Spectral and color prediction for arbitrary halftone patterns: a drop-by-drop, WYSIWYG, “ink on display” print preview

Peter Morovič, Ján Morovič, Xavier Fariña, Pere Gasparin, Michel Encrenaz, Jordi Arnabat, Hewlett Packard Company, Spain

Abstract

Accurately previewing the appearance of a print job can make the difference between producing saleable output and wasting expensive materials and is a challenge to which a host of solutions already exist. However, what the majority of these have in common is that they base their predictions on the inputs to a printing system (e.g., continuous-tone data in ink channels) instead of its outputs (i.e., the halftone data that is then printed) and that they are only valid for a given set of choices already made in the printing system (e.g., color separation and halftoning). Alternatively, attempting to make appearance predictions using general-purpose models such as Kubelka Munk, Yule Nielsen and Neugebauer results in limited performance on systems whose behavior diverges from these models’ assumptions, such as inkjet printing. As a result of such constraints, the resulting previews either work only under limited conditions or fail to predict some artifacts while erroneously predicting others that do not materialize in print. The approach presented here takes advantage of the flexibility of the HANS framework and the insights into spectral correlation to deliver a print preview solution that can be applied to any printing system, that allows for the variation of fundamental imaging choices without the need for re-computing model parameters and that delivers ICC-profile-level accuracy.

Introduction

With digital print production being characterized by large numbers of short-run jobs, the bottleneck in the end-to-end process becomes operator intervention. Web-to-print workflows further increase the cost of manual overrides and heighten the need for accurate print preview before a printed product is ordered – not only for the sake of avoiding a job having to be re-printed but also as a means of providing the print-buying customer with a sense of assurance before they trigger a print order.

A key challenge in any workflow where the cost of an unsaleable print is high (either due to the cost of the materials used or even just due to the delay and frustration that such an event is likely to cause), is to make the print preview closely resemble the final printed output. In the absence of a close match, the decisions made on the basis of the preview are unreliable and its inaccuracies can further aggravate the process of getting to the right print rather than improving it.

The key obstacle to providing an accurate print preview has been the lack of two components: First, an accurate model of print color formation that is applicable to technologies like inkjet, where ink-media interactions are complex, highly non-linear, and not well represented by models that apply more directly to analog printing technologies. Second, a framework for applying scaling to a print preview that reflects the nature of color formation in it, instead of assuming that of displays.

The approach presented here addresses both these challenges by basing the print preview on the result of using the same image processing as used for printing and by taking advantage of recent insights into spectral correlation between neighboring spectral bands. The resulting preview is immediately identifiable as being that of a print (as opposed to only of print-like colors) and displays every feature and artifact of the print it simulates, down to its level of grain and the subtle consequences of the specific halftoning algorithm used. Viewing the on-screen preview results in an experience resembling the viewing of an ink-on-substrate print, including when viewing it at different magnification levels.

Background

Before proceeding to an exposition of the approach introduced here for simulating the appearance of print, two areas will be reviewed first: print spectral and color models and their challenges, and the limitations of current soft-proofing solutions, since it is against their background that a new solution will be presented.

The literature on the prediction of printed spectra and color is extensive, has had good review papers (Wyble and Berns, 1999) and book chapters (Bala, 2002) written about it, and spans the prediction of both ink overprinting and the effect of halftoning on the resulting stimulus. In terms of the former, the Kubelka Munk model (1931) is widely used to predict the properties of multiple layers of ink overlaid at a given location, given information about each constituent ink’s reflectance and opacity. From these, their absorption (K) and scattering (S) coefficients can be computed:

$$\frac{K(\lambda)}{S(\lambda)} = \frac{\left(1-R_\infty(\lambda)\right)^2}{2R_\infty(\lambda)} \tag{1}$$

where R_∞ is the reflectance of an infinitely thick sample and the prediction of scattering and absorption from reflectance is made at a given wavelength λ. Having the K and S coefficients for multiple inks then allows for the prediction of their combined K and S coefficients as follows:

$$K(\lambda) = K_B(\lambda) + \sum_{i=1}^{l} c_i K_i(\lambda) \tag{2}$$

where B indicates the substrate, l the number of ink layers, c_i the concentration and K_i the absorption coefficient of the i-th layer. Computing S(λ) analogously and inverting Equation 1 then predicts the reflectance of the total ink layer stack and, e.g., allows for the prediction of all the ink overprinting combinations of a set of n inks – i.e., its Neugebauer primaries, wavelength by wavelength.

With the ability of predicting ink overprinting, the reflectance of an entire halftone pattern can be computed using the Yule-Nielsen modified Neugebauer model:

$$R(\lambda) = \left(\sum_{j=1}^{p} w_j P_j(\lambda)^{1/n}\right)^n \tag{3}$$

where R(λ) is the reflectance of a halftone pattern neighborhood that is optically integrated as it is being viewed, w_j is the relative area coverage of the j-th Neugebauer Primary – P, and n is the Yule-Nielsen non-linearity that accounts for optical dot gain. Note that the prediction of overall reflectance from component reflectance is again done on a wavelength-by-wavelength basis.

Looking at the entire set of mechanisms for predicting color formation, it can be seen that they have specific ways for predicting specific phenomena: scattering, absorption, optical dot gain and optical integration, that these come with specific assumptions about the properties of the constituent parts (e.g., the homogeneity of ink layers, the nature of their interfaces, independence of layer properties, an infinite substrate thickness, etc.), and that knowledge of certain component properties is a prerequisite (reflectance of infinitely thick layers of ink, scattering of substrate, concentrations of inks, magnitude of optical dot gain, relative areas covered by Neugebauer Primaries).

In practice, this results in significant challenges when attempting an application of the above models and frequent recourse is made to estimating fundamental properties from measurements of printed patterns. While these could be dismissed as purely practical constraints, another way to interpret such an approach is to think of models like the above merely as computational mechanisms (rather than as attempts at modeling phenomena).

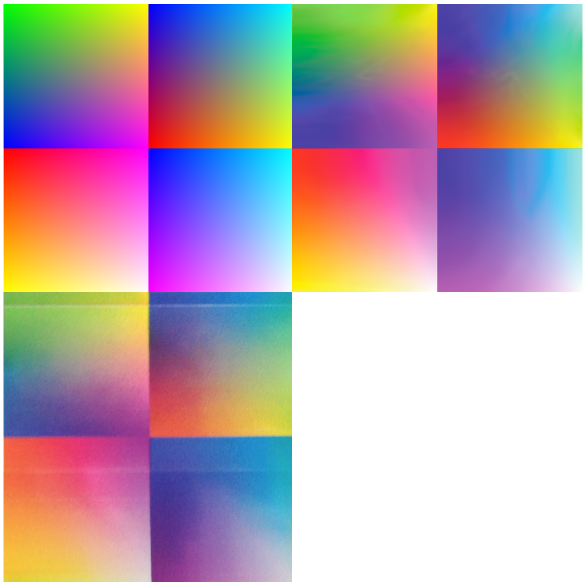

Figure 1: *Input–based print simulation: source sRGB (top left), ICC based soft–proof (top right), photo of print (bottom).*

The second background topic to discuss are the limitations of typical soft-proofing approaches. Here, a common approach is to compute print color simulations (soft-proofs) by building ICC profiles of the display and the printing system that is to be simulated on it, which effectively means treating the printing system as a black box. In other words, the device color inputs (e.g., CMYKs) to the printing system and color measurements of corresponding prints for a set of color patches are taken and a look-up table (LUT) is built from them (indexed in device color space and containing colorimetry at its nodes). With such an ICC profile, source content is transformed to the color gamut of the printing system in a way where, at best, the result is color accurate at the LUT's nodes. In-between these, assumptions of linearity are made and the process is blind both to spatial features such as grain and to the actual nature of transitions between the profile LUT's nodes. Here, the latter include both changes to the mapping from device color space to the printing system's colorant space (i.e., color separation) and the consequences of the specific halftoning used (e.g., when dots start overlapping).

As a consequence, the kind of preview provided using ICC color management based soft-proofing has significant limitations. As can be seen in Figure 1, while the ICC approach does give a sense of overall color changes from sRGB to print, it does not convey a sense of the print's spatial features, makes transitions have a characteristically display-like look and mismatches in terms of artifacts (both missing some from the print and introducing others that do not correspond to print features and instead flow from the ICC profile itself).

An end-to-end print simulation approach

To address the challenges of ICC-based soft-proofing, the approach presented here builds on three key elements:

1. Predicting print appearance from halftone data instead of continuous-tone data input to a printing system after color management.
2. Extending printer modeling by a mechanism that enforces spectral correlation.
3. Printer model based scaling, instead of device color space linear interpolation.

Instead of treating the internal image processing data pipeline of a printing system as a black box and only being concerned about the color that corresponds to specific device color inputs, the starting point here is to follow the exact same steps as used at print time. To generate the preview of a print job, the first step is to pass it through the same color management, color separation, linearization and halftoning as would be used if it were being printed.

Here, the end result is the halftone that could in fact be directly printed, but that will be used for on-screen simulation instead.

The RONT model

Given the halftone that is to be used for simulation, the next challenge is to predict its spectral reflectance or colorimetry in an accurate way. Here, it is possible to apply the approach outlined in the Background section, which consists in using the Kubelka Munk model for predicting Neugebauer Primary colorimetry, from ink reflectance and opacity, and the Yule-Nielsen modified Neugebauer model to predict color integrated over a local neighborhood, from additionally knowing the n exponent representing optical dot gain. In terms of its defining parameters, such a combination can be though of as relying on reflectance (R), opacity (O) and the optical dot gain non-linearity (N) – i.e., RON.

Since ink jet printing systems in particular are not well represented by RON parameters only, and taking advantage of the freedom to think of printer color formation models as computational mechanisms, the next step is to extend them.

Looking at the above KM+YNN model, all of its processing is performed on a wavelength-by-wavelength basis. However, recent work on spectral reflectance analysis has pointed to the presence of strong correlation between adjacent spectral bands (e.g., Singh *et al.*, 2003; Morovic *et al.*, 2014). As a consequence, a useful computational mechanism to add to the KM+YNN models, which relies on RON parameters, is a matrix transformation of the reflectances predicted using those models alone. Conceptually this matrix transformation gives the computational mechanism an opportunity to adjust the reflectance predictions in a given band based on predictions made for their neighbors. This both takes advantage of data from a spectral neighborhood to make predictions in each band and allows for an enforcing of the correlation found in natural spectra that a per-band model may end up breaking.

To compute such a transformation matrix T, the following L2 norm between intermediate reflectance predictions based on RON parameters (R_I) and measured reflectances (R_M) needs to be minimized:

$$\|f(\boldsymbol{R}_I) * \boldsymbol{T} - \boldsymbol{R}_M\| \tag{4}$$

where the two $\boldsymbol{R}$ matrices have m rows (for m samples) and s columns (for s spectral bands) and where *f()* is a function specifying what *t* terms to consider in the error minimization and resulting in the dimensions of $f(\boldsymbol{R}_I)$ being $m \times t$ and those of T being $t \times s$.

For example, *f*() may add a constant term, cross terms and power terms, where in a simple hypothetical case where s=2, the result of *f*() could be the following for a second order case:

$$f(\boldsymbol{R}_i) = \begin{Bmatrix} 1 \\ \boldsymbol{R}(\lambda_0) \\ \boldsymbol{R}(\lambda_1) \\ \boldsymbol{R}(\lambda_0) * \boldsymbol{R}(\lambda_1) \\ \boldsymbol{R}(\lambda_0)^2 \\ \boldsymbol{R}(\lambda_1)^2 \end{Bmatrix} \tag{5}$$

Having computed a T matrix by minimizing Equation 4, its result is applied to the predictions based on RON parameters as follows:

$$\boldsymbol{R}_P = f(\boldsymbol{R}_I) * \boldsymbol{T} \tag{6}$$

The final reflectance prediction therefore starts with KM+YNN, based on reflectance, opacity and optical dot gain nonlinearity (RON), and is then extended by a spectral correlation matrix (T).

Print optical integration-based scaling

The final piece for an accurate print color preview is the ability to simulate the optical integration that takes place when print is viewed at different distances. Starting from a full, print-resolution halftone, where spectral reflectance can be predicted for each pixel using the RONT model outlined above, it is now necessary to add to it a scaling mechanism that faithfully mirrors the optical integration taking place when the print is viewed from different distances.

Here, the key is to carefully express the relationship between halftone pixels that form the basis of the simulation and display pixels for which device RGB content is to be computed (Figure 2).

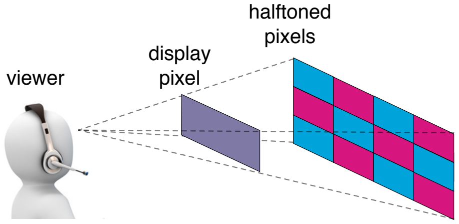

Figure 2: *Correspondence between halftone and display pixels when viewing a print simulation.*

Since each display pixel represents the optical integration of a halftone pixel neighborhood, its color needs to be computed using the full RONT model for that neighborhood, instead of being the mean of the per-halftone-pixel predictions. The first step, therefore, is to determine the neighborhood's Neugebauer Primary (NP) area coverage (NPac) vector, by counting its pixels and establishing how much relative area is covered by each of the system's Neugebauer Primaries. E.g., in the example in Figure 2, 50% relative area is covered using one NP (the cyan primary) and the remaining 50% is covered using another NP (the magenta primary), with there being no overprinting or blank substrate in that neighborhood.

Instead of computing the mean of the cyan and magenta primaries' reflectances, the NPac's reflectance is obtained by applying the YNN model to the KM predictions of its NPs, followed by a transformation by the T matrix. When the viewer then wishes to zoom in our out, then the new simulation colorimetry will be based on first re-computing NPacs corresponding to display pixels and then predicting their reflectances and then colorimeters using RONT.

Test setup

The accuracy of the RONT model introduced here has been tested on the four printing systems shown in Table 1. Note that the details listed there are not necessarily the full capabilities of those systems and can instead just be their subsets.

Label	Z6200-4	PWT-4	L310-4	L310-6
Printer (HP)	Desigjet Z6200	PageWide Technology	Latex 310	Latex 310
Substrate	HP Heavy-weight Coated	HP Heavy-weight Coated	Self-Adhesive Vinyl	HP Photo-realistic Paper
Inks	CMYK aqueous	CMYK aqueous	CMYK latex	CMYKcm latex
Max. drops per ink per pixel	2	2	3	2
Neuge-bauer Primaries	3^4=81	3^4=81	4^4=256	3^6=729
Training samples	756	620	3257	4808
Test samples	459	620	3257	4808

Table 1: *Test printing systems*

Prints made under the four sets of conditions were measured using a 0/45° XRite spectrophotometer, taking measurements at 20 nm intervals between 400 and 700 nm. Colorimetry was then computed under D50 and color difference expressed using the CIE DE2000 color difference equations.

Results

Table 2 and Figure 3 show the colorimetric accuracy of the RONT versus RON models on the two test systems.

System	Model	Median	95th %tile	Maximum
Z6200-4	RON	8.5	23.2	35.4
	RONT	1.5	4.8	8.2
PWT-4	RON	3.2	8.1	15.6
	RONT	0.9	2.2	3.4
L310-4	RON	4.1	11.2	29.9
	RONT	1.3	2.8	6.5
L310-6	RON	1.7	4.9	10.6
	RONT	0.9	2.2	7.7

Table 2: *Colorimetric accuracy of KM+YNN (RON) versus T-matrix extended (RONT) models.*

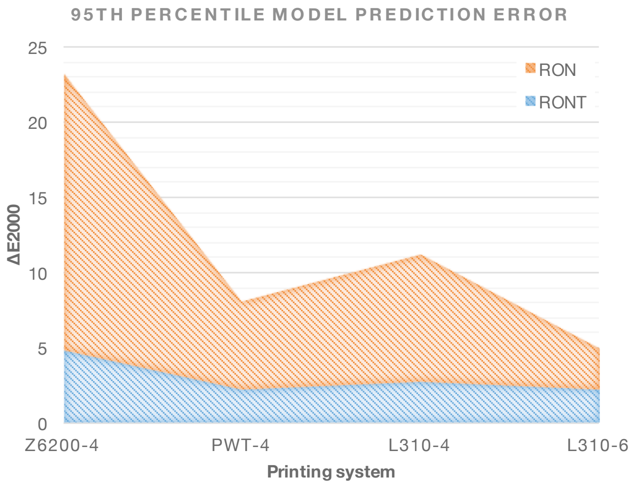

Figure 3: *95th percentile model prediction errors in ΔE2000 for the RONT versus RON models as applied to the four test printing systems.*

As can be seen, the addition of the spectral correlation matrix takes the model from having significant errors on these inkjet systems, using different ink technologies but being poorly represented by the combination of the Kubelka Munk and Yule-Nielsen modified Neugebauer model, to providing predictions that are of a similar level of accuracy to what a specific ICC profile could deliver for them.

While the reduction of prediction errors to a quarter or less of the RON model errors is significant, it also needs to be noted that the result is the ability to predict any halftone pattern printed on a printer-substrate-inks combination, as opposed to only applying for a given set of color separation, linearization and halftoning choices, as is the case for ICC profiles.

Since the purpose of the color model is in simulating complex content instead of individual color patches, Figure 4 shows the result of both printing and photographing a set of halftone data and deriving a display simulation from it too. The sequence of transformations to get to displayable results involved applying the RONT model to get to colorimetry (via reflectance) and subsequently applying an XYZ to sRGB transformation in order to get to display sRGBs that can be used directly on a calibrated, sRGB capable monitor. Since in this case the gamut of colors addressed by the printers did not exceed that of sRGB no additional gamut mapping is needed. In such cases however a gamut-warning akin to that found in software packages such as Adobe Photoshop may be employed, highlighting the areas of the previewed content that exceed the capabilities of the display device, or alternatively applying a gamut mapping algorithm (while potentially loosing some of the information present at the spectral or colorimetric representation).

Looking at Figure 4, note first the close match in overall look and feel. The simulation looks like a print, with its characteristic grain and tonal resolution. The simulation also anticipates closely a variety of artifacts, including contouring (e.g., in the bottom right of the large squares), the hue variations in the gray sphere and grain varying with color (e.g., look at the yellows, versus magentas).

Figure 4: *Simulated print (top) and photo of print (bottom) derived from the same halftone data.*

Conclusions

Successful print simulation can make the difference between wasting time and materials and frustrating (or even losing) customers, and delivering a saleable product efficiently and with good customer experience. While existing techniques of print simulation are able to anticipate high-level color changes to content as it is adapted for printing, they inherit limitations both from the workflow used to generate them and from their being considered like any other display content too early in the process.

Basing simulations on the same data that drives printing systems – i.e., colorant channel halftone data – and ensuring that simulated halftone data is scaled for display in a way that respects its local properties, results in simulations that are both more accurate colorimetrically and that represent print-specific features and limitations faithfully.

In terms of next steps, the framework presented here will be tested on a broader variety of printer-substrate-ink configurations and extended as necessary.

References

Bala R. (2002) Device characterization, in *Digital Color Imaging Handbook*, Sharma G. (ed), CRC Press, 269–384.

Kubelka P. and Munk F. (1931) Ein Beitrag zur Optik der Farbanstriche, *Zeitschrift für technische Physik*, Germany, 12:593–601.

Morovic P., Morovic J., Brill M. H. and Walowit E. (2014) Compression of Reflectance Data Using An Evolved Spectral Correlation Profile, *Color and Imaging Conference*, Boston, MA, 259-264.

Neugebauer H. E. J. (1937) Die theoretischen Grundlagen des Mehrfarbendrucks, *Zeitscrift für wissenschaftliche Photographie*, Germany, 36:73– 89.

Singh B, Freeman W, Brainard D (2003) Exploiting Spatial and Spectral Image Regularities for Color Constancy, *Proc. Workshop Statistical and Computational Theories of Vision*.

Wyble D. R. and Berns R. S. (1999) A Critical Review of Spectral Models Applied to Binary Color Printing, *Color Research and Application*, 24/1:4-19.

Yule J. A. C. and Nielsen W. J. (1951) The penetration of light into paper and its effect on halftone reproductions, *Proc. TAGA*, 3:65-76.