EUROGRAPHICS 2012 / P. Cignoni, T. Ertl
(Guest Editors)

Volume 31 (2012), Number 2

Black is Green: Adaptive Color Transformation For Reduced Ink Usage

L. Shapira[1] and B. Oicherman[2]
[1]Hewlett Packard Labs, Israel [2]Independent consultant

Figure 1: *In preparing an image for print (left), we reduce ink usage by detecting areas which are active (middle, brighter is more active) i.e. contain detailed textures, edges and structure. In these areas we are able to use less cyan, magenta and yellow ink while increasing use of black ink. We are able to reduce ink usage (right) by up to 25% (10% in this example) without degradation in printed image quality.*

Abstract

A vast majority of color transformations applied to an image in the digital press industry are static and pre-calculated. In order to achieve the best quality on a wide variety of different images, these transformations tend to be highly conservative with respect to the use of black ink. This results in excessive use of inks, which has a negative economic and environmental impact. We present a method for dynamic computation of color transformation based on image content, with the aim to reduce ink usage. We analyze the image, and predict areas in which quality artifacts that may result from such a reduction will be masked by the image content. These areas include detailed textures, noisy areas and structure. We then replace the image CMYK values by a new combination with increased black. Our algorithm ensures negligible color shifts in the resulting image, and no visible reduction in quality. We achieve an average of over 10% ink savings.

Categories and Subject Descriptors (according to ACM CCS): B.4.2 [Input/Output And Data Communications]: Input/Output Devices—Printers I.4.3 [Image Processing and Computer Vision]: Enhancement—Filtering

1. Introduction

Printing images and documents on paper is a highly prevalent practice, despite the development of image reproduction alternatives such as tablets and digital picture frames. A printer reproduces images by applying dots of ink arranged in halftone patterns on paper. The cost of inks in a substantial part of the total printing cost, and there is a constant drive on the side of printer's manufacturers and users to reduce it. Another aspect of ink usage is environmental: removing ink from paper is a substantial part in the recycling process.

A vast majority of printers use four inks to reproduce color: cyan, magenta, yellow and black (CMYK). The CMY inks reproduce the color shades, whereas the black ink serves a dual purpose: enhancing the image quality and saving inks. For a given color, different combinations of CMY and K may be used. However, in practice, the amount of black added to the CMY mixture is determined beforehand.

The realization that drives this research is that if the color transformation is computed according to the content of the image, black ink can be introduced more aggressively, and a significant amount of ink can be saved. Our contribution is an algorithm which predicts the effect of increasing use of black in different areas of the image. We use that knowledge to compute an optimized local transformation which minimizes the amount of ink used, without degrading image quality (Figure 1).

2. Related Work

In [FSPG97], Ferwerda et al develop a computational model of visual masking, based on psychophysical data. The model predicts how one visual pattern affects the detectability of another. They demonstrate the effectiveness by predicting how masking provided by a texture map affects the visibility of shading artifacts caused by polygonal tessellation. In [WPG02] a similar measure, based on the JPEG discrete cosine transform, is used to process textures and accelerate rendering algorithms such as irradiance caching and path tracing for global illumination.

There have been several works attempting to characterize image quality artifacts in half toned prints. ISO standard 13660 [ISO01] defines an objective means of communications about basic image quality parameters and measurement methods. ISO Standard 19751 [RDN*06] presents a standard for evaluation of the perceptual image quality of color printers. [Miz00] develops a metric for measuring graininess, the subjective perception of noise in an image, particularly for clustered dot color images.

There are several commercial color-pipeline products, which claim to reduce ink without visibly changing the printed product [Fus11, Puz11, One11]. However we could not find white-papers, publications or patents detailing these products and therefore could not compare our algorithm to theirs.

In the last decade, the importance of energy saving and protecting the environment has become evident. In [CWM09], energy aware color sets are developed, composed of distinguishable iso-lightness colors. These colors sets save up to 40% power consumption, when used for 2D and 3D visualization.

We are interested in reducing the amount of ink used in printed images, which will lead to savings in ink production and an easier recycling process. However, saving ink cannot come at the expense of image quality. Therefore, we employ the concept of visual masking to estimate the visibility of image quality artifacts (e.g. graininess).

3. CMYK Process and Printed Image Quality

As noted, in the CMYK printing process [Kan97], the CMY inks reproduce the color shades, whereas the black ink serves a dual purpose: enhancing the image quality and saving inks. When applied in proper quantities, the combinations of CMY inks reproduce faithfully gray shades and also black. However, the CMY-only images suffer from two problems: excessive use of inks and reduced dynamic range. The excess of ink results from the fact that many CMY combinations contain certain amount of "gray component": a combination of inks which, if printed alone, will produce a neutral gray. If this CMY component would be replaced with an achromatic ink (black or gray) the total amount of ink in the image would be greatly reduced without any change in color. In addition, the darkest possible color achievable with the combination of CMY inks in most of the printing processes is still not dark enough for good contrast image reproduction. Addition of black ink in the dark areas of the image helps achieve darker blacks and richer details, thus enhancing the dynamic range of the image.

The amount of black ink that is added to the CMY mixture is computed by the process known as gray Component Replacement (GCR). In this process, in every pixel the amounts of CMY inks that make up the gray component are identified. Consequently, part of the CMY inks in the gray component is replaced by black ink. Just what part of the gray component is to be replaced by black is the matter of trade-off between the image quality and ink saving. The more CMY inks are replaced by black the lower is the printing cost. On the other hand, excessive usage of black creates visual artifacts.

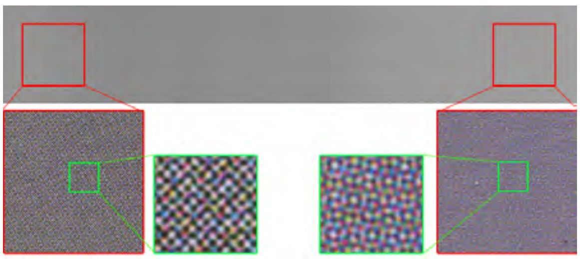

Figure 2: *In order to demonstrate graininess, we print a smooth gray rectangle. However, from right to left, we change the gray component balance from CMY only with no K, to maximum K use. Shown here is a 1200 DPI scan, demonstrating distinct differences in the half-toning texture. The GCR maximized patch on the left shows more graininess.*

Due to the high contrast between the ink and the paper, the black dots are more visible then the dots of other inks. Thus, excessive addition of black ink produces quality artifacts known as graininess. Graininess describes the subjective perception of noise in an image, and is mostly comprised of high frequency a-periodic fluctuations, which cause visual non-uniformity on the printed image (Figure 2). The visibility of graininess in an image varies, depending on the properties of the image area: in smooth image areas such as sky or skin the graininess is most visible; "active" areas, which contain texture or structure such as vegetation or stone visually mask the graininess and render it invisible. In addition, the visibility of graininess depends on the frequency of the halftone screen: higher frequency screens will produce smaller ink dots and therefore will exhibit less graininess.

The standard method of encoding the transformation from device independent color space (e.g. CIELAB) to device color space (e.g. CMYK) is the ICC (International Color Consortium) color profile [ICC06]. An ICC profile is static, computed once for every printing condition, and stores the color transformation in look-up tables or in matrix transformation parameters. Since a given profile should satisfy the quality requirements of any possible image, the approach to the GCR process taken in the profile creation process is highly conservative. In order to avoid any quality artifacts only the minimum of black ink is used in areas of high and medium lightness.

4. Image Activity Map

In order to apply localized GCR to an image, without degrading image quality, we must be able to estimate whether an area will appear grainy for a given amount of black ink. We can then, locally, set the maximum GCR level. We define a new function, **image activity**, that estimates graininess visibility in a pixel neighborhood. The algorithm for creating the activity map for an image is as follows (example results for each step can be seen in Figure 3):

1. Scale the image to a standard resolution.
2. Calculate a weighted entropy filter.
3. Calculate a skin probability map.
4. Combine mapping of entropy and skin probability into activity map.
5. Apply joint-bilateral upsampling, to upscake the activity map and force edges to adhere to original image edges.

We scale the input image to a standard resolution, such that details in different images will be of the same size. This resolution is dependent on characteristics of the printing engine. In our case we scale all images to 240 DPI, the “natural” resolution for the printers used.

Weighted entropy. Image entropy is a measure of the uncertainty associated with the image, its value being higher the closer the image is to white noise. A localized entropy filter is calculated on a gray level image. For each pixel p we build a gray level histogram H_p of 256 bins ($h_0..h_255$) for the $9x9$ neighborhood N_p around the pixel. We normalize the histogram bins (such that the total area is 1) and define the entropy to be

$$E = -\sum_{i\in[0,255]} h_i \cdot log(h_i)) \tag{1}$$

Entropy, however, is insensitive to the magnitude of difference between gray levels (Figure 3b), i.e areas with high variance might produce the same entropy result as areas with small variance. We experimented with using gray level co-occurrence matrices (GLCM) but found them computationally expensive without substantial benefits. We define a weighted entropy variant $\hat{E}$

$$\hat{E} = -\sum_{i\in[0,255]} w(i) \cdot h_i \cdot log(h_i)), \tag{2}$$

where $w = 1 - G(mean(N_p), 8)$, G is a Gaussian around $mean(N_p)$ with variance 8. $\hat{E}$ cancels out small variations around the neighborhood mean so that only significant gray level variations affect the entropy measure (Figure 3c).

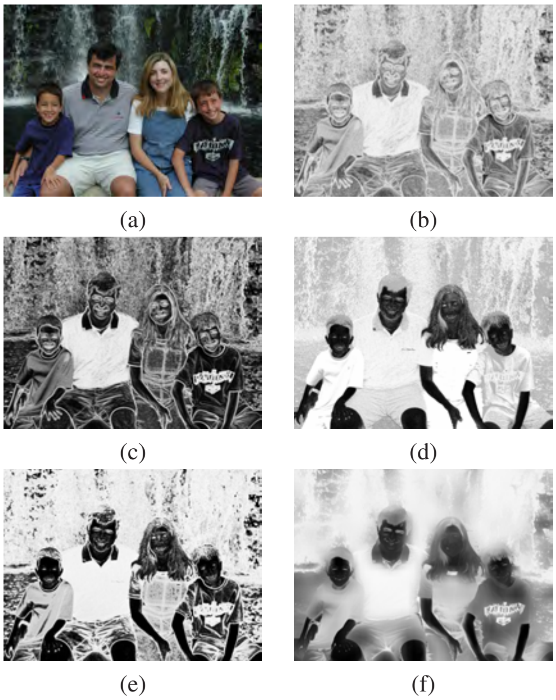

Figure 3: *Calculating activity: (a) Original image (b) Entropy filter (c) Weighted entropy filter (d) Skin tone detection (e) Combining weighted entropy with skinmap (f) Final activity map after applying joint-bilateral upsampling.*

Skin tones. Humans are sensitive to artifacts on skin. Therefore we wish to dampen activity in skin toned areas. We use a skin probability estimator based on [NBG*09]. Skin probability for a pixel is calculated using its chroma and hue:

$$P(skin,c,h) = e^{\frac{1}{3}\frac{-(c-\mu_c)^2}{2\sigma_c^2}+\frac{2}{3}\frac{-(h-\mu_h)^2}{2\sigma_h^2}} \tag{3}$$

where (c,h) are the chroma and hue coordinates of the pixel, and $\mu_c = 33, \mu_h = 53, \sigma_c = 25, \sigma_h = 25$. The resulting anisometric Gaussian is visualized in figure 4.

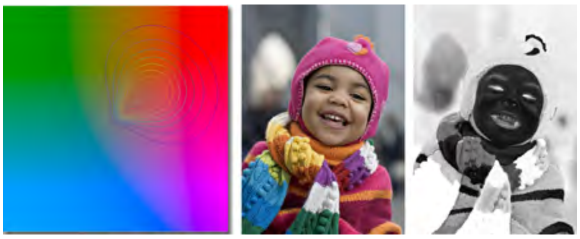

Figure 4: *(left) Skin tone probability is estimated using a modified Gaussian over hue and chroma. (right) Example skin probability map (darker is higher skin probability).*

Combined Activity Map. We developed a test which consists of printing a series of modulated gray-level patches (Figure 5). Each patch shows a horizontal modulated signal, defined by frequency, luminance, amplitude and GCR level. We printed a complete set of patches over the following parameter grid (overall 980 patches):

- $Frequency = [10,20,30,40,50,60,100]$
- $Luminance = [32,64,96,128,160,192,224]$
- $Amplitude = [4,8,16,32]$
- $GCR = [0,0.2,0.5,0.75,1.0]$

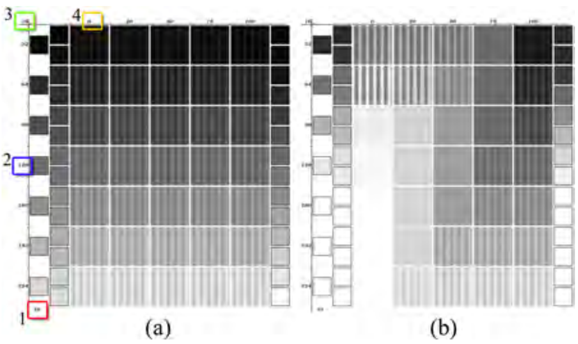

Figure 5: *(a) Series of test patches varying over frequency <1>, luminance <2>, amplitude <3> and GCR <4> intended for building a graininess measure. (b) The difference in GCR levels can only be seen in the K channel of the test patches.*

where frequency is in cycles per inch (CPI) and GCR levels signify the amount of K used (1.0 is maximal). For each triplet of values $(Freq, Luminance, Amplitude)$, we observed the minimal GCR level at which graininess becomes visible. We fit a curve mapping weighted entropy values into activity levels, according to these observations (Figure 6a). The curve is defined as

$$a = \frac{1.061 \cdot \hat{E} - 0.05}{w + 0.98} \tag{4}$$

where a is the activity level, calculated from a weighted entropy value $\hat{E}$. In the future we intend to replace this curve with data from a psychophysical experiment.

We then map skin probability levels into an activity level multiplier so that in areas with high probability we completely dampen the activity (Figure 6b). This mapping is described by a sigmoid curve as follows

$$s_m = 0.45 + \frac{-0.5972 \cdot \hat{s}}{(1+\hat{s}^2)^{0.43}} \tag{5}$$

where s_m is the multiplier dependent on the normalized skin probability levels $\hat{s}$ (with mean 0.71 and standard deviation 0.19), restricted to the range $[0,1]$. The combined entropy and skin values can be seen in figure 3e.

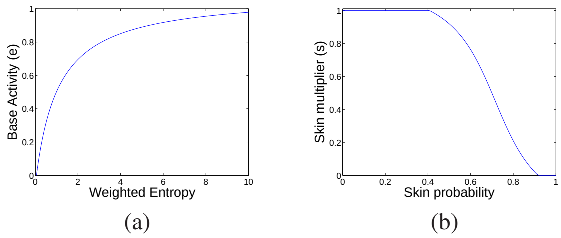

Figure 6: *(a) Weighted entropy values are mapped to activity levels (b) Skin probability values are mapped to an activity multiplier.*

Joint bilateral Upsampling The resulting activity measure (Figure 3e) has blurry edges, due to the neighborhood-based action of the weighted entropy filter. We would like the activity measure to have distinct edges, corresponding to the original image edges. Moreover, at this stage, we want to upsample the activity map to the original image resolution. Therefore we apply a joint-bilateral upsampling [KCLU07], using the edges of the original image. The result (Figure 3f) is sharp and scaled back to the original image size.

Additional examples of activity map can be seen in figure 7.

5. Dynamic Color Transformation

Given an input image I, an activity map A, and a target ICC profile (created for the output printer), we aim to produce a

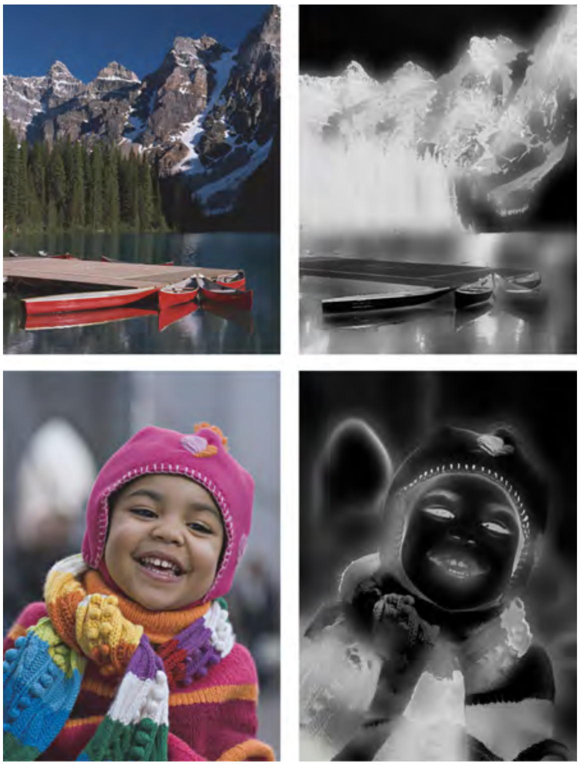

Figure 7: *Additional examples of activity map generation. Active areas appear as white, while inactive areas appear as black.*

ready to print image I_{out}, in CMYK device space. We maximize the usage of black ink in areas with high activity, such that visual quality is preserved, and total ink usage is significantly reduced. Our algorithm is composed of the following steps:

5.1. Conversion and Gamut Mapping

Given an image I, we convert it to CMYK using the device ICC profile. We save the K channel as $K_{profile}$. We then convert the CMYK values back to CIELAB using the inverse transformation. The resulting LAB values I_{lab} are within the printer gamut.

5.2. Calculating Target K

Given an activity map A, the default $K_{profile}$ map, and the input image, we generate a new K_{target} map. We note that light or highly saturated pixels have a small gray component. Therefore, in these pixels, the k levels cannot be overly increased. We define two distance functions

- d_g - the distance between the color of the pixel and the gamut boundary (calculated using Segment-Maxima [Mor08]).
- d_l - the distance between pixel luminance and the profile white point.

  We then define the k_{target} value for each pixel

$$k_{target} = (1-a)\cdot k_{default} + a \cdot \min(M_g(d_g), M_l(d_l)) \quad (6)$$

where a is the activity level of a pixel. M_l maps pixel luminance values using the following sigmoid function

$$M_l(\hat{d}_l) = 0.51 + \frac{-0.97\cdot \hat{d}_l}{(1+\hat{d}_l^{\,2})^{0.93}} \quad (7)$$

where $\hat{d}_l$ is a normalized value (with mean 0.57 and standard deviation 0.34). We restrict the values of M_l to the range $[0,1]$. M_g maps d_g using a similar function

$$M_g(\hat{d}_g) = 0.48 + \frac{1.03\cdot \hat{d}_g}{(1+\hat{d}_g^{\,2})^{0.95}} \quad (8)$$

where $\hat{d}_g$ is a normalized value (with mean 12.2 and standard deviation 11.37). We restrict the values of M_g to the range $[0,1]$. These mappings are illustrated in figure 8.

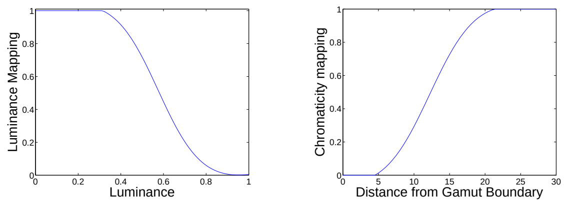

(a) (b)

Figure 8: *Maximum k value for a pixel is determined using pixel luminance (a) and distance from gamut boundary (b).*

5.3. Generating a Reduced Ink Image

Given (l,a,b,k_{target}) values, we seek to find (c,m,y,k) coordinates, which minimize

$$||T(c,m,y,k)-(l,a,b)|| \quad (9)$$
$$||k_{target}-k|| \quad (10)$$

where T is the ICC profile transformation. In practice, we rely on the printer ICC profile in order to build this function. As noted before, the profile contains forward and backward transformations of type

$$T_{fwd} : LAB \rightarrow CMYK \quad (11)$$
$$T_{back} : CMYK \rightarrow LAB \quad (12)$$

These transformations are based on colorimetric measurements, however these measurements are usually not directly available. We wish to build a transformation $T_{inksave} : LABK \rightarrow CMYK$. As a pre-processing stage (once per profile), we build a data structure as follows:

- Sample CMYK space uniformly, using 40^4 samples (40 samples along each dimension).
- Transform CMYK values into LAB using T_{back}.
- Store the transformed *LAB* values in a data structure *D* which supports nearest-neighbors (NN) queries. Each *LAB* entry is linked to its originating *CMYK* coordinates.

Given (l, a, b, k_{target}), we perform a NN search in D, within a small color shift tolerance ($0.5\Delta e$). The result is a group of colorimetrically nearly identical (l, a, b) values. All *LAB* values in this group satisfy the first constraint. We select a value, such that its originating (c, m, y, k) coordinates minimize the second constrinat $||k - k_{target}||$. If a near enough k value does not exist (beyond a certain tolerance), we interpolate between the two nearest values.

We implemented this data structure using FLANN [ML09], a highly efficient data structure supporting approximate nearest-neighbor queries. The pre-processing stage of the algorithm takes about 5 seconds on an Intel Core i7 2.67 GHZ machine, with 8GB RAM. Each image is processed in blocks, such that we perform a query for 2000 pixels at once. The total running time for a 5 mega-pixel image is about 3 minutes.

6. Psychophysical Study

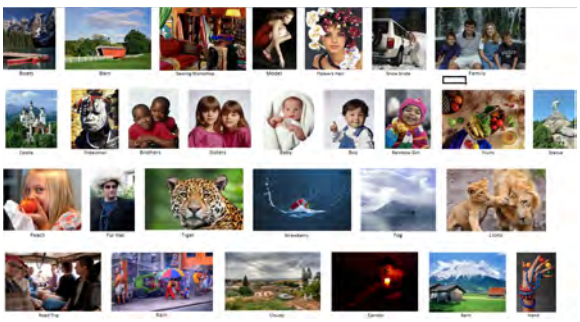

Figure 9: *Thumbnails of the images appearing in the user study*

In order to validate the results of our algorithm we conducted a forced-choice pair comparison psychophysical study. In each pair, one image was prepared for print using a static ICC profile color transformation. The other was prepared using the adaptive ink saving algorithm described in this paper. Our hypothesis is that overall, observers will not favor one version of the image over the other, since we do not change the image quality. Therefore, when an observer is forced to choose we expect her to choose a preferred image at random. Hence, over a large set of observations we expect a 50/50 ratio of choice between image versions. This paradigm was verified by performing a forced-choice experiment with identical image pairs.

We prepared a set of 28 image pairs, covering a range of scene types (Figure 9). Each image was printed on a single sheet, cut out and mounted in gray matboards identified with a generic code. Samples were presented to each observer in random order, and were shuffled before the next observer evaluated them. This procedure distributes any start-up, sequential, or fatigue effects over different samples. The evaluation was done in a room operated under industry standardized conditions for visual testing - namely neutral gray surround, D65 diffuse lighting, and 2000 lux illuminance at the sample viewing surface.

Each participant was required to evaluate all pairs of images, and select an image *she would like to take home with her*. We conducted the test in two groups, the first consisted of 19 professional observers, with previous experience of evaluating the quality of printed images. The second group consisted of 20 naive observers i.e. with no such experience.

The results for the two groups of observers were not significantly different, therefore the two sets of data were combined. The results show slight preference for images prepared with the standard ICC profile: they were preferred 56% of times, with the 95% confidence interval at 6%.

When viewing the results per image (Figure 10), it can be seen that for some images, the ICC version was preferred almost 20% more frequently then the ink-saving one. We attempted to analyse the causes for the apparent failure of our approach in these instances by examining the "problematic" prints, and also in conversation with some of the observers after the experiment. To our surprise, the pairs of images seem identical, and not ourselves nor the observers could point to any defect, or indeed to any difference between the images that could cause such a result. We conducted an additional experiment with a different set of images and a smaller pool of observers, and received very similar results. At the time of this writing, we cannot point to any possible cause for the apparent experimental anomaly, so we consider understanding this phenomenon an ongoing task.

7. Results and Discussion

Our algorithm saves ink by increasing use of black ink over cyan, magenta and yellow, in areas of high activity (Figure 11). The printed result is of the same quality as an image printed with a standard static ICC profile. In figure 12 we

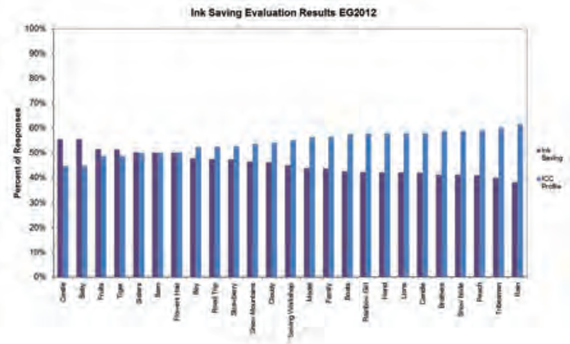

Figure 10: *Psychophysical experiment results by image*

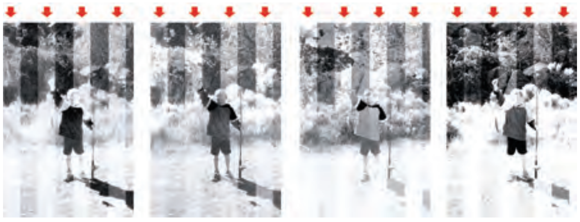

Figure 11: *After applying the ink saving algorithm, the CMY channels of the image become brighter (less ink is used), while the K channel (black) becomes darker, in active areas. The channels are shown in stripes showing ink saving results (marked with arrows) and default profile conversion. In this image our algorithm reduces ink usage by 17%.*

show a distribution of ink saving levels, tested on 100 images, of different catagories. Average ink savings was 10% while most natural images (non-studio) showed savings of higher than 15%.

Although we produce images without compromising quality, some changes in half-toning patterns may be noticeable between the ICC profile print and our ink saving print. Moreover, the base ICC profile which we use may cause artifacts which are replicated in our results.

We have found that in a high activity area with a large range of k values, our algorithm may compress this range, and produce posterization in the k channel (Figure 13). This effect can be noticeable in print, giving a flattened feeling to the area, but occurs rarely.

Finally, a psychophysical study we conducted shows that when forced to choose, users show a slight but consistent preference for the ICC profile version of the image. Image examination and conversation with observers however, did not reveal the cause of the preference: the images indeed look identical. We are investigating the source of this experimental anomaly.

8. Future Work

Our algorithm is highly parallelizable and we are working on a faster, multi-threaded implementation. Additionally, we are exploring implementing it on a GPU.

Currently, given an activity map, the transformation is done per pixel, regardless of its neighbors. In future work we intend to create a spatially aware algorithm, which can help solve the k posterization problem among other advantages.

We intend to increase the amount of ink savings, by exploiting several visual phenomena, such as allowing some minor color shifts in noisy areas. We also intend to extend the work to graphics, which are characterized by smooth and gradiently changing areas of color. We believe that a dynamic color transformation may reduce artifacts resulting from using a static profile. For example, if an image contains a smooth gradient and k is introduced into the ICC color transformation within this range, graininess will be highly evident. We on the other hand, could detect the gradient area and change the black ink starting level to cover the whole gradient range (saving ink and improving quality).

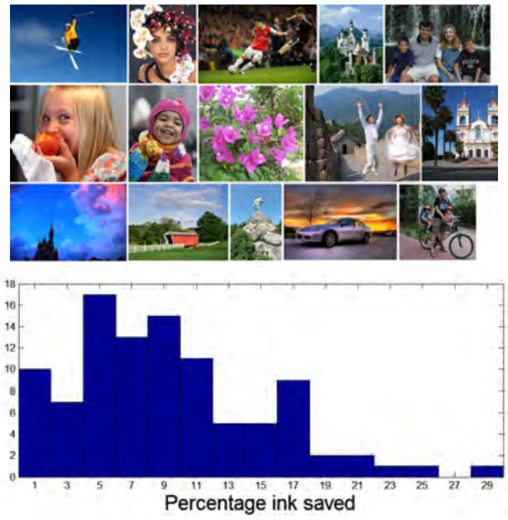

Figure 12: *We have tested our ink saving algorithm over a diverse set of 100 images, a sampling of which (ordered by ink save percentage from left to right) can be seen on top. The histogram of ink saving percentage is shown on the bottom.*

References

[CWM09] Chuang J., Weiskopf D., Möller T.: Energy aware color sets. *Comput. Graph. Forum 28*, 2 (2009), 203–211. 2

[FSPG97] Ferwerda J. A., Shirley P., Pattanaik S. N., Greenberg D. P.: A model of visual masking for computer

Figure 13: *(a) Image processed using static ICC profile K channel with enlarged patch (b) Image processed using dynamic ink saving algorithm K channel, enlarged patch shows K posterization effect (bottom row shows k-channel only).*

graphics. In *Proceedings of the 24th annual conference on Computer graphics and interactive techniques* (1997), SIGGRAPH '97, ACM Press/Addison-Wesley Publishing Co., pp. 143–152. 2

[Fus11] FUSION SYSTEMS INTERNATIONAL: Alwan cmyk optimizer eco, 2011. `http://bit.ly/pos7gW`. 2

[ICC06] ICC (Ed.): *Specification ICC.1:2004-10*. International Color Consortium, 2006. `http://www.color.org/icc_specs2.xalter`. 3

[ISO01] ISO STANDARD: *13660: Measurement of image quality attributes for hardcopy output – Binary monochrome text and graphic images*. ISO (International Organization for Standardization), 2001. 2

[Kan97] KANG H.: *Color Technology for Electronic Imaging Devices*. SPIE Publications, 1997. 2

[KCLU07] KOPF J., COHEN M. F., LISCHINSKI D., UYTTENDAELE M.: Joint bilateral upsampling. *ACM Trans. Graph. 26* (July 2007). 4

[Miz00] MIZES H.: Graininess of color halftones. In *Proceedings of the 16th NIP International Conference on Digital Printing Technologies* (November 2000), IS&T, pp. 396–399. 2

[ML09] MUJA M., LOWE D. G.: Fast approximate nearest neighbors with automatic algorithm configuration. In *International Conference on Computer Vision Theory and Application VISSAPP'09)* (2009), INSTICC Press, pp. 331–340. 6

[Mor08] MOROVIC J.: *Color Gamut Mapping*. Wiley, 2008. 5

[NBG*09] NACHLIELI H., BERGMAN R., GRIEG D., STAELIN C., OICHERMAN B., RUCKENSTEIN G., SHAKED D.: *Skin-sensitive Automatic Color Correction*. Tech. rep., HP Labs, 2009. `http://bit.ly/oJYuKy`. 3

[One11] ONEVISION: Inksave, 2011. `http://bit.ly/pl2vH0`. 2

[Puz11] PUZZLEFLOW: Inksaver, 2011. `http://bit.ly/rctAup`. 2

[RDN*06] RASMUSSEN D. R., DONOHUE K. D., NG Y. S., KRESS W. C., GAYKEMA F., ZOLTNER S.: Iso 19751: Macro-uniformity. *Image Quality and System Performance III 6059* (January 2006), 188–199. 2

[WPG02] WALTER B., PATTANAIK S. N., GREENBERG D. P.: Using perceptual texture masking for efficient image synthesis. *Comput. Graph. Forum 21*, 3 (2002). 2