HANS print smoothness optimization and continuous control

Ján Morovič, Peter Morovič, HP Inc., Spain

Abstract

A print property closely related to color and with a strong dependency on the choices made for the sake of color separation is the smoothness of a printed pattern. While certain obvious choices are simply related to grain, such as the degree of black ink use in colorant-channel imaging pipelines, the domain of all possible printable patterns that HANS (Halftone Area Neugebauer Separation) [1] provides access to eludes such simple rules of thumb. At the same time, the magnitude of the new control domain gives access to new patterns that allow for a reduction of grain and an increase in smoothness beyond what is possible using conventional techniques. In this paper, a framework for optimizing smoothness is presented first, followed by a mechanism for varying it in a controlled and continuous way, when the aim is a given trade-off between smoothness and other attributes such as colorant use efficiency or robustness to perturbations.

Introduction

While color is the most prominent variable controlled in printing, there is redundancy in achieving it, which opens the door to varying other print attributes at the same time. Due to the inherent color-redundancy of printing systems, there are entire sets – metamer sets – of either colorant combinations (in traditional imaging pipelines) or of Neugebauer Primary area coverage vectors (in HANS pipelines) that result in a given color. Some of these will use less ink than others while some will yield smoother-looking prints and others will result in grainier ones.

Considering grain, the main control mechanisms used for varying it have until now been choices in terms of halftoning (i.e., the use of green versus blue noise halftone matrices [2] would result in clusters that increase the graininess of a print) and black generation (i.e., increasing the strength of gray component replacement would result in more black colorant being used and therefore in more visible, grainier printed patterns). While this yielded variation in grain, the relationship between the choices and the resulting level of grain has been indirect, often discrete (i.e., the use of one halftoning matrix versus another) and usually global (i.e., a single set of choices had to be made for all colors).

In this paper, an approach to smoothness optimization and control is presented that allows for direct choices to be made about it, where control can be exercised in a continuous way and that benefits from an access to the vast variety of printable patterns that the HANS framework provides.

Before proceeding with grain prediction, optimization and continuous control using HANS, it is worth reviewing the basic principles of this imaging pipeline framework that enable them.

The key insight of HANS is that print can be controlled at the level of its atomic states, i.e., the contents of individual halftone pixels. Looking at such pixels, they can specify any combination of a printing system's colorants. E.g., for a three-colorant, binary system, there are eight possible halftone pixel states: blank substrate, one colorant at a time (three states), two colorants at a time (another three states) and all three colorants combined in a single pixel. For a CMY system, the states are B, C, M, Y, CM, CY, MY and CMY. These states are also known as the Neugebauer Primaries [3] and a halftone pattern can be specified by assigning relative area coverages to each one of them. The resulting Neugebauer Primary (NP) area coverage (NPac) vector expresses the probability with which each of a printing system's NPs is to be used within a spatial neighborhood.

An important consequence of specifying NPacs instead of colorant amounts is that they represent a convex combination of at-pixel states and that their color therefore is the result of convexly combining the color of their constituent pixel colors (in an appropriate color space, such as Yule-Nielsen (YN) corrected XYZ [4]). This also means that NPacs can be convexly combined while remaining in that same convex domain. Such associativity [5] means that the convex combination of relative area coverage weighted NP colors can also be seen as the convex combination of two or more constituent sub-patterns, i.e., individual patterns that, when convexly combined, give rise to a continuum of patterns between them (Figure 1).

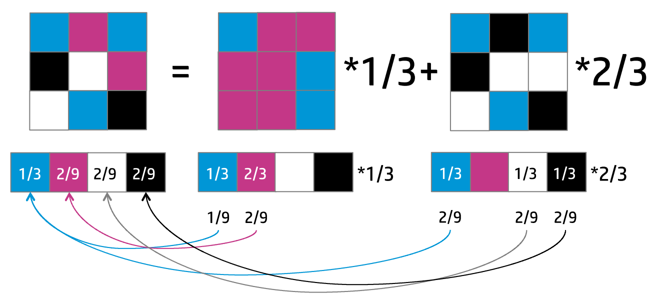

Figure 1. Associativity of pattern convex combination.

Since what varies as two patterns are combined are only convex weights that represent area coverages, a transition obtained by varying such weights occurs in the optical additivity domain resulting from the human visual system's limited spatial resolving power and is therefore smooth. A second feature of convexly combining NPacs is that only NPs that have non-zero coverage in those NPacs are used in the transition between them. E.g., if the end points only contain colorants by themselves, no overprinting between colorants will take place in a transition between them. This property of convex combinations can be thought of as "closedness." Given arbitrary patterns, determined individually, their convex combinations will therefore form a smooth transition in visual terms, which will be relied on when optimizing color to NPac mappings, such as the ones described here, which are aimed at controlling a print's smoothness.

The following sections will introduce the modeling and optimization of print smoothness, followed by a mechanism that allows for continuous control over it, including results about the success of the approach on a test system.

Smoothness prediction

There are a number of ways in which the grain of a halftone pattern can be quantified, all of which share the basic feature of linking it to the standard deviation of local colors within a neighborhood [6, 7]. The greater such a standard deviation of local component colors, the more visible the pattern they form and the greater the resulting perception of grain. However, existing methods typically involve psychometric evaluation and/or a scanning of prints whose grain is then predicted from the resulting images. Since grain optimization requires a metric that is applied before any printing takes place, it is first necessary to identify both a suitable metric for application to digital data and an appropriate domain to which to apply it. A further requirement for a good grain-optimization metric will be simplicity and computational efficiency, since it will need to be applied to 10s or 100s of millions of candidate printable patterns.

To test candidate metrics, a small-scale experiment was first conducted in which six observers with color-normal vision were asked to view 16 color patches, printed using CMYK inks on a matte paper substrate, and judge their level of grain in terms of three categories: low, medium and high. The result, using the mode of individual category judgments, was a categorization where eight patches were judged to have low grain, five patches were identified as having medium grain and three as having high grain (Figure 2). A candidate grain optimization metric will be considered successful if its predictions separate low from high grain.

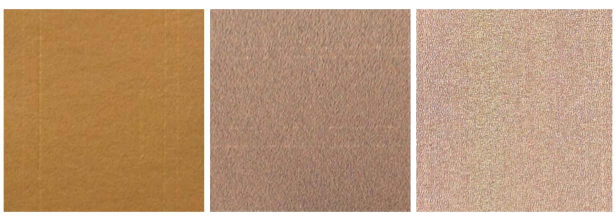

Figure 2. Scans of example color patches with (left) low, (center) medium and (right) high levels of perceived grain.

Since existing metrics all share the use of standard deviation, a starting point was to apply it directly to NPacs. Because an NPac specifies area coverages of a printing system's NPs, its standard deviation can be obtained from the area coverage weighted color difference between each NP's color and the NPac's color:

$$g = \sum_{i=1}^{n} a_n \left\| NP_n - T \right\| \tag{1}$$

where g is the grain metric, a_n is the area coverage of the n-th NP, NP_n is its color, T is the color of the NPac (i.e., the area-coverage weighted sum of its constituent NPs' colors) and || || is a color difference metric. In this simple framework, the best results (Figure 3) were obtained by using YN-corrected XYZ to compute T and ΔE2000 as the color difference metric. As can be seen, there is a significant overlap between the high and low grain patches here, which does not make this a suitable solution. This, in fact, is not surprising, since the very simplistic approach used here is entirely agnostic of the halftone matrix used by the PARAWACS algorithm [8] to provide print-ready halftones for the 16 NPacs in question and would yield the same predictions regardless of whether halftone structure was clustered or not, which however would result in very different levels of grain of the halftoned patches.

The next step, therefore is to go from an NPac to generating a corresponding halftone pattern and to then derive a metric from that halftoned image. If all that was done were to compute a standard deviation of the halftoned image's pixels, then the exact same result as before would be obtained, since the pixels' statistics are identical to the NPac. Two alternatives present themselves here: first, to take into account optical dot-gain by filtering the halftones in a YN-corrected XYZ domain, and, second to take into account the contrast sensitivity function (CSF) of the human visual system. Since for the kinds of printed content where smoothness of halftone patterns is important, prints are often inspected up-close, the benefit of taking CSFs into account is limited. This however, would be a useful extension of what follows, in cases where smoothness optimization would be applied to prints made for viewing at greater distances.

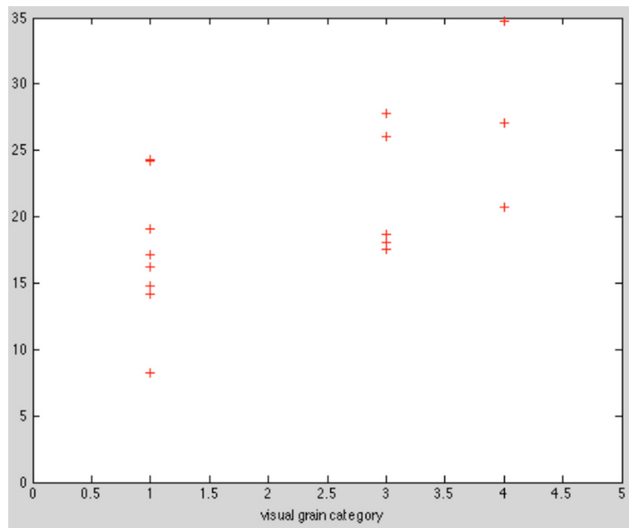

Figure 3. Grain predictions using simple NPac-applied grain metric.

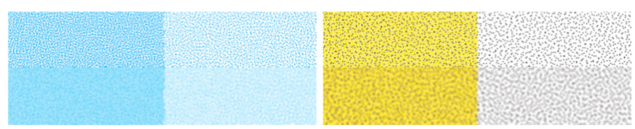

Figure 4. A pair of (left) low and (right) high grain halftones, (top) before and (bottom) after filtering in Yule-Nielsen XYZ over 3×3 local neighborhoods.

In the case of interest here, where prints are inspected up-close, halftones obtained from NPacs using the PARAWACS halftoning algorithm will only have the effect of optical dot-gain simulated. This is done by taking the halftoned image and replacing the NP at each pixel with its corresponding colorimetry, predicted using the RONT model, which extends Kubelka-Munk and Yule-Nielsen by a T matrix that reinforces spectral correlation [9]. NP colorimetry computation is here performed by taking NPacs that use only a single NP at a time (i.e., representing some unit area with that NP at each location) and then pass through the three stages of color prediction to arrive at spectral reflectance and from there colorimetry. Then, each pixel's colorimetry is replaced by the mean of its' $r \times r$ neighborhood's pixels, computed in YN XYZ and then transformed back to XYZ (Figure 4). This processing is akin to a convolution filtering with a mean filter and therefore the result is, as expected, a blurring of the original halftone, which is broadly consistent with the effect of dot-gain.

Finally, the grain metric is computed as in equation 1 over the halftone patch's filtered pixel colorimetries instead of over area-coverage-weighted NP colorimetries, resulting in the following grain predictions (Figure 5).

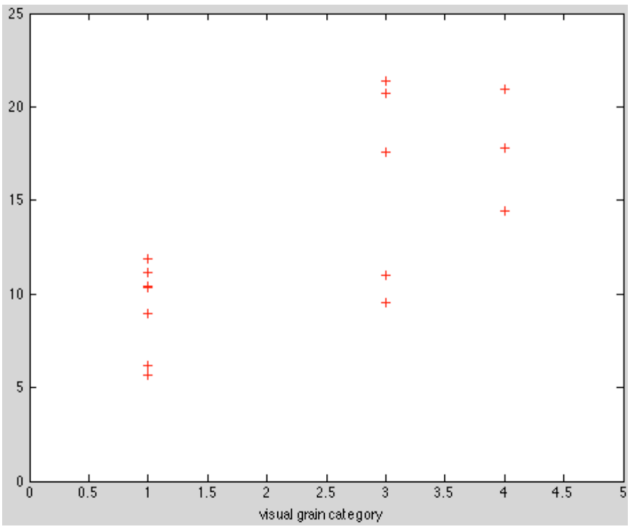

Figure 5. Grain predictions from metric applied to YNN-filtered halftone.

As can be seen, color patches judged to have low and high grain are successfully separated by the grain metric derived from YN-filtered halftone images. However, patches judged to have medium grain have predicted grain values interspersed with the predictions of both high and low grain patches. To understand whether this level of performance was acceptable, the 16 NPacs were printed in the order of their predicted levels of grain and again inspected by the original six observers in comparison with their manual ordering and with the ordering obtained using the NPac-only metric (Figure 6). Here the judgment of the observers indicated a similar level of performance for their original grain ordering and for the ordering predicted from YN-filtered halftone colorimetries.

Figure 6. Printed halftone patterns ordered left to right from low to high by (top) perceived grain, (middle) grain predicted from NPacs and (bottom) grain predicted from YN-filtered halftones.

To test the suitability of this, simple, digital grain metric, it was then used in the context of optimizing a HANS color separation. While the grain metric derived here is based on a clustering into categories of grain (low-medium-high), it may not be suitable for a ranking of arbitrary patches in order of grain. However, it does allow for distinguishing of the aforementioned categories. This in turn also allows for optimization, since here in every neighborhood the halftone that has the lowest (or highest) metric will be selected, hence this too is a categorization problem and does not strictly require the ability to strictly order by grain.

Smoothness optimization

Since the NPac domain of HANS provides the ability to smoothly transition between arbitrary NPacs, it enables an optimization of color separation, where individual nodes of a colorimetry to NPac look-up table (LUT) may freely and independently have their content determined. At high-level, the approach here is one of mining NPac space, followed by an evaluation of mined NPac properties, by a grouping of mined NPacs by color and finally by a selection of one NPac per color group.

More specifically, the following process can be used to optimize a color separation for smoothness:

1. Generate NPacs using a RANSAC approach [10], where random values are assigned to a printing system's NPacs and where these are subsequently adjusted to sum to one. To make this NPac synthesis more useful, the process can be performed for specific choices of NPs having non-zero area coverages, since the process would otherwise favor NPacs with many, or even all, NPs of a system being used. Given a printing system with *n* NPs, NPacs with 2, 3, 4, … and up to *n* NPs can be generated in turn by generating the appropriate number of random values and then assigning them randomly to NP indices.
2. For each of the NPacs generated in step 1, predict their colorimetry using the RONT model and their grain using the standard-deviation computation applied to YN-filtered halftones generated from an NPac.
3. Subdivide color space into bins of given dimensions, e.g., in L*, a* and b* and assign each NPac from step 1 to its corresponding bin based on its predicted color (from step 2). The bins here can be thought of as paramer sets. As bin size reduces, these paramer sets approach being metamer sets.
4. For each bin from step 3, select the NPac that has the lowest grain score.

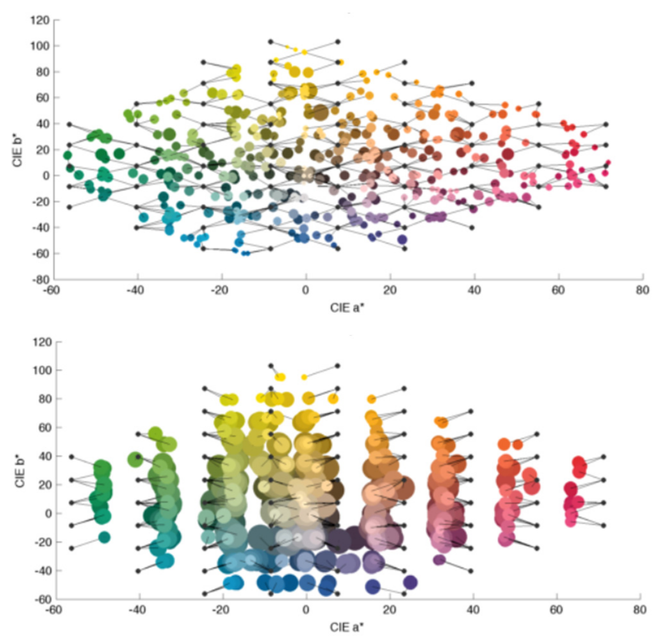

Figure 7. NPacs with (top) lowest and (bottom) highest grain, where the chosen NPac per bin's marker shows its color and the marker's size indicates the grain score. The black lines connect the NPac's predicted color to the centroid of the bin that contains it.

The result of the above process is one NPac per color space bin, where that NPac has the least amount of grain in that part of color space, given the sampling used.

The above approach was used on a printing system with four inks, CMYK, using HP Heavyweight Coated paper and having the ability to print up to one drop per ink per halftone pixel. In this case there are 3^4=81 NPs since up to two drops could be specified for each ink at a halftone-resolution pixel, which results in three levels per ink per pixel. The RONT model tuned to this printing system had a mean accuracy of color prediction of 0.9, a 95th percentile of 2.2 and a maximum error of 3.4 ΔE2000. An oblong enclosing the colors of the 81 NPs was subdivided into 25^3 bins and a total of 13 171 174 NPacs were generated, evaluated and binned as described above. The result were 1257 non-empty bins, grouping the above NPacs by color (Figure 7).

Finally, the NPacs with both lowest and highest predicted grain were kept for each bin. Delaunay-tessellating each of the two sets of NPac-color sets in color terms then provides the means to take color inputs, gamut mapped to the system's color gamut, and interpolating NPacs for them both for grain minimization and maximization. Using the two LUTs for processing image content then resulted in prints with very different degrees of grain and in both cases with consistent levels of grain within each of the prints (Figure 8).

Figure 8. Photos of two prints made on the same printing system using the (left) grain minimizing and (right) grain maximizing color to NPac LUTs.

Even with the very limited ground truth from which the simple grain metric was derived here, the NPac choices made on the basis of their predictions resulted in consistently low versus high grain, depending on the choices made per color bin. These per-bin choices can be thought of each as akin to the simple experiment conducted initially to derive the metric, whereby all NPacs in a single bin are clustered into the low-medium-high categories by the metric and the lowest and highest are kept. Again, since the purpose of the metric is not an absolute ordering of all NPacs by grain, but a relative clustering, and this being performed for a uniform sampling of color space, the end result is a consistent choice of low and high grain NPacs that jointly constitute low-grain and high-grain LUTs.

Continuous control

In addition to enabling the per-bin optimization of NPacs described above, the convexity, associativity and closedness of the unconstrained NPac domain also provides access to the continuum between individual optimized LUTs. Since the content of each bin is chosen independently of the other bins' contents, it can also be adjusted independently and, as long as the resulting NPac has the correct color associated with it, it can be tessellated with other NPac-color pairs to form a valid color to NPac mapping.

While it is in some cases desirable to maximize the smoothness of a print (and therefore minimize its grain), in other cases some grain may be desirable since it may provide some degree of robustness to perturbations in a printing system. E.g., a missing print head nozzle or a change in drop size may be less noticeable when printing a grainier than a smoother print. As a result, it may in some cases be desirable to choose a level of grain that is not at the minimum possible.

Here the most direct solution is to run the optimization presented above and to select not the minimum grain at each bin, but a grain level that is at a certain, higher level. While this is certainly a valid approach, it comes with two challenges: the need to re-compute an optimization (or to re-select NPacs per color bin from having kept all of the NPacs as an optimization proceeded) and the need to know what specific level of grain is desired. Here the first challenge is about compute resources while the second introduces the risk of having to iterate until the correct level of grain is arrived at.

A simple alternative that the HANS NPac domain enables is to blend LUTs, since the NPac domain is associative [5]. However, before proceeding with the blending, it is beneficial to perform a regular resampling of the minimum and maximum LUTs obtained by the optimization. Taking a regular grid in CIELAB and interpolating NPacs for each of its nodes gives us two NPac sets where corresponding members match in color and differ only in grain.

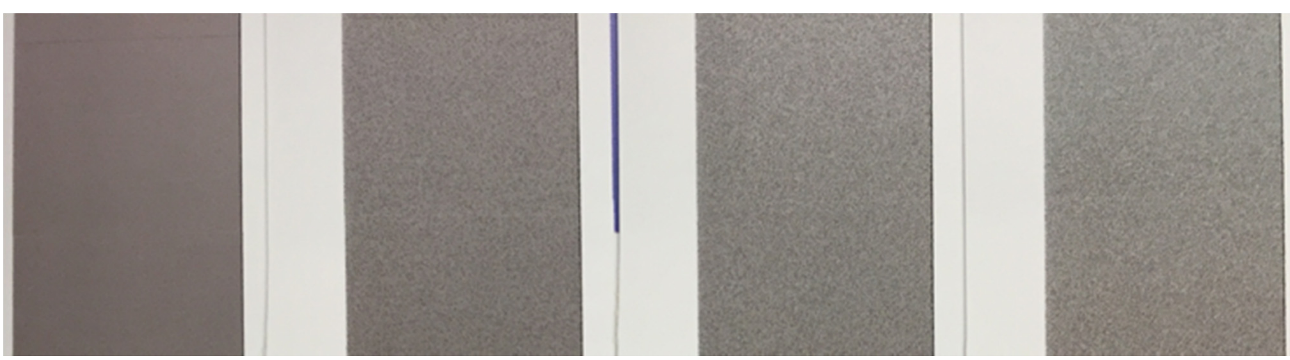

Figure 9. Prints made using (from left to right): grain minimization (i.e., w=0), w=0.25, w=0.5 and grain maximization (i.e., w=1) LUTs.

From such colorimetry-regular minimum and maximum grain LUTs, it is then possible to compute LUTs in-between by assigning weights to the two extreme LUTs that add up to one. This results in LUT members of an intermediate, blended LUT being computed from the minimum and maximum LUTs as follows:

$$ac_{\mathrm{blended},n,i} = w{*}ac_{\max,n,i} + (1\text{-}w){*}ac_{\min,n,i} \quad (2)$$

where ac represents the area coverage of the n-th NP in the i-th NPac and w is the weight of the maximum-optimized LUT. Note that more LUTs could be blended in this way, using a set of convex weights. E.g., the minimum and maximum grain LUTs could be combined with the minimum and maximum ink use ones.

Looking at the result of LUT blending of the minimum and maximum grain-optimized LUTs in Figure 9, it can be seen that there is a continuous increase of grain as we proceed from minimum to maximum. It can also be seen from Figure 9 that the change is not linear, which can be solved by characterizing the relationship between w and perceived grain. Again, using the small panel of 6 observers, the w-grain relationship shown in Figure 10 was determined, which then allows for linear control over grain between the extremes accessible on the test system used here.

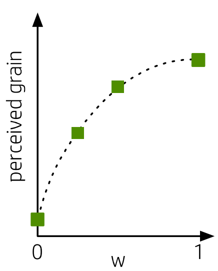

Figure 10. Relationship between perceived grain and w used to blend minimum and maximum grain LUTs.

Finally, it is worth noting that all of the above variation in grain was achieved solely by varying NPac composition and that there is another, important vector that could be adjusted, which is the use of halftoning matrix, where greater clustering can increase grain. To use the above method with a different matrix however requires a new optimization since the grain prediction metric takes the halftone matrix into account and would not preserve the categorization or choices made for another matrix.

Summary

Color separation choices provide control over a variety of print attributes and while this is the case already with colorant-space approaches, using the HANS NPac domain for affecting them provides a host of additional benefits. Choices for individual colors can be made in isolation, due to the convex, closed nature of transitioning in this domain. Making choices about NPacs also benefits from their close proximity to printed properties, which allow for effective predictions to be made on their basis.

In this paper, results were shown for making predictions about grain on the basis of digital halftones and for using such models in the process of optimizing color separations for grain. Finally, it was also demonstrated that the HANS NPac domain allows for continuous transitioning between individually-optimized separations, which, in turn, allows for interactive user control over print properties.

The framework presented here has in the past been used for controlling ink use [1], and will in future work be applied to other print properties.

Acknowledgements

The authors would like to thank their colleagues at HP Inc. for their support, in particular Africa Real, Albert Serra, Annarosa Multari, Hector Gomez, Joan Enric García, Joan Manel García Reyero, Jordi Arnabat, Martí Rius, Pere Gasparin, Victor Diego and Xavi Fariña.

References

[1] J. Morovič, P. Morovič, J. Arnabat "HANS – Controlling Inkjet Print Attributes Via Neugebauer Primary Area Coverages," *IEEE Transactions on Image Processing*, 21, 2, (2011)

[2] R. Ulichney "The Void-and-Cluster Method for Generating Dither Arrays", *Human Vision, Visual Processing, and Digital Display IV, Proc. SPIE,* 1913, 332-343, (1993)

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

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

[5] P. Morovič, J. Morovič, J. M. García–Reyero "HANS: Meet The CMY Metamers," IS&T/SID Color and Imaging Conference '11, 7–11 November 2011, San Jose, CA, 229–233, (2011)

[6] K. Y. Lee, Y. Bang, H. K. Choh "Characterization of '2D noise' print defect," *Proc. SPIE 7242*, 72420M, (2009)

[7] ISO "13660:2007 (e) standard. Information technology – office equipment – measurement of image quality attributes for hardcopy output – binary monochrome text and graphic images." (2001)

[8] P. Morovic, J. Morovic, J. Gondek, M. Gaubatz, R. Ulichney "PARAWACS: Color Halftoning with a Single Selector Matrix," *24th IS&T Color and Imaging Conference*, San Diego, CA, USA (2016)

[9] P. Morovič, J. Morovič, X. Fariña, P. Gasparin, M. Encrenaz, J. Arnabat "Spectral and color prediction for arbitrary halftone patterns: a drop-by-drop, WYSIWYG, "ink on display" print preview," *23rd IS&T Color and Imaging Conference*, Darmstadt, Germany, (2015)

[10] M. A. Fischler, R. C. Bolles "Random Sample Consensus: A Paradigm for Model Fitting with Applications to Image Analysis and Automated Cartography," *Comm. of the ACM*, Vol 24, pp 381-395, (1981)

Author Biography

Ján Morovič received his Ph.D. in color science from the University of Derby (UK) in 1998, where he then worked as a lecturer. Since 2003 he has been at Hewlett-Packard in Barcelona as a senior color scientist and later master technologist. He has also served as the director of CIE Division 8 on Image Technology and Wiley and Sons have published his 'Color Gamut Mapping' book. He is the author of over 100 papers and has filed 100+ US patents (26 granted).

Peter Morovič received his Ph.D. in computing sciences from the University of East Anglia, UK in 2002 and holds a B.Sc. degree *in theoretical computer science from Comenius University, Slovakia. He has been a Color and Imaging Scientist at HP Inc. since 2007, has published 50+ scientific articles and has 100+ US patents filed (25 granted) to date. His interests include 2D/3D image processing, color vision, computational photography, computational geometry and his Erdős number is 4.*