Spectral Vector Error Diffusion - Promising Road or Dead End?

Ole Norberg[1] and Daniel Nyström[2]
[1] Voxvil AB, Strandgatan 12E, SE-891 33 Örnsköldsvik, Sweden
[2] Linköping University, Dept. of Science and Technology (ITN), SE-601 74 Norrköping, Sweden

Abstract

The interest for spectral color reproduction has increased with the growing field of multispectral imaging and the increasing use of multi-colorant printing systems. Spectral color reproduction, i.e. aiming at reproducing the spectral reflectance of an original, first requires a colorant separation for a multi-colorant printing system, followed by halftoning of each the color channels. Spectral vector error diffusion, sVED, has previously been introduced as a tempting alternative for spectral color reproduction, since the method combines the colorant separation and the halftoning in a single step. Only the spectral properties of the Neugebauer primaries are needed as input, and there is no need to invert a complex printer model for the colorant separation. Previously, spectral vector error diffusion has been positively evaluated for simulated prints, assuming a perfect printer and no dot gain. In this study, we evaluate the performance of sVED in practice, for real prints.

Spectral vector error diffusion has been used to reproduce 1000 spectral targets, all within the spectral gamut of the printing system. The resulting color patches have been printed in various print resolutions, using a 10-colorant inkjet printing system. The experimental results reveal a remarkably large difference between the reproduction errors for the printed samples compared to the simulated spectra from the digital halftones. The results show a strong relation between the print resolution and the magnitude of the reproduction error, with lower resolutions giving smaller errors, due to the effect of dot gain in the printing process. The experimental results imply that in its current form, without compensation for physical and optical dot gain, spectral vector error diffusion produces unacceptable spectral and colorimetric reproduction errors, for any print resolutions used in practice.

The results further show that the sVED method in many cases produces color patches that appear noisy and visually unpleasant. By replacing the spectral RMS difference with the ΔE_{94} color difference as criterion in the sVED algorithm, the graininess as well as the resulting color difference was decreased. However, the improvements in colorimetric performance and more visually pleasant reproductions, comes at the cost of an increase in spectral reproduction errors.

Introduction

The interest for spectral color reproduction has increased with the growing field of multispectral imaging. Indeed multispectral imaging offers the great advantage to dispose of the full spectral color information of a surface. When a conventional color acquisition system records the color of a surface under a given illuminant, a multispectral acquisition system can record the spectral reflectance of a surface, and allows us to simulate its color under any illuminant. A spectral color workflow is made of several steps: a spectral color acquisition, a spectral data storage and, finally, spectral color reproduction. Where a conventional color reproduction in the best case can match the colors of an original under a given illuminant, a spectral color reproduction, having the same spectral properties as the original, will match under any illuminant.

In recent years, multi-colorant printing systems, adding extra colorants to the conventional four, have been introduced on the market. The additional colorants included in such systems have been added primarily for increasing the printer gamut and for producing more visually appealing images. However, the extra colorants of such multi-colorant systems also open up the possibility to use them for spectral color reproduction.

The workflow for spectral color reproduction usually involves two sequential steps: a colorant separation, followed by halftoning of each of the different color channels. The colorant separation involves the transformation between the spectral image and the printer colorant combination, which can be achieved by inverting the spectral Yule-Nielsen modified Neugebauer model [1]. After the colorant separation, halftoning is applied for each of the different colorant separations, to create a number of binary images that is sent to the printer. This workflow is similar to conventional color printing, but with an increase in dimensionality due to the additional color channels.

An interesting alternative to this workflow is spectral vector error diffusion, sVED [2]. The sVED method performs directly the transformation from the spectral image to the multi-binary colorant image, thus combining color separation and halftoning in a single step. The technique also requires less knowledge and modeling of the characteristics of the printing system, since it only uses the spectral reflectances of the Neugebauer primaries as input.

Previous studies of sVED have mainly been focused on the halftoning part, only simulating the actual printing, assuming a perfect printer and no dot gain [2, 3]. However, dot gain, both physical (when the size of the printed dots differs from their nominal size), and optical (originating from light scattering within the substrate), is inevitable in the printing process. Therefore, the aim of this study is to challenge spectral vector error diffusion, applying the method to produce real prints, finding out if the method is useful in practice.

Spectral Vector Error Diffusion

Spectral vector error diffusion, sVED, is an extension of the conventional scalar error diffusion, to reproduce spectral images [4]. Just as in conventional error diffusion, sVED is performed in a raster scan mode, where each pixel is binarised and the resulting error is diffused to neighboring pixels.

Spectral vector error diffusion first requires spectral data for the Neugebauer primaries, NPs, of the printing system, i.e. all combinations of the primary inks. For each pixel in the spectral image, the spectral root mean square, RMS, difference between the target spectrum and each of the NPs is calculated. The Neugebauer primary giving the minimum RMS-difference, i.e. having the spectrum closest to the target spectrum, is then selected as output for that pixel. The spectral difference, or the spectral error, between the selected NP and the target spectrum is then calculated and diffused to surrounding pixels, wavelength by wavelength. The weights when diffusing the spectral error are defined by an error distribution filter, in the same way as for conventional error diffusion. In this study, the error distribution filter by Floyd and Steinberg, with the dimension 2 x 3 pixels, is used [5].

Experimental setup

Inkjet printer

The printer used is a Canon ipf 5100 multi-colorant inkjet printer, using pigmented inks. The 10 utilized inks corresponds to the conventional primary colors: cyan, magenta, yellow and black, the lighter colorants: light cyan, light magenta and gray, and the complementary colors: red, green and blue. The spectral reflectances for the 10 colorants are displayed in fig. 1. The set of available primary inks is the original set for the printer, designed for conventional color reproduction. With the aim of spectral color reproduction, however, this set of primary inks are not optimal [6].

The printer can be controlled directly, overriding the internal RIP, using binary bitmap files for each color channel as input. The maximal print resolution corresponds to 1200 dpi, but to assure good registration, 600 dpi is used as the highest print resolution. All prints have been made using the same matt, coated inkjet paper.

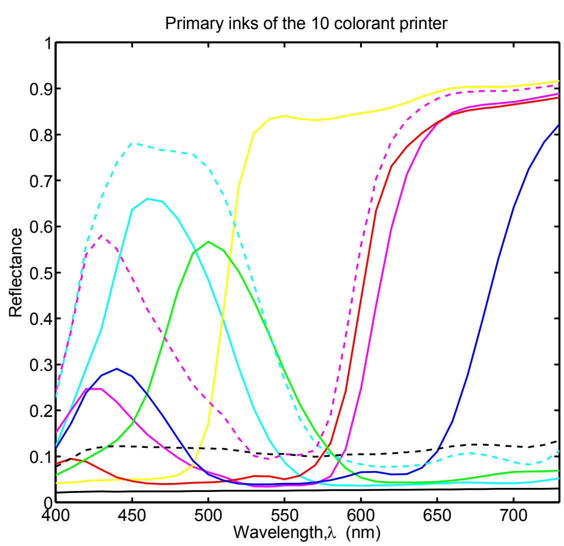

***Figure 1**. Spectral reflectance of the primary inks for the 10-colorant printer.*

Neugebauer primaries

The number of Neugebauer primaries, NPs, for a printing system is defined as 2^n, where n is the number of color channels. This implies 2^{10} = 1024 NPs for the current 10-channel system. However, paper has limited ink receiving capacity and printing many inks on top of each other results in a very dark color. Initial tests indicated that more than 3 ink layers did not result in any additional primaries that would be beneficial in spectral reproduction. Furthermore, any ink combination including black ink will still result in black. In this study, 176 different NPs have been selected out of the 1024 possible candidates, all with the maximum ink coverage 300%.

Spectral measurements

Spectral measurements of the printed color patches were made using a Gretag Machbeth Spectrolino spectrophotometer, using the 45°/0° measurement geometry. The spectral data are in the interval 380 to 730nm, in steps of 10 nm. All colorimetric computations were performed using the CIE 1931 standard observer, and CIE standard illuminant D65.

Target colors

To evaluate the performance of spectral vector error diffusion in practice, a number of target colors are needed. Since the performance of the spectral reproduction for targets outside the spectral gamut of the printing system will be highly dependent on spectral gamut-mapping, we decided to only use target colors within the printer gamut. A test chart of 500 random target colors was created. For each target, a random combination of maximum 4 colorants was created, each with random ink coverage in steps of 25%, with the additional criterion of maximum 300% total ink coverage. Each color channel was then halftoned independently, using scalar error diffusion.

However, after initial tests, an additional set of 500 lighter colors was created. This time with random ink coverage in the range 0:1:30%, for each of maximum five colorants combined. The two test charts give totally 1000 target colors, all within the printer gamut. Used as spectral targets in the experiments are the measured spectra of the 1000 color patches, thus including the physical and optical dot gain originating from the printing process.

Experimental results

The spectral vector error diffusion method, sVED, has been used to reproduce the 1000 spectral targets. For each target color, a patch of the size 90 x 90 pixels was generated and printed in 300 and 150 dpi, respectively. The results for the digital halftones are computed by a simple linear model, averaging the known reflectances of the selected NPs for each pixel in the halftone patch, thus representing an ideal printer without the effect of physical or optical dot gain.

Histograms of the distributions of spectral and colorimetric errors for the 1000 targets are displayed in fig. 2. Table 1 lists the mean and maximal errors, in terms of the spectral RMS difference and the CIE 1994 color difference, ΔE_{94}, for the total set, as well as for the two subsets.

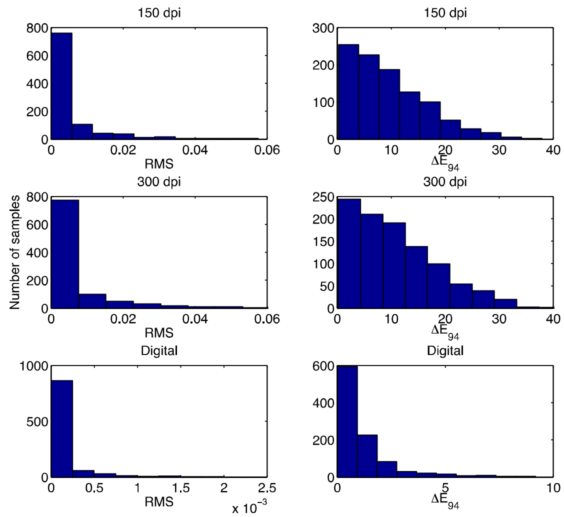

Figure 2. The distributions of spectral RMS difference and the colorimetric error ΔE_{94}, between the target spectra and the reproductions of 1000 targets.

Table 1. Error metrics for 1000 color patches, haltoned using sVED and printed in 300 and 150 dpi.

Test chart	Print res.	RMS		ΔE_{94}	
		Max	Mean	Max	Mean
All 1000	300 dpi	0.076	0.0061	41.5	10.7
	150 dpi	0.058	0.0049	37.9	9.49
	Digital	0.0025	$1.3 \cdot 10^{-4}$	9.18	1.23
500 Dark	300 dpi	0.049	0.0013	31.3	6.17
	150 dpi	0.047	0.0011	29.3	5.28
	Digital	0.0025	$1.1 \cdot 10^{-4}$	8.81	1.38
500 Light	300 dpi	0.076	0.0108	41.5	15.3
	150 dpi	0.058	0.0087	37.8	13.7
	Digital	0.0022	$1.5 \cdot 10^{-4}$	9.18	1.07

From fig. 2 and table 1, it is clear that there is a large difference between the results for the digital halftones compared to the actually printed samples. The printed targets produce large errors, with only small improvements when reducing the print resolution from 300 to 150 dpi. Also noticeable is the surprisingly large difference between the dark and the light test charts, where the 500 lighter targets produces a mean RMS error about 8 times larger than the 500 darker targets. However, this difference is less pronounced for the digital halftones, where the errors are in the same order of magnitude. The explanation for this is that the dot gain is larger for the lighter test chart, where all the patches have low or intermediate ink coverage. In the dark chart, many patches include full ink coverage for at least some of the inks, which reduce the effect of the dot gain.

Print resolution

To further investigate the relationship between the print resolution and the reproduction errors, a set of 40 patches, known to be difficult to reproduce, were selected and printed in 75 and 37 dpi print resolution. The 40 target colors, selected from both the dark and light test charts, are displayed in fig. 3. The corresponding patches after halftoning by sVED are displayed in fig. 4, at the resolution 75 dpi. Notice that the colors in figs. 3 and 4 are simulated by converting the known spectral reflectances of the printer NPs to sRGB, under standard illuminant D65. Hence, they are not representative to the actual spectral reproductions, using the multi-channel printer.

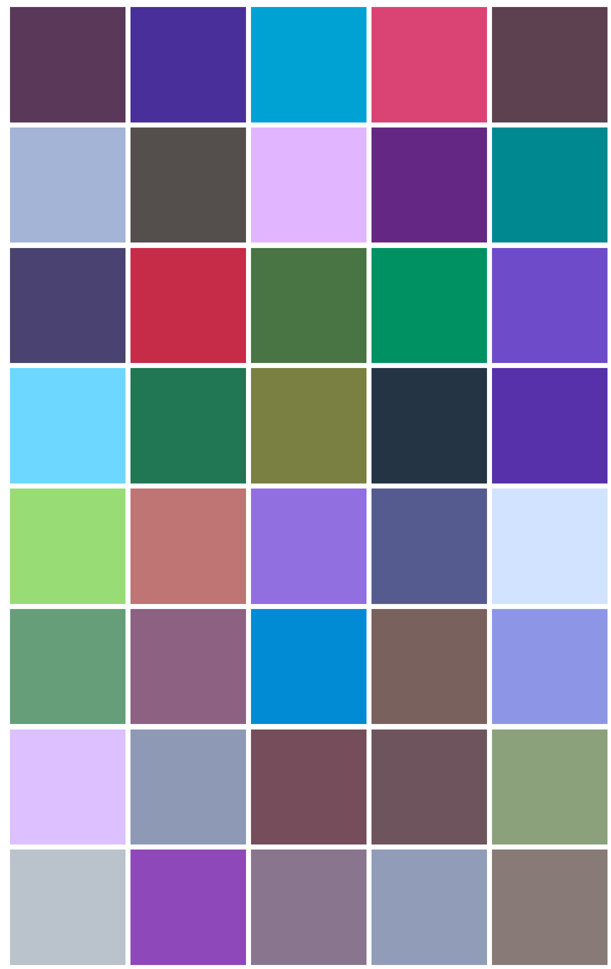

Figure 3. Visualization of the 40 target patches used to evaluate the performance of sVED for different print resolutions, rendered under standard illuminant D65.

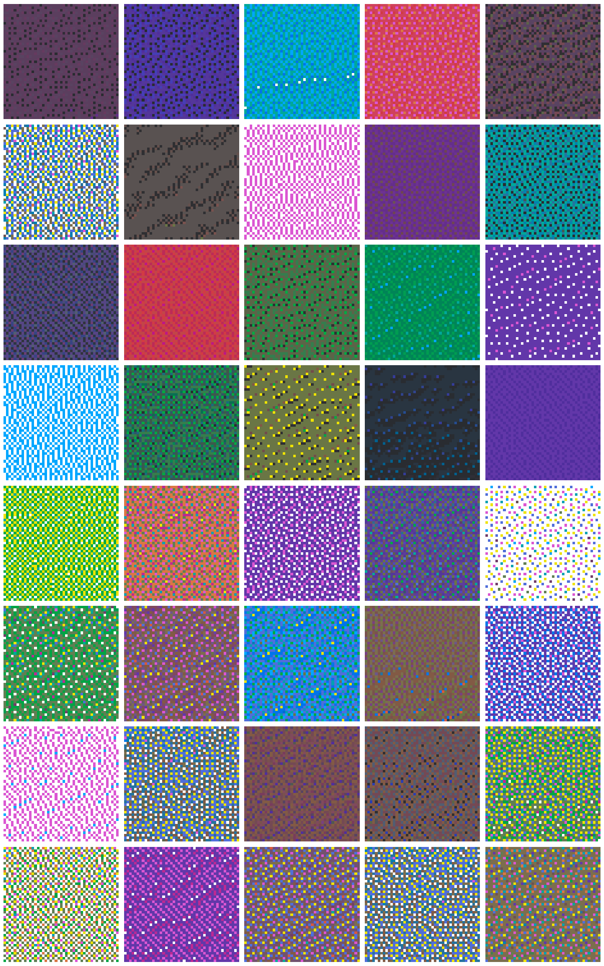

***Figure 4**. Visualization of the 40 patches from fig. 3, after haltoning by sVED, under standard illuminant D65. The resolution corresponds to 75 dpi.*

Table 2. Error metrics for 40 color patches, haltoned using sVED and printed in the resolutions 300, 150, 75 and 37 dpi.

Resolution	RMS		ΔE_{94}	
	Max	Mean	Max	Mean
300 dpi	0.064	0.011	34.1	13.8
150 dpi	0.054	0.0091	31.7	12.1
75 dpi	0.040	0.0063	24.7	9.54
37 dpi	0.017	0.0028	14.1	6.28
Digital	0.003	0.0006	7.19	2.68

The resulting spectral and colorimetric errors from the reproduction of the 40 target colors, printed in different resolutions, are given in table 2. The 40 patches that were selected out of the 1000 are among the ones that are most difficult to reproduce, which explains the increase of the mean errors for this smaller set, when compared to table 1. The results clearly illustrate the strong relation between the print resolution and the magnitude of the error, with lower resolutions giving smaller reproduction errors. This fact is further illustrated in fig. 5, depicting the max and mean errors for increasing print resolutions (where 0 refers to the digital halftones). Figure 6 displays the target spectra compared to the prints in different resolutions, for 4 of the 40 targets. The result for the digital halftone (blue) is close to the target spectra (dotted), while the reflectance for the printed patches gets lower (i.e. darker) with increasing print resolution. The explanation is simply that the effect of the dot gain is reduced with lower print resolution. Both the physical and optical dot gain increases when the size of the halftone dots decreases, i.e. with a higher print resolution [7].

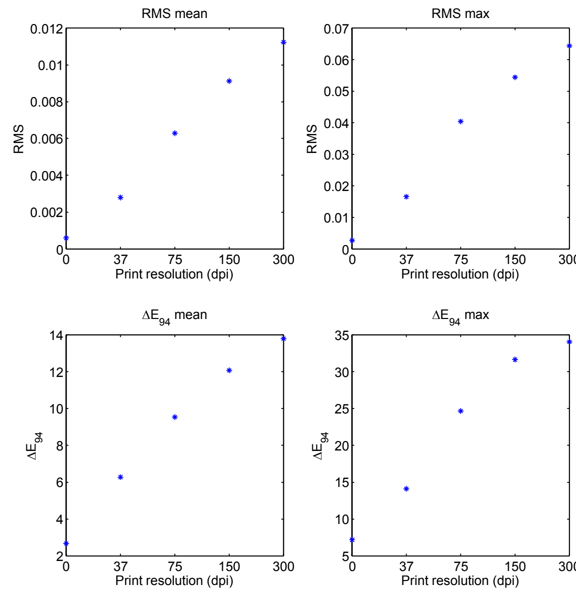

***Figure 5**. Error metrics for the 40 color patches printed in various resolutions.*

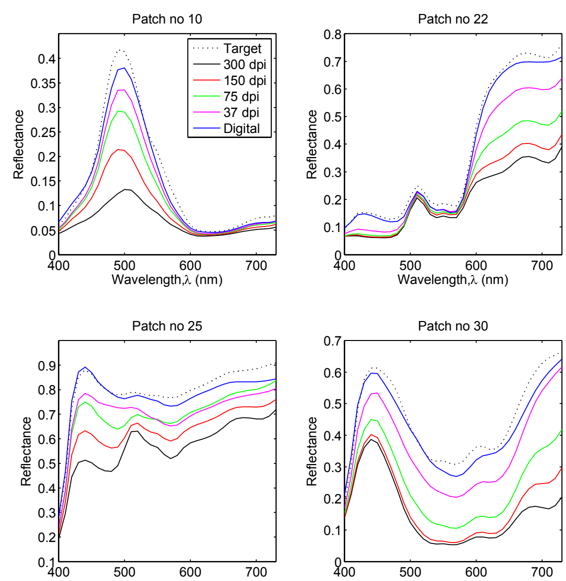

***Figure 6**. Target spectra (dotted line) compared to the printed reproductions in different print resolutions.*

Reducing the graininess in sVED

For many of the color targets, the sVED method produces patches that are noisy and visually unpleasant. To explain the reason for this, fig. 7 demonstrates the sVED process for a sequence of 4 pixels, where the target is a neutral gray color (the left patch in fig. 8). The dashed blue line represents the target spectrum (including the spectral errors diffused from previous pixels), the full red line is the selected NP, and the dotted line is the resulting spectral error, which is diffused to the surrounding pixels.

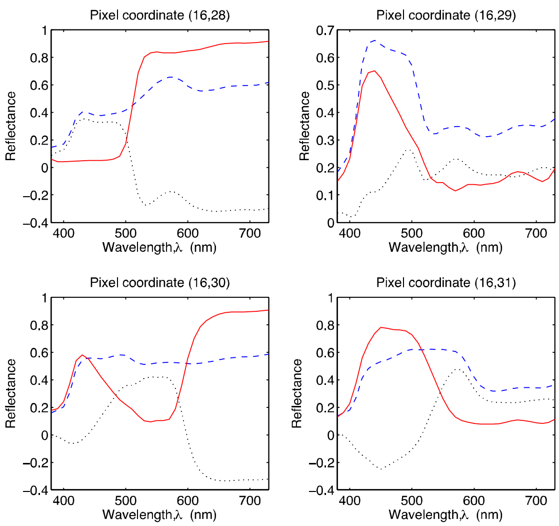

Figure 7*.The sVED algorithm for a 4-pixel sequence. The dashed line is the target spectra, the full line is the selected NP and the dotted line is the spectral error that is diffused to surrounding pixels.*

In the 4-pixel sequence displayed in fig. 7, the selected NPs corresponds to: 1) yellow, 2) the combination of light magenta and light cyan, 3) light magenta, and 4) light cyan. In each pixel, the selected NP is the one with the spectrum closest to the target, in terms of RMS difference. However, producing a neutral gray color by combinations of cyan, magenta and yellow may produce low RMS when averaged over a larger area, but it will also result in a noisy and unpleasant color reproduction. In fig. 8, the target gray is displayed to the left and the result after halftoning by sVED is displayed in the middle column. The top row visualizes the results in 75 dpi print resolution, and the magnification in the lower row corresponds to 15 dpi. It is clear that the sVED method in this case has produced a noisy pattern, and also a visible color shift.

The results from figs. 7 and 8 implies that the spectral RMS difference may not be the optimal criterion when selecting the NP in the sVED method, at least not if the visual impression of the reproductions are to be taken into account. The spectral RMS difference simply sums up the difference for each wavelength, which does not necessary results in a spectrum visually similar to the target. To overcome this, possibly a criterion more closely related to the human visual system could be used when selecting the NP in the sVED algorithm. In order to achieve a better visual match, we applied the CIE 1994 color difference, ΔE_{94}, as criterion, hoping to avoid the situation illustrated in fig. 7. However, since the aim is to achieve spectral color reproduction, not colorimetric, it is still the spectral error that is diffused to the surrounding pixels, and included in the computations for those pixels.

For the gray patch in fig. 8, the RMS-error increases by a factor of 4 when ΔE_{94} is used as criterion in sVED, but the color difference decreases by a factor of 10 (for the digital case, assuming no dot gain). The impression is also that the reproduction is visually closer to the target, and less noisy (fig. 8, right patch).

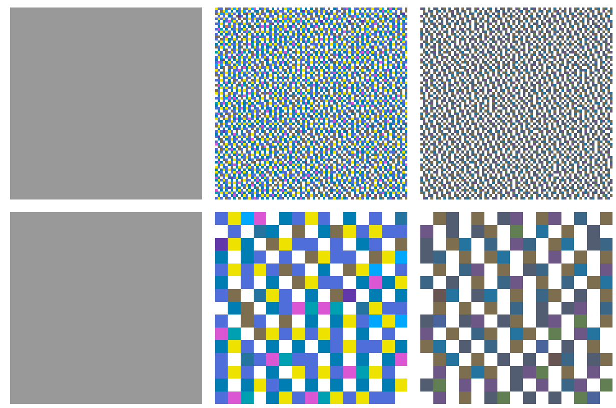

Figure 8*. Halftoning of the gray target to the left by sVED, selecting the NPs using RMS (middle) and ΔE_{94} (right). The resolution for the upper row corresponds to 75 dpi and the magnification below to 15 dpi.*

Table 3 lists the mean and max of the spectral and colorimetric errors for all 1000 targets, when the CIE 1994 color difference, ΔE_{94}, has be used as criterion in the sVED algorithm. Compared to the corresponding results in table 1, the alternative criterion produces lower color differences, at the expense of increasing the spectral RMS errors. The differences in performance are most noticeable in the digital case, without the effect of the dot gain.

Table 3. Error metrics for 1000 color patches, haltoned using sVED and minimizing the ΔE_{94} color difference.

Test chart	Print res.	RMS		ΔE_{94}	
		Max	Mean	Max	Mean
All 1000	300 dpi	0.092	0.0070	39.8	8.79
	150 dpi	0.074	0.0062	38.3	8.12
	Digital	0.027	0.0014	8.19	0.77
500 Dark	300 dpi	0.073	0.0031	33.2	5.09
	150 dpi	0.057	0.0028	31.9	4.67
	Digital	0.027	0.0015	7.01	0.86
500 Light	300 dpi	0.092	0.011	39.7	12.5
	150 dpi	0.074	0.0095	38.3	11.6
	Digital	0.024	0.0014	8.19	0.67

Figure 9 displays a comparison between RMS (middle) and ΔE_{94} (right) as criterion in sVED, for three additional patches rendered in 75 dpi under D65. For all of them, the RMS error increases for the ΔE_{94} method, at the same time as the color difference is decreased. Besides producing a better visual match to the target (left), the reproductions using the ΔE_{94} method usually appear less noisy, with a more pleasant spatial distribution, compared to the RMS method.

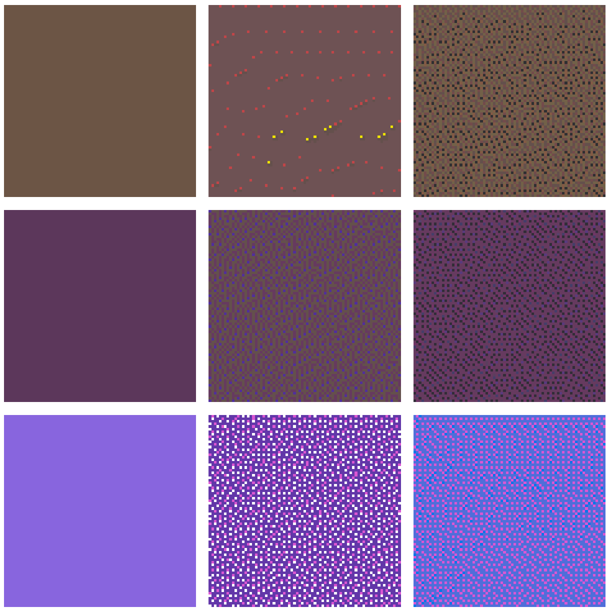

Figure 9. *Left column is the target color; middle column is the result after sVED using RMS, and right column after sVED using ΔE_{94}.*

Discussion and continuation

In this study we have applied the ideas of spectral vector error diffusion, which has previously been used in simulations, in practice. Evaluation of the method by producing and printing a large number of target colors in different print resolutions, has revealed that the spectral and colorimetric errors are unacceptable for all print resolutions that can be used in practice, due to dot gain in the printing process. In previous work it has been shown that the performance of sVED can be improved by clipping or scaling the modified signal to control the error, or by designing new filters for error diffusion [3]. However, for the method to be of any real use in practice, compensation for the dot gain must be included.

For many target colors, the resulting reproduction after sVED appears noisy and visually unpleasant. By using the ΔE_{94} color difference as criterion when selecting the NP for each pixel, reduces the noisy impression as well as the resulting ΔE_{94} color difference. However, the improvements in colorimetric performance and more pleasant reproductions, comes at the cost of increased spectral error, which contradicts the aim of spectral reproduction.

What might have appeared like a promising road when simulating the results, assuming an ideal printer, turns out to be far more problematic under real conditions. In a conventional color reproduction workflow, each colorant separation is individually compensated for the dot gain before the halftoning is applied. In a sequential method such as sVED, combining the colorant separation and the halftoning in a single step, it is not straightforward to compensate for the dot gain. To modify the Neugebauer primaries to include the effect of the dot gain is difficult, since the characteristics of the optical dot gain will depend on the surrounding area. One idea that will be evaluated in the continuous work is to incorporate the Yule-Nielsen *n*-factor in the sVED method [8].

However, the authors' belief is that in order to make the concept of sVED successful in practice, essential modifications are needed. Accurate spectral reproduction will require the development of new halftoning methods, incorporating models for the effect of both physical and optical dot gain. Besides the halftoning, future development in the field of spectral color reproduction will also require research and improvements in many related areas, such as spectral gamut mapping, new error metrics, optimal ink selection and optical modeling. These areas, and other aspects of spectral color reproduction, will be thoroughly studied in the recently started project *Colour Printing 7.0*: www.cp70.org.

References

[1] D. R. Wyble and R. S. Berns, "A critical review of spectral models applied to binary color printing," Color Research and Application 25(1), (2000).

[2] J. Gerhardt and J. Y. Hardeberg. Spectral colour reproduction by vector error diffusion. Proc. CGIV, pg 469, (2006).

[3] J. Gerhardt and J. Y. Hardeberg. Controlling the error in spectral vector error diffusion, Proc. SPIE 6493, (2007).

[4] M. M. Kouzaki, et al., Evaluation of spectral color reproduction by vector error diffusion method, Proc. SPIE 3963, pg 429, (1999)

[5] R. W. Floyd and L. Steinberg. An adaptive algorithm for spatial greyscale. Proc. Society for Information Display,17 (2): 75–77, (1976).

[6] A. Alsam and J. Y. Hardeberg, Optimal colorant design for spectral colour reproduction, Proc. IS&T SID Twelfth Color Imaging Conference , pg. 157, (2004).

[7] D. Nyström and L. Yang, Physical and Optical Dot Gain: Separation and Relation to Print Resolution. Proc. Iarigai, pg. 337, (2009).

[8] J. Gerhardt and J. Y. Hardeberg, Simple Comparison of Spectral Color Reproduction Workflows, Proc. SCIA, pg. 550, (2009).

Author Biography

Ole Norberg received his MSc. in Physics from Umeå University, Sweden in 1996 and his Ph D in Media Technology from Linköping University, Sweden in 2006. Ole has long experience from R&D work in the paper Industry with special interest in paper optics, digital printing and color science. At present his is working as a Senior Research Scientist at Voxvil AB and part time at the Digital Printing Center at Mid Sweden University.