Adaptive Re-Rendering of CMYK Image Data

Johan Lammens[1]*, Ján Morovic*[1]*, Mary Nielsen*[2]*, and Huan Zeng*[3]
Hewlett-Packard Company
[1]*Barcelona, Spain;* [2]*Boise ID, USA;* [3]*Vancouver Washington, USA*

Abstract

The move to digital commercial printing is happening, and is expected to become a major new market for digital printing products and services. The paragon of quality and productivity in the commercial printing world is analog offset printing, which has been perfected over almost a century. While home and office printing is RGB-centric, offset printing is CMYK-centric. Commercial printing workflows, tools, processes, and service providers are geared towards generating and manipulating CMYK content only, with a "generic" analog offset printing press as the (often implicit) reference. For a digital printing device to operate successfully in this context, it must deal with offset CMYK content, do so without requiring major changes in workflows, tools, or processes, and while producing excellent output quality. We describe a CMYK re-rendering method that guarantees high output quality and consistency with offset and other digital printers. At the same time it allows digital devices to make optimal use of their native color gamuts, often larger than offset.

Introduction

The long awaited and predicted move of commercial printing from analog to digital is finally happening. For digital printing to really make inroads into the commercial world, it is not sufficient to provide the basic printing technology. Digital printing equipment needs to be "plugged into" sources of content and electronic design and publishing workflows. A side effect of the move to digital is expected to be a democratization of full color commercial printing and publishing, i.e. a lowering of the access threshold, much as happened two decades ago for monochrome typesetting with the advent of desktop publishing and PostScript. This in turn implies that digital printing equipment will source content from traditional high-end prepress shops or large in-house corporate design departments, as well as directly from the desktops of less sophisticated office users. Even consumer content can be expected, perhaps in the form of EXIF digital camera images. In terms of color spaces this has its implications also: traditionally the commercial printing world has been based on CMYK device color spaces, while the office and consumer worlds have been almost exclusively RGB based. Hence digital presses and their "front ends" will need to be able to deal with color content defined using RGB and CMYK color spaces. While printing from RGB sources in the home and office contexts is fairly standardized (around sRGB mostly), printing from CMYK sources is not. The name derives from the Cyan, Magenta, Yellow, and blacK (or Key) inks that are typically used in commercial printing – but unfortunately there are as many different CMYK color spaces as there are different CMYK inks, printing papers, and printing devices on the market. This is true for a given marking technology, e.g. offset printing, but even more so across technologies like inkjet, dry EP, liquid EP, and offset. While there are sophisticated color management systems available in the high-end market, typically based on the International Color Consortium device profile standard, they are neither easy to use nor widely available to office or home users. We are looking to establish a way of exchanging CMYK color data that is unambiguous and compatible across a large range of digital devices, while at the same time allowing each device to make the most of its own color gamut. To achieve this goal we have defined and implemented adaptive re-rendering algorithms to transform "standard" offset press CMYK into arbitrary device CMYK maintaining relative lightness and hue but mapping chroma in an adaptive and non-linear way to maximize the use of each output device's native gamut, while preserving black.

Overview

The solution consists of two equally necessary and important parts: an implicit colorimetric reference to unambiguously define the exchanged CMYK data, and a way of re-rendering reference color data into each device's gamut. Our current default choice of colorimetric reference is the ANSI/ISO TR-001 data set, which is a colorimetric description of what is otherwise known as SWOP (Specifications for Web Offset Publishing). We have purposely chosen to adopt an existing and well-established colorimetric standard to promote market acceptance, and to prevent confounding the issues by proposing yet another standard. Many professional graphics and even office applications implicitly assume a SWOP-like color space when generating CMYK separations (conversions from RGB to CMYK), hence this is a clear choice for compatibility with existing workflows. Nevertheless, the re-rendering algorithms are flexible enough to be able to work with any reasonable CMYK space as input (SWOP, EuroScale, Toyo, etc.).

But the colorimetric reference is only half of the solution. The other half, and the one where the true novelty resides, is how we interpret that data for each printing device. We re-render the CMYK color data with implicit colorimetric reference into something that preserves compatibility and a common "look and feel", while at the same time allowing each device to optimize usage of its own intrinsic color gamut. In technical terms this first of all implies a tone curve remapping from reference colorimetry to device colorimetry, adapted to the output device's white and black points. In addition we apply a non-linear transformation to reference chroma, while maintaining reference hues. The net result of this transformation is something which comes as close as possible to the seemingly unreachable goal of being "the same as offset, only better", which is one way customers and marketing folks alike tend to express their ideal of CMYK color compatibility. The purpose of all this is to enable *best first print* for the vast majority of digital commercial printing jobs in unspecified CMYK space, generated with common graphic arts applications using their default configurations. For the remaining jobs we adhere to the established sRBG and/or ICC-based color workflows, which ours nicely complements (not substitutes).

Re-Rendering Press Color Appearance

Due to the great variety of digital printing systems, differing significantly in terms of their color gamuts, a re-rendering from press appearance to appearance for a wide range of digital output media needs to be highly adaptive in its behavior. Currently typical digital printing systems have color gamut volumes from around 400,000 to 700,000 cubic CIELAB units in media-relative colorimetric terms and, for comparison, the gamut volumes of press standards like SWOP and Euroscale on coated paper are 450,000 and 480,000 units respectively. As such the re-rendering algorithm needs to be able to deal with significant gamut compression and also address the potential for dramatic gamut expansion.

Furthermore, differences between press and digital printer gamuts are not uniform in all directions in color space. Digital printers can have a more limited lightness range on some substrates than a press standard, while having primaries that are significantly more chromatic. Alternatively, the destination gamut can exceed the source at some hues, while being inside the source at others.

In addition to this heterogeneity of gamut differences, a further challenge for the re-rendering algorithm is to maintain consistency with press appearance across the range of destination gamuts. While for gamut compression there are numerous existing studies that provide a good basis,[1] the gamut expansion case is far less well understood in general. Furthermore the re-rendering of CMYK images in gamuts larger than those of press standards is a special case of gamut extension, as it needs to maintain a strong link to press appearance, while in general gamut expansion is only concerned with optimizing the re-rendering of an image in a larger gamut.

Given the above preamble, the CMYK adaptive re-rendering technique proposed here consists of the following elements. Note that CMYK interpreted in terms of a press standard will be referred to as the source color and CMYK for an output device will be referred to as destination color.

First, source hue is kept unchanged. Doing this is the basis of providing a strong relationship between the appearance of the re-rendered CMYK and press appearance, as hue is the perceptual color attribute that has been found in previous color reproduction studies to tolerate least change.[2,3] Providing the hue match with press appearance also results in consistency of re-rendering across different printing systems. As the primaries of these systems can be of different hues, natively rendering CMYK instead of doing a re-rendering like the one described here, could result in significantly different hues for a given CMYK input.

Second, the source colors are mapped adaptively in terms of their lightness depending on the relationship between the dynamic range of the press standard and that of the destination device. This mapping is such that it provides a smooth, monotonic transition between the extreme cases of heavy compression and significant expansion. Its effect on lightness is shown in Figure 1. The expansion and compression provided by this mapping are both an attempt at enhancement in the former and feature preservation in the latter case.

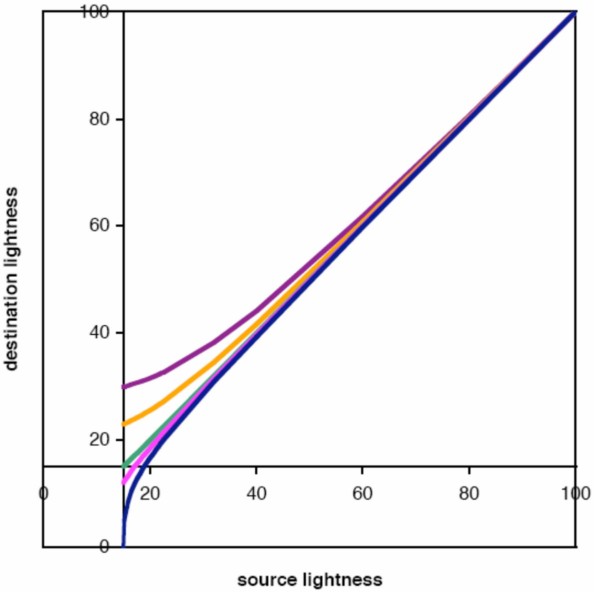

Figure 1. Dynamic range mapping for 15-100 source lightness range and different destination ranges.

Third, once source colors are mapped to have the same dynamic range as the destination, their chromas are mapped using an adaptive method. Here a core region is first defined by taking a scaled-down version of the intersection of the press and destination gamuts. Inside the core region no change is made to chroma, which results in the preservation of the nature of neutral and near-neutral colors as well as of skin tones. Then, if the source color's chroma is outside the core region, it is either expanded or compressed depending on the relationship between the press and destination gamuts in its part of color space.

The complete re-rendering process is illustrated in Figure 2, for an sRGB image processed through different workflows for two different types of digital printers (inkjet and liquid electrophotography).

Figure 2. An sRGB image processed through different workflows for the HP DesignJet 10ps inkjet printer with semi-glossy proofing paper (top row) and the HP-Indigo 3000 LEP printer with glossy paper (bottom row). Left to right: offset CMYK data rendered as SWOP Relative Colorimetric, re-rendered (see text), native (no color correction), and RGB data printed with Perceptual rendering intent. The re-rendered hues are consistent with offset (SWOP). Color saturation is selectively expanded, and closer to direct RGB rendering. The amount of expansion (or compression) depends on the native color gamut of the digital printer being used. Hue differences are especially noticeable in the native case.

Invariability with Respect to Separation Profiles Used

If we want to implement our re-rendering technique as a simple one-button solution we have to use a single invariant interpretation of incoming CMYK data. Such data may have been generated directly in CMYK space, as often happens for text and vector page elements, or have been converted from RGB or other printer-independent spaces, as is normally the case for raster image data. In an ICC based workflow the RGB to CMYK conversion would normally be done using a specific output profile, which we will refer to as the *separation profile*. Our re-rendering algorithms are flexible enough to deal with a rather large range of separation profiles.

In brief, we maintain input hues but adapt lightness range and chroma values from input to digital printer. Interpreting colorimetric characterization data in media-relative terms already reduces some of the apparent differences between various offset press standards.

Figure 3 illustrates the lightness and chroma differences for a number of common offset standards.

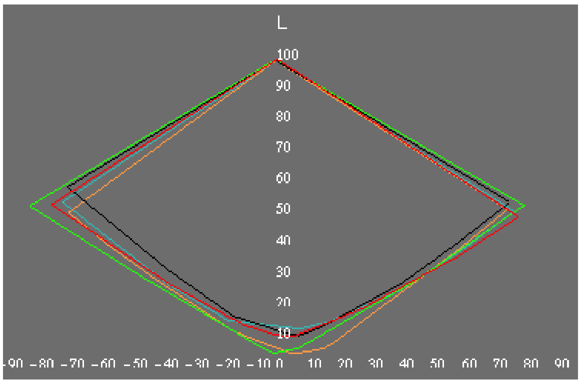

Figure 3. Comparison of different offset standards in the L/a* plane using media-relative colorimetry for SWOP, Gracol, Toyo, EuroScale, and ISOCoatedsb.*

Although there are remaining differences in terms of lightness, they are within the range that our re-rendering algorithms can deal with. Figure 4 illustrates the chroma and hue differences for the same offset standards as Figure 3. As is the case with lightness differences, these chroma and hue differences are dealt with successfully by our re-rendering algorithms.

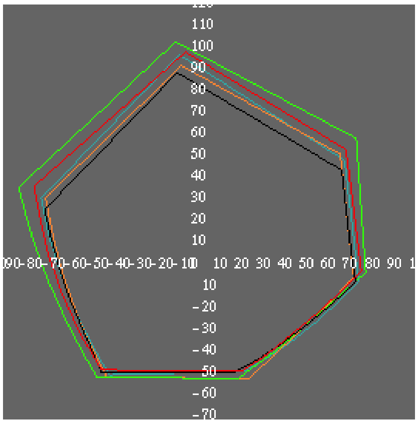

Figure 4. Comparison of different offset standards in the a/b* plane using media-relative colorimetry for SWOP, Gracol, Toyo, EuroScale, and ISOCoatedsb.*

Gray Axis Treatment

Probably the most sensitive aspect of color reproduction in the perceptual sense is gray balance, or the four color reproduction of neutral image areas without any perceived color cast.

When using different separation profiles to turn grayscale images into their four color grayscale CMYK equivalents, there is typically some remaining chroma (non-neturality) after processing the latter through the basic version of our re-rendering algorithms. To evaluate the grayscale cross-rendering properties of different CMYK spaces, the following procedure was followed:

Rendered as	dE*ab Mean			Max			dC* Mean			Max		
	RC	P	all	RC	P	all	RC	P	all	RC	P	all
Euroscale (c)	4.12	4.20	4.16	17.91	8.87	17.91	1.79	1.83	1.81	5.37	5.31	5.37
Toyo	4.81	6.81	5.81	18.94	13.03	18.94	1.45	1.48	1.47	5.55	5.36	5.55
ISOcoatedsb	4.92	6.12	5.52	18.76	11.35	18.76	2.51	2.66	2.59	7.78	7.77	7.78
SWOP	4.04	4.05	4.05	17.51	10.04	17.51	1.77	1.87	1.82	5.11	5.08	5.11
USSheetfed (c)	7.07	3.90	5.48	23.94	10.26	23.94	1.31	1.32	1.31	3.27	3.28	3.28
USSheetfed (uc)	3.99	5.51	4.75	15.30	17.16	17.16	0.89	0.93	0.91	2.30	2.30	2.30
USWeb (uc)	3.99	5.51	4.75	15.30	17.16	17.16	0.89	0.93	0.91	2.30	2.30	2.30
Euroscale (uc)	6.35	8.72	7.54	20.40	22.52	22.52	1.11	1.14	1.12	3.15	3.11	3.15
Gracol	5.05	3.46	4.25	19.71	6.66	19.71	2.06	2.13	2.10	5.82	5.83	5.83

Best
second best
worst

Figure 5. Color and chroma difference means and maxima for cross-rendering CMYK gray scales (see text).

Figure 6. Cross-rendering a CMYK grayscale image separated for five CMYK spaces using re-rendering and GAT (see text). From left to right: separation profiles EuroScale, Gracol, SWOP, Toyo, and US Sheetfed Coated.

1. Generate a 21 step scale along the L* axis (a*=b*=0) from L*=20 to L*=100.
2. Compute CMYK values for the lightness scale using each of nine ICC profiles and using the relative colorimetric (RC) and perceptual (P) rendering intents.
3. Take each of the 18 CMYK scales from step 2 and compute CIE LAB for it using each of the nine CMYK ICC profiles and the relative colorimetric rendering intent.
4. For each of the 162 CIE LAB scales from step 3 compute ΔE^*_{ab} and ΔC^* differences from the original lightness scale from step 1.
5. Report the means and maxima of differences for each of the scales from step 3.

The result of the above procedure are mean and maximum differences for 162 combinations of input and output CMYK space which show what happens if a gray scale computed using each of the CMYK ICC profiles is rendered using simple relative colorimetric intent in each of the CMYK spaces. For each of the nine spaces in which the gray scales are rendered the means of the means and the maxima of the maxima of color and chroma differences are shown in Figure 5. The dE differences are due in part to the dynamic range differences between the spaces, so we are more interested in the chroma (non-neutrality) differences. Some spaces are clearly better than others for grayscale cross-rendering, but in general the resulting non-neutrality is not acceptable.

To improve these results, we have implemented a *Gray Axis Treatment* (GAT) algorithm which acts as a pre-processing step before the actual re-rendering takes place. It applies a transformation to CIE LAB values where chroma values below a first threshold are collapsed to zero, chroma values in the next interval are heavily compressed and a gradual transition is then made to identity. With GAT, the neutrality of re-rendered prints is improved considerably for a wide variety of separation profiles. In images with very near-neutral colors some desaturation occurs as a side effect, but nevertheless transitions from neutrals to near-neutrals are smooth, and the result is deemed acceptable. The process is illustrated in Figure 6. Neutrals are rendered as neutrals for all separation profiles tested. There are differences in tone reproduction, but those are intended.

Results and Conclusion

We have implemented the solution described for inkjet printers, dry EP ("color laser printers"), and liquid EP printers. An important constraint on the implementation is that it must be black preserving, i.e. the incoming black ink image (or separation) should be preserved as is. Black ink (K) is used extensively in graphic arts applications for more than just making dark colors: black text and line art is printed as K only to ensure maximal sharpness and to avoid color halos, the black separation is used to enhance contrast in images and/or reduce total ink usage, to ensure a stable gray balance, to create "rich black" and drop shadows, and so on. In practice the solution is embodied as a 4 dimensional lookup table combined with standard interpolation techniques, which can be embodied in an ICC device link profile for use in RIPs. Preliminary internal and external image quality testing shows good results, and a general preference for re-rendered results versus the ones obtained with traditional means (media-relative colorimetric or "native" rendering).

None of the existing or possible alternative solutions would seem to meet the requirements. To summarize, the solution needs to enable *best first print* in the vast majority of cases, be compatible with legacy workflows as well as sRGB and newer high end ICC workflows, maintain control over black, and provide cross-device "look and feel" compatibility while optimizing the usage of each device's "native" color gamut. If offset CMYK data is rendered colorimetrically we end up with a color proof which limits the color gamut to that of the press. Offset color gamuts are relatively small compared to many modern digital printers. If each device is driven directly in device color space (without any intervening transformations), we may end up with very large differences from one device to another because of the differences in marking technology, colorants, and fundamental ink/media interactions. If we were to define a new colorimetric reference and promote it as a digital color standard, we would add to the already considerable confusion in the color standards arena, complicate the solution, break compatibility with existing workflows, and under-utilize some devices' capabilities.

References

1. Morovic, J. (2002) Gamut Mapping. In *Digital Color Imaging Handbook*, Sharma G. (ed.), CRC Press.
2. Katoh N. and Ito M. (1996) Gamut Mapping for Computer Generated Images (II), *Proceedings of 4th IS&T/SID Color Imaging Conference*, 126–129
3. Ebner F. and Fairchild M. D. (1997) Gamut Mapping from Below: Finding the Minimum Perceptual Distances for Colors Outside the Gamut Volume, *Color Research and Application*, 22:402–413

A Note on the Figures

Due to the nature of the work described here, it is difficult to do justice to the color and grayscale illustrations in a single-printer hardcopy of any kind. The illustrations are therefore best viewed in electronic form on a reasonably calibrated monitor. Please contact the authors for electronic copies in PDF format.

Biography

Johan M. Lammens was born in 1962 in Belgium, where he studied linguistics and computational linguistics at the University of Antwerp (1980-1985). At the Institute for Perception Research in Eindhoven and the University of Utrecht (the Netherlands), he worked on speech synthesis and text-to-speech conversion (1985-1988). At the State University of New York at Buffalo (USA), he obtained his MS and PhD degrees in computer science (1988-1994). His dissertation research was concerned with a computational model of color perception and color naming. At the S. Anna institute of the University of Pisa (Italy) he held a post-doctorate position, working on high-speed stereo foveal vision and motion control (1994-1995). Since 1995 he has been with the R&D department of Hewlett-Packard in Barcelona (Spain), currently in the position of Senior Color Scientist working on color imaging for digital publishing systems.