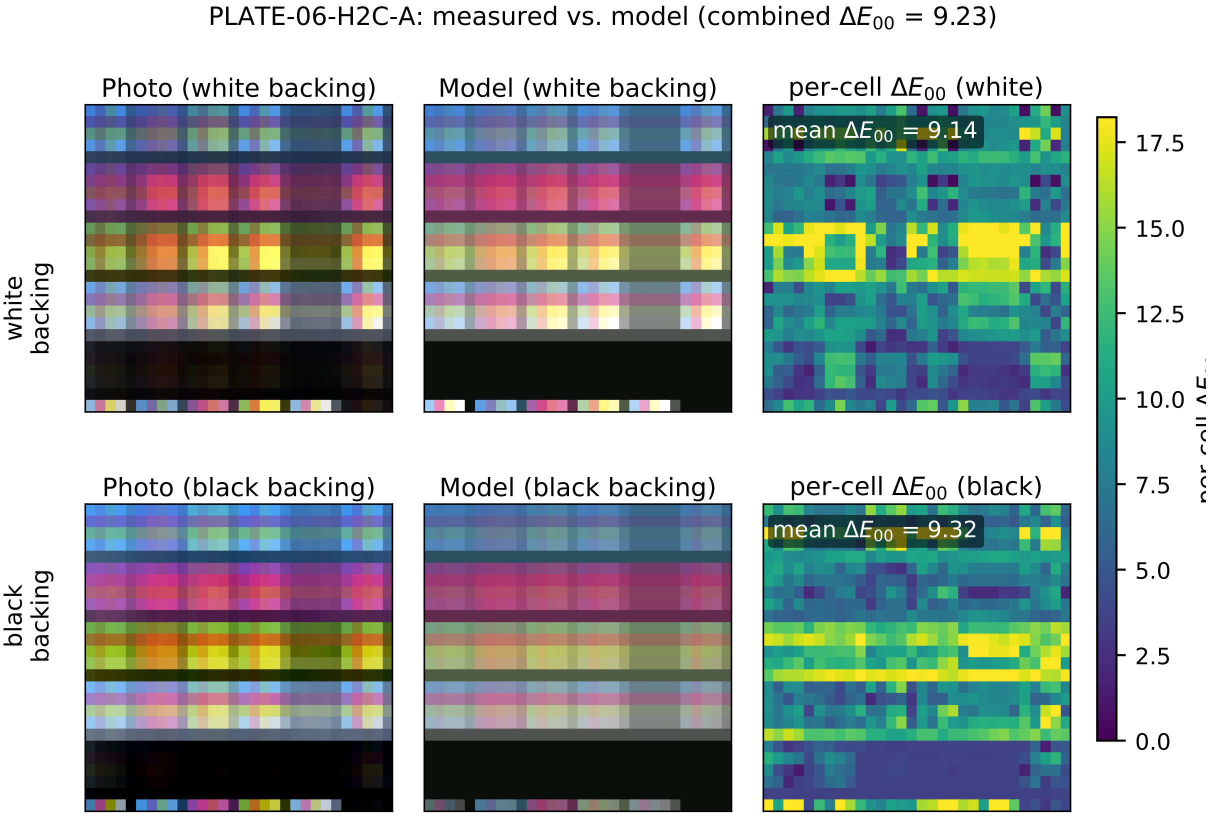

Figure 5: PLATE-06-H2C-A under both backings: photographed per-cell RGB, model reconstruction, and per-cell ΔE_{00} map on a unified scale. Combined mean ΔE_{00} 9.23 over all 780 cells (median 8.5; 65% of cells below 10).

Figure 6: Application print. Source image (left), the printed translucent layered CMYW panel front-lit (center), and the same panel backlit (right), where subtractive transmission through the stacked filaments yields saturated stained-glass color.

CMYWK filaments recover physically coherent wavelength-dependent absorption under the same protocol (magenta: td_G 0.17 mm while transmitting red; Table 5). The monochromatic scattering term, exact for broadband scatterers, also carried the tinted scatterers (ice-blue, beige) in the same test.

Our evaluation defines a clear operational envelope: (1) across repeated prints within a batch, characterization exhibits strong repeatability (error standard deviation 0.21 ΔE_{00} and k_K stability of 23.66 ± 0.51); (2) the recovered coefficients serve as process-conditioned effective optical quantities anchored to the print and capture setup (with camera-pipeline colorimetry anchored at a ΔE_{00} residual of 2.64); and (3) the forward formulation captures the macroscopic transmissive behavior of stacked layers with computational efficiency suitable for interactive slicer workflows.

7 Conclusion

This paper presented a forward model mapping independently measured filament properties and stack geometry to arbitrary stack color. Staircase measurements predict unseen transparent stacks at 18.6 versus 47.2 ΔE_{00}, 42–60% below scalar-TD error on all six cases (Table 4). Five arbitrary non-CMYWK filaments follow at 6.9 ΔE_{00} with eight fitted scalars (Fig. 4); physical prints lead the community TD-table formula by 14.2 ΔE_{00} (Table 7). A staircase, a dual-backing plate, and a phone camera supply what single-point sensors cannot. Within the process-conditioned scope of its coefficients, a single shared model predicts arbitrary stacks. Future work extends the model to volumetric color arrangements [35].

Data and Code Availability

Plate designs, print configurations, photographs, fitted-coefficient reports and the analysis scripts for every table are in the paper artifact directory of the project repository, together with the supplementary material.

References

[1] Vahid Babaei, Kiril Vidimče, Michael Foshey, Alexandre Kaspar, Piotr Didyk, and Wojciech Matusik. Color contoning for 3D printing. *ACM Transactions on Graphics*, 36(4):124:1–124:15, July 2017. doi: 10.1145/3072959.3073605. Proc. SIGGRAPH 2017.

[2] Bambu Lab. Bambu Studio color mixing / mixed filament: Release notes. Product release notes for Bambu Studio, 2026. Accessed 2026-07-17.

[3] Alan Brunton, Can Ates Arikan, and Philipp Urban. Pushing the limits of 3D color printing: Error diffusion with translucent materials. *ACM Transactions on Graphics*, 35(1):4:1–4:13, December 2015. doi: 10.1145/2832905.

[4] Oskar Elek, Denis Sumin, Ran Zhang, Tim Weyrich, Karol Myszkowski, Bernd Bickel, Alexander Wilkie, and Jaroslav Křivánek. Scattering-aware texture reproduction for 3D printing. *ACM Transactions on Graphics*, 36(6):241:1–241:15, November 2017. doi: 10.1145/3130800.3130890. Proc. SIGGRAPH Asia 2017.

[5] Oskar Elek, Ran Zhang, Denis Sumin, Karol Myszkowski, Bernd Bickel, Alexander Wilkie, Jaroslav Křivánek, and Tim Weyrich. Robust and practical measurement of volume transport parameters in solid photo-polymer materials for 3d printing. *Optics Express*, 29(5):7568–7588, 2021. doi: 10.1364/OE.406095.

[6] Shannon M. Hamp, Riley D. Logan, and Joseph A. Shaw. Optical transmittance of 3D printing materials. *Applied Optics*, 60(22):6573, June 2021. doi: 10.1364/AO.427525.

[7] Raed Hlayhel, Sony George, Jon Yngve Hardeberg, and Aditya Suneel Sole. White core thickness effect on 3D-printed color reproduction. *The International Journal of Advanced Manufacturing Technology*, 138:541–553, April 2025. doi: 10.1007/s00170-025-15515-w.

[8] Horn & Rhode LLC. HueForge: Filament painting software for FDM 3d printers. Commercial desktop application, 2023. URL `https://shop.thehueforge.com/`. Accessed 2026-07-17.

[9] International Organization for Standardization and International Commission on Illumination. ISO/CIE 11664-6:2022: Colorimetry — part 6: CIEDE2000 colour-difference formula. International Standard, 2022. URL `https://www.iso.org/standard/82662.html`.

[10] Tomáš Iser, Tobias Rittig, Emilie Nogué, Thomas Klaus Nindel, and Alexander Wilkie. Affordable spectral measurements of translucent materials. *ACM Transactions on Graphics*, 41(6):1–13, 2022. doi: 10.1145/3550454.3555499.

[11] Tomáš Iser, Tobias Rittig, and Alexander Wilkie. Scattering-aware color calibration for 3d printers using a simple calibration target. *ACM Transactions on Graphics*, 44(6):1–14, 2025. doi: 10.1145/3763293. Proc. SIGGRAPH Asia 2025.

[12] justinh rahb. FilamentMixer: Physics-based color mixing for 3D printer filaments, 2025. URL `https://github.com/justinh-rahb/filament-mixer`. Accessed 2026-07-17.

[13] Paul Kubelka. New contributions to the optics of intensely light-scattering materials. Part I. *Journal of the Optical Society of America*, 38(5):448–457, May 1948. doi: 10.1364/JOSA.38.000448.

[14] Paul Kubelka and Franz Munk. Ein beitrag zur optik der farbanstriche. *Zeitschrift für Technische Physik*, 12:593–601, 1931.

[15] Tim Kuipers, Willemijn S. Elkhuizen, Jouke Verlinden, and Eugeni L. Doubrovski. Hatching for 3D prints: Line-based halftoning for dual extrusion fused deposition modeling. *Computers & Graphics*, 74:23–32, August 2018. doi: 10.1016/j.cag.2018.04.006. Special Section on Computational Fabrication.

[16] Eammon Littler, Bo Zhu, and Wojciech Jarosz. Automated filament inking for multi-color fff 3d printing. In *Proceedings of the 35th Annual ACM Symposium on User Interface Software and Technology*, pages 1–13. ACM, 2022. doi: 10.1145/3526113.3545654.

[17] Lumina Studio contributors. Lumina Studio / Lumina-Layers: Physics-based multi-material FDM color system, 2025–2026. URL `https://github.com/lumina-layer-studio/Lumina-Layers`. Accessed 2026-07-17.

[18] Neroued and ChromaPrint3D contributors. ChromaPrint3D: Photo-calibrated multi-color 3D print generator, 2026. URL `https://github.com/neroued/ChromaPrint3D`. Accessed 2026-07-17.

[19] pama1234 and UniverseEmbedded. OpenColor: Four-flux adding-doubling stacked-color FDM pipeline, 2026. URL `https://github.com/UniverseEmbedded/OpenColor`. Accessed 2026-07-26.

[20] Alina Pranovich, Morten Rieger Hannemose, Janus Nørtoft Jensen, Duc Minh Tran, Henrik Aanæs, Sasan Gooran, Daniel Nyström, and Jeppe Revall Frisvad. Digitizing the appearance of 3d printing materials using a spectrophotometer. *Sensors*, 24(21):7025, 2024. doi: 10.3390/s24217025.

[21] Prusa Research. ColorMix for PrusaSlicer and EasyPrint, 2026. URL `https://blog.prusa3d.com/our-new-open-source-colormix-model-in-prusaslicer-and-easyprint_136079/`. Accessed 2026-07-26.

[22] Prusa Research. prusa-fdm-mixer: Calibrated color-prediction model for layer-interleaved FDM, 2026. URL `https://github.com/prusa3d/prusa-fdm-mixer`. Accessed 2026-07-26.

[23] ratdoux. OrcaSlicer-FullSpectrum: Filament color-mixing fork of OrcaSlicer, 2025. URL `https://github.com/ratdoux/OrcaSlicer-FullSpectrum`. Accessed 2026-07-17.

[24] Tim Reiner, Nathan Carr, Radomír Měch, Ondrej Stava, Carsten Dachsbacher, and Gavin Miller. Dual-color mixing for fused deposition modeling printers. *Computer Graphics Forum*, 33(2):479–486, May 2014. doi: 10.1111/cgf.12319. Proc. Eurographics 2014.

[25] Tobias Rittig, Denis Sumin, Vahid Babaei, Piotr Didyk, Alexey Voloboy, Alexander Wilkie, Bernd Bickel, Karol Myszkowski, Tim Weyrich, and Jaroslav Křivánek. Neural acceleration of scattering-aware color 3D printing. *Computer Graphics Forum*, 40(2):205–219, May 2021. doi: 10.1111/cgf.142626. Proc. Eurographics 2021.

[26] J. L. Saunderson. Calculation of the color of pigmented plastics. *Journal of the Optical Society of America*, 32(12):727–736, December 1942. doi: 10.1364/JOSA.32.000727.

[27] Gaurav Sharma, Wencheng Wu, and Edul N. Dalal. The CIEDE2000 color-difference formula: Implementation notes, supplementary test data, and mathematical observations. *Color Research & Application*, 30(1):21–30, 2005. doi: 10.1002/col.20070.

[28] Liang Shi, Vahid Babaei, Changil Kim, Michael Foshey, Yuanming Hu, Pitchaya Sitthi-Amorn, Szymon Rusinkiewicz, and Wojciech Matusik. Deep multispectral painting reproduction via multi-layer, custom-ink printing. *ACM Transactions on Graphics*, 37(6):271:1–271:15, December 2018. doi: 10.1145/3272127.3275057. Proc. SIGGRAPH Asia 2018.

[29] Haichuan Song, Jonàs Martínez, Pierre Bedell, Noémie Vennin, and Sylvain Lefebvre. Colored fused filament fabrication. *ACM Transactions on Graphics*, 38(5):1–11, 2019. doi: 10.1145/3183793.

[30] Rainer Storn and Kenneth Price. Differential evolution – a simple and efficient heuristic for global optimization over continuous spaces. *Journal of Global Optimization*, 11(4):341–359, 1997. doi: 10.1023/A:1008202821328.

[31] Denis Sumin, Tobias Rittig, Vahid Babaei, Thomas Nindel, Alexander Wilkie, Piotr Didyk, Bernd Bickel, Jaroslav Křivánek, Karol Myszkowski, and Tim Weyrich. Geometry-aware scattering compensation for 3D printing. *ACM Transactions on Graphics*, 38(4):111:1–111:14, July 2019. doi: 10.1145/3306346.3322992. Proc. SIGGRAPH 2019.

[32] Théo Phan Van Song, Christine Andraud, and Maria V. Ortiz Segovia. Implementation of the four-flux model for spectral and color prediction of 2.5d prints. *NIP & Digital Fabrication Conference*, 32(1):26–30, 2016. doi: 10.2352/ISSN.2169-4451.2017.32.26.

[33] Pauli Virtanen, Ralf Gommers, Travis E. Oliphant, Matt Haberland, Tyler Reddy, David Cournapeau, Evgeni Burovski, Pearu Peterson, Warren Weckesser, Jonathan Bright, Stéfan J. van der Walt, Matthew Brett, Joshua Wilson, K. Jarrod Millman, Nikolay Mayorov, Andrew R. J. Nelson, Eric Jones, Robert Kern, Eric Larson, C. J. Carey, .Ilhan Polat, Yu Feng, Eric W. Moore, Jake VanderPlas, Denis Laxalde, Josef Perktold, Robert Cimrman, Ian Henriksen, E. A. Quintero, Charles R. Harris, Anne M. Archibald, Antônio H. Ribeiro, Fabian Pedregosa, Paul van Mulbregt, and SciPy 1.0 Contributors. SciPy 1.0: Fundamental algorithms for scientific computing in Python. *Nature Methods*, 17(3):261–272, 2020. doi: 10.1038/s41592-019-0686-2.

[34] vycdev. Kromacut: Open-source HueForge-compatible auto-painting, 2024. URL `https://github.com/vycdev/Kromacut`. Accessed 2026-07-17.

[35] Youjia Wang, Ruixiang Cao, Teng Xu, Yifei Liu, Dong Zhang, Yiwen Wu, and Jingyi Yu. DreamPrinting: Volumetric printing primitives for high-fidelity 3d printing. arXiv preprint arXiv:2503.00887, March 2025.

[36] Günter Wyszecki and W. S. Stiles. *Color Science: Concepts and Methods, Quantitative Data and Formulae.* Wiley, New York, 2nd edition, 1982. ISBN 0-471-02106-7.