https://doi.org/10.2352/issn.2169-2629.2019.27.33

A psychovisual study of print smoothness via metameric samples

Sergio Etchebehere, Peter Morovič and Ján Morovič
HP Inc, Barcelona, Spain

Abstract

The smoothness of a print is one of its main image quality attributes. Here smoothness can refer to the level of unexpected changes or discontinuities in color transitions (at a macro scale) or the level of local variation (at a micro scale), sometimes also described as grain. This paper starts with presenting an approach to building a first-ever set of metameric printed samples that match in color but vary in grain, followed by a psychovisual study of smoothness perception based on a large number of evaluations by both experts and non-experts. This data shows high levels of intra- and inter-observer correlation and can therefore serve as a robust ground truth for understanding and modelling the print smoothness phenomenon. Then, a previously published predictive smoothness model is revisited, that estimates smoothness from a digital halftone before it is printed, and it is shown to result in high degrees of correlation between observer assigned smoothness judgments and computationally predicted scores. The paper also reports the results of tuning the smoothness metrics parameters to further enhance is alignment with the psychovisual ground truth.

Introduction

In general, there are a number of key considerations in terms of the goodness of a print, which include its naturalness or impactfulness, its color gamut, its contrast, the level of detail it shows, the continuity and gradualness of color transitions and the smoothness of its halftone patterns. Where such smoothness is lacking, the result can be described as grain, mottle or a lack of purity and while for some applications of color printing the expectations here are very high, for others there is more leniency towards imperfection. In all cases though, being able to assess and adjust the level of smoothness is important since it allows for choice when it comes to how that print is made, both in terms of the speed of printing and the quantity of ink used.

A lack of smoothness, or presence of grain, is considered to be one of the artifacts that most affect the image quality (IQ) of a print, together with mottle, banding, streaks and a variety of nonuniformities that may follow from the specific nature of a printing system (e.g., from the air-flows in a non-impact system to cylinder eccentricities and alignment variations in an impact one). When it comes to grain, its key contributors are the color separations and halftoning, which determine the digital pattern to be printed. Dot size variation and placement, as well as printhead alignments and other printing system variations, also play an important role, since they affect how that pattern is deposited onto a substrate and both the physics of the deposition process and the chemistry of the ink and substrate materials then impact the final prints properties. Even a well-chosen pattern that under ideal conditions is smooth, may result in a grainy print if large variations exist in a printing system. And, conversely, a pattern that, when viewed on a display, looks less smooth, may be so in ways that counterbalance the printing systems deviations and result in a smoother looking print.

Significant research has gone into finding a good way of estimating the grain of a given print, with a broad variety of approaches, most being based on capturing prints with imaging systems (e.g., by scanning samples). While this is of great value when it comes to evaluating grain of existing, printed samples, its applicability is limited when it comes to predicting the grain of digital patterns before it is printed and a physical sample becomes available. Existing approaches are therefore often not suitable for, e.g., making choices about alternative patterns before printing. A further challenge is that studying grain using psychovisual methods requires both a rich dataset that covers as many levels and types of grain as possible and the ability to use it in an extended psychophysical experiment.

This paper first takes advantage of the control offered by the Halftone Area Neugebauer Separation (HANS)[1] domain, which allows to construct Neugebauer Primary (NP) area coverages as probabilities of ink-combinations, followed by a deterministic halftoning approach [2] which can be used to construct a rich dataset that varies in grain [3]. Prints of such grain-varying halftone patterns are then used to get a rich psychovisual data set from a large number of subjects. Finally, it corresponds to existing predictive metrics is evaluated and used it to further optimize their performance.

Related Work

A lot of work has already been carried out into model and predict printed grain with significant advances. Here most approaches are based on some variant of the standard deviation of colorimetries within a specific neighborhood, while others revolve around spectral power analysis. ISO/IEC 13660 was standardized in 2001 [4] as a common way of defining and computing mottle and grain separately. However, there are also studies that show a lack of correlation with the human perception of grain for this metric [5]. There, the authors propose a modification of the ISO/IEC metric based on a previous wavelet filtering approach, that resulted in better correlation with human perception for both grain and mottle, becoming the new ISO/IEC TS 24790 standard [6]. Other studies tried using a wavelet-based approach [7, 8] while yet others based grain prediction on frequency power spectrum analysis /[9, 10], used more colorimetric approaches based on S-CIELAB [11] or involved halftone textures [12].

All these methods have one common attribute, which is that they require for a print to exist already and for this print to be captured using an colorimetrically-characterised imaging system such as a scanner or a camera. In a previous paper [3] a method was introduced that allows for grain prediction prior to printing. That method starts by taking an existing halftone, a spatial distribution of NPs over the area of a patch to be evaluated. Then, having the ability to predict the colorimetry of each of the NPs [13], the information from which grain can be predicted becomes accessible. That grain metric, which showed promising results, was based on estimating the halftone colorimetric values of a given patch, apply a blurring filter and then computing the resulting bitmaps standard deviation.

In all cases, however, the correlation of any given metric with human perception is crucial, and in turn, highly depends on having a good ground truth data set. Both previously described standards were validated on a dataset of 10 printers that printed 10 different patches varying only in L*. Later, these methods were also validated with 26 samples provided by the ISO/IEC TS 24790 committee that covered a wider range of grain artifacts [14]. The aforementioned method [3] was validated over 16 NPac patches inspected by 6 observers, which is a very small data set. Looking at other previous methods, some use 7 patches for each of the Black, Magenta, and Cyan channels [8], 11 levels for each ink channel [15], and some use Monte Carlo simulation in order to test their model accuracy [7]. Yet other studies entirely lack psychophysical evaluation [9, 10].

These limitations in the studies that correlate grain predicting metrics with psychophysical data might be related to the absence of a rich grain sample data set, on which experiments can be run, and on the difficulty of building such a data set in a systematic and scalable way. Compared with other areas of imaging research, such as the study of image quality, there are, e.g., many and rich image data sets [16, 17, 18, 19]. This might be due to the complexity of controlling the halftone domain and therefore the difficulty of generating a good variety of samples. Printing patches that have different levels of grain is not a trivial task, and there is either full control of ink drop placement or the need to simulate grain by adding external noise patterns. The following section provides details of how dataset was constructed that varies in grain by controlling halftone level properties.

Experiment

In this section, the experiment aimed at identifying grain perception will be introduced: the generation of the dataset, the set-up of the psychovisual experiment, methodology, and analysis of the data, a brief review of the grain metric and an analysis of how it correlates with the perceptual grading as well as work on tuning its parameters to improve the predictions will be shown.

A Metameric Dataset

In order to design a test for studying the perception of print smoothness, a data set in the form of samples that vary in levels of grain is required, even before having a robust way to estimate grain *a priori*. As described earlier, sources of grain can vary from those related to the content being printed (i.e. the halftones) to those that relate to printer behavior and stability. Since it is much harder to have controlled conditions of printer behavior (i.e. how well drops are placed, aerodynamic effects, etc.) the approach taken here is that of designing halftone patches that sample levels of grain as much as possible.

A halftones overall color and image quality will depend both on the composition of its atomic states the Neugebauer Primaries as well as how these are distributed. The HANS (Halftone Area Neugebauer Separation) print control paradigm exercises control precisely in this domain: which NPs to use (and therefore which inks to use), at what proportions to have them combine over a unit area or patch and how to distribute them spatially. The first two choices relate to how NP area coverages (NPacs) are designed, while the last contribution is dictated by halftoning parameters.

Furthermore, the dataset should also be such that dependencies on the overall color can be detected or discounted, hence the samples should also span a variety of colorimetries but at the same time vary at those different colorimetries. Hence metamers samples that match in color under the reference conditions but differ in other aspects under other conditions are required.

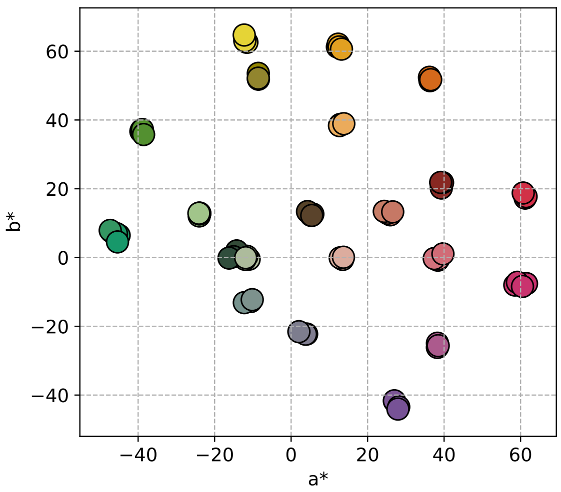

Figure 1. Distribution of grain data set color anchor-points after printing and measurement: near-metameric samples.

Three strategies will be used to generate the metameric dataset: an initial and partial grain optimization to find two sets of NPacs that differ significantly in grain, a second order, per-NPac optimization to further tune attributes known to affect grain and finally interpolation between samples of differing grain but matching color, to build intermediate cases. The test set-up on which these computations are done uses cyan, magenta, yellow and black inks on an HP DesignJet Z3200 using glossy media. On a 4-ink system that, in this case, can deposit 0, 1 or 2 drops (three states) of each ink, the domain of all the possible NPs results in a $3^4 = 81$-dimensional space. It is this domain over which the following process operates.

The starting point for is a set of 20 sparsely distributed anchor points in CIE L*a*b* space (see Figure 1). For each of these colors, an initial grain optimization was computed according to a previously described approach [3]. This process involved generating NPacs (following different strategies, from the simplest tetrahedral ones, all the way to randomly combining NPs at random coverages) that result in colorimetries close to the anchor set (as predicted by a color model [6]) as well as computing their grain metric approximation. Once the NPacs, their colorimetries, and grain metrics have been computed, in each neighborhood (or bin), the best and worst grain NPacs are selected. Finally, the anchor colorimetries are interpolated in both sets in order to match them and the result are two sets of NPacs that are expected to vary in grain levels. This is the first level of processing and results in an initial dataset.

The second step is to compute intermediate steps between the two most and least grainy set of NPacs. These are simple convex combinations of the two data sets, that preserve colorimetry but transition in grain terms. Finally, the last approach to add yet more variety of grain is to apply per-NPac optimization order. The approach follows that described in [13] where a linear programming (LP) optimization is performed, computing new NPacs that use the same ink amounts, but distribute NPs in different ways (e.g. controlling the amount of white coverage, promoting inks placed side-by-side, such as C, M, CC, or trying to overprint them, such as CM, CY, CCM). These LP-based approaches act as proxies to the intuition of grain being related to the local contrast of NPs within a halftone, such that a halftone containing white (blank media) would be grainier, while halftones that use NPs whose colorimetries (and especially lightnesses) are closer among themselves.

Finally, all NPacs are halftoned using the Parallel Random Area Weighted Coverage Selection (PARAWACS) halftoning algorithm [2], where a single halftone matrix determines the spatial distribution of the NPacs. Here different halftone matrices can be used and since it is used as a spatial selector to determine where each NP will be placed, it determines the pattern of the halftone and is therefore closely related with grain as well; for this data set the focus is on a default blue-noise matrix that for appropriate NPacs results in low grain, as well as some samples with a more clustered, green-noise one (which produces grainer prints) and control set using a white noise matrix (that makes prints very grainy, regardless of the NPacs).

While the above process can generate many different metameric samples (different NPacs) for a given colorimetry, a pre-selection was done in order to reduce the number of samples to those that are at least barely distinguishable.

The result of this process is then a dataset of 80 patches, having colorimetric values centered around the 20 different locations of the CIE L*a*b* space (Figure 1) and where each of these locations has between 3 to 5 patches differing in grain. The colorimetric difference between the patches considered to have the same color were on the average of 0.61 CIE ΔE 2000, with the 95^{th} percentile of 1.65 for CIE illuminant D50. This means that the data set is composed of having a set of 20 different printed colors, such that for each of the colors there is a set of near-metameric samples, closely matching in color but varying in terms of the grain a unique feature of this data set.

Figure 2 shows an example of three different metameric patches, that have a significantly different NPac structure: the one in the top is composed of non-overlapping inks, mainly NPs of Y, C, and M by themselves, the patch in the middle overprints CM, resulting in lower contrast of the NPs within the patch, therefore its perceived grain should be reduced. Finally, the patch at the bottom is made up of 68 different NPs and has a white (blank) area coverage of almost 50% of the patch, leading to having high contrast with white, creating a higher grain perception.

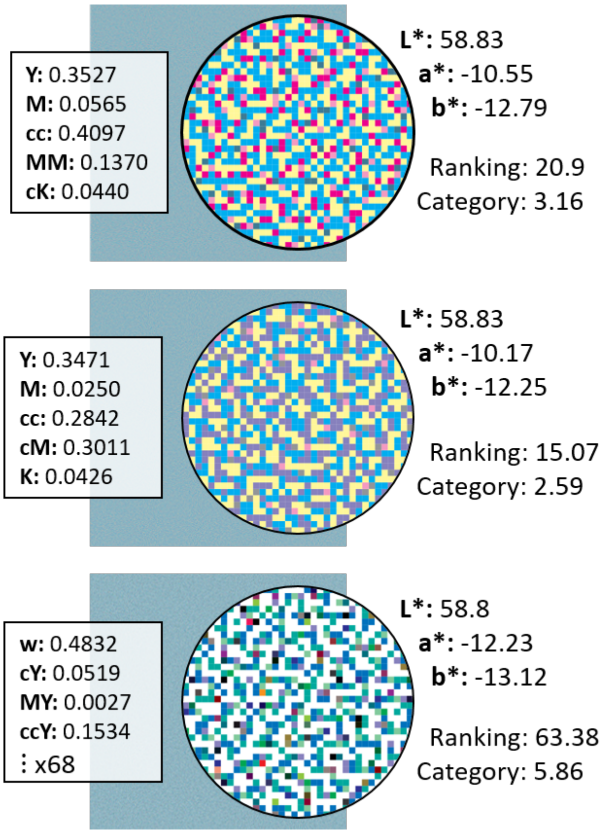

Figure 2. *Three examples of NPacs from the process described in this section showing both a variety of NPacs and a variety of grain while being near-metameric in colorimetric terms.*

Psychophysical Experiment

Set-up

The 80 samples of the dataset, defined following the process described in the previous section, were halftoned at 1200 dpi to result in patches of 1 inch square and were cut to show no blank substrate around its edges. During the experiment, the patches were placed on a neutral gray surface, under controlled lighting conditions using a CIE illuminant D50 simulator in a VeriVide viewing booth. The observers were presented with the patches at a starting position of approximately 50 cm viewing distance, meaning that the patch occupied approximately 3^o of visual angle. However, they were free to inspect patches from closer-up. The observers did not have any other visual cues in the environment except for a white patch of blank substrate that was used to facilitate complete adaptation. A total of 20 observers participated in the experiment, 15 of which were labeled as experts at evaluating printing attributes and artifacts (engineers and scientists who routinely perform visual IQ evaluations), with the remainder considered non-experts. All participants showed normal color vision when tested with Ishiharas colorblindness test and normal or corrected-to-normal visual acuity.

Methodology

For each participant, the 80 patches were randomly separated into 4 groups of 20 patches each. Each participant was asked to perform the experiment 6 times, with two groups of patches being explored together and with each viewing session being one of 6 possible combinations of the 4 groups. This resulted in each viewing session involving a set of 40 patches to be evaluated by an observer at once.

Observers were asked to perform two tasks: first, to sort all patches from least to most grainy, with the possibility of judging two or more patches as equally grainy. Second, to rank each patch into a grain category: from 1 (no grain) to 7 (most grain). In order to have an anchor patch, at each viewing session, a patch with no grain (100% coverage of a single ink) was included. The time taken for the observers to complete the task of one viewing session was approximately 15 minutes. The 6 viewing sessions were split over two days (3 + 3 sessions each) and a 5 minutes break was taken in-between sessions of the same day. Figure 3 shows an example of the experimental set-up, methodology and reporting of the participants. This methodology, for 20 observers and 80 patches, resulted in 20x6x40 = 4800 grain judgments in total.

$$RankingOrder = \frac{\sum_{i=1}^{N} S_i}{N} + \sum_{i=1}^{N} D_i \quad (1)$$

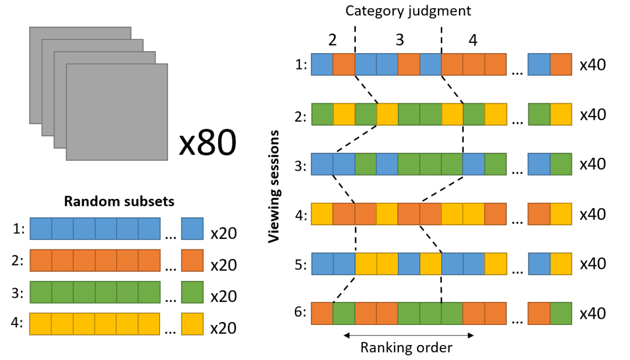

Figure 3. Diagram showing the experimental methodology; the distribution of patches into four random subsets and all combinations explored at each viewing session. Also shown is an example of the reporting of the participants (category judgment and ranking order).

Table 1: Summary statistics

	Ranking Order		Cat. Judgment	
	mean	95^{th}	mean	95^{th}
Inter-observer corr.	0.87	0.75	0.82	0.66
- *Experts only*	0.87	0.79	0.82	0.71
Intra-observer std	-	-	0.45	0.94
- *Experts only*	-	-	0.45	0.94
Patches std	7.26	11.57	0.77	1.17
- *Experts only*	6.79	12.87	0.73	1.30

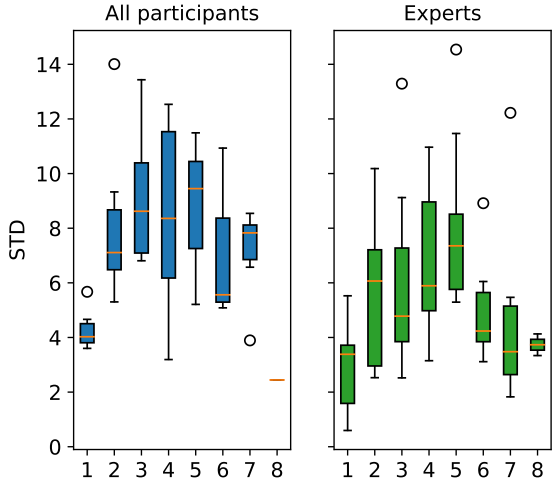

Figure 4. Boxplots of the standard deviation of the patches grouped by their average mean ranking order given by all the observers (left) vs only the experts (right). Group 1 goes from values of 0 to 10, group 2 10-20, until group 8 70-80. Boxes show the upper and lower quartile, an orange line shows the median, whiskers extend to show the range of data within $1.5 \times IQR$, rest marked as points are considered outliers.

After all viewing sessions were concluded, the overall, absolute ranking order could be obtained from the sub-set orderings as shown in equation 1, where S_i is the relative order of a patch within the set of patches of the same group in the i-th viewing session, D_i is the order relative to the samples of different subsets and $N = 3$. Since every participant had to evaluate a grain category three times for each patch, intra-observer repeatability can also be evaluated.

Analysis of results

Two ways of looking at the data allows for different types of insights to be gained; category judgment is an absolute value that relates directly to an observers perception of grain, while ranking order has a greater discrimination of the samples, but is relative with respect to the sample set shown.

The Pearson correlation coefficients between the three different scores (the ranking, category judgment and combined ones) of an observer are in all case greater than 0.99, indicating a high degree of consistency between the responses to the three tasks. Table 1 then shows summary statistics of the data in terms of inter- and intra-observer behavior as well as the standard deviations of patches. Overall, there is good correlation between the different observers in each of the score types, as well as small standard deviations between repetitions of the same patch by the same observer; taking in consideration the range in ranking order is within 80 and in the category judgment within 7

When looking at the difference of the data from all the observers versus isolating the expert observers, there is no clear signal in intra- and inter-observer repeatability, but an improvement in the repetition of patches (average of standard deviarion of 6.79 vs 7.26 when looking at the ranking order and 0.73 vs 0.77 at the category judgment). Meaning expert subjects were more consistent with their choices, as would be expected.

When looking at the distribution of repeatability versus the average score, the decisions with the highest variation are clustered around the middle of the score range (see Figure 4, with much better repeatability for patches scoring lowest or highest (the left-most and right-most data points in both subplots). This effect can indicate an ambiguity related to different kinds of grain; while the most and least grainy patches are clearly identified as such by all observers, the patches that score in-between can present different types of grain, and depending on the observer the preference for different types may vary, a comment made by some observers during the experiments. Figure 4 also shows how experts have overall higher consistency among them, but maintain the same relative tendency of ambiguity for mid-level grain patches.

Grain Prediction

Given this rich psychovisual data, the next step is to evaluate how well a predictive metric can estimate it. In previous work, a grain metric was introduced that allowed to predict grain levels from digital input, without the need to first print and then scan corresponding patches [3], having as its inputs the halftone and a prior characterization of the printing system in terms of its NP colorimetries. The basic principle of the approach is to replace halftone states (NPs) with their colorimetries (i.e. at every location in the halftone, the NP is replaced by its Yule-Nielsen modified XYZ values), then a convolution of a given window-size is performed over the size of the halftone patch and the standard deviation is computed over the entire patchs pixels. This standard deviation of the NP colorimetries over the filtered patch is taken as the grain metric. The basic intuition behind this metric (as behind many others) is that the lower the standard deviation, the lower the color contrast among different NP colorimetries and hence the lower the perceived grain.

However, this metric has only been validated on a very limited data set, both in terms of the number of patches and the number of observers and observations. Furthermore, subjects were asked to perform a much simpler and coarser three-category judgment: to discriminate between low, medium and high levels of grain. Therefore, in this section, a validation of this metric will be performed on the new, richer and finer data set. An exploration will also be presented of how some modifications to the metric and tuning of its parameters can result in even closer agreement with the ground truth data.

Figure 5a shows the performance of the previously described metric and its default parameters over 66 of the patches that used the same halftone matrix. The method predicts the different grain levels with a Pearson correlation of 0.7132. Comparing the predictions against the scores obtained from the expert observers alone, the correlation is similar at 0.7454. However, when estimating the grain level of all the patches, including the ones that use different halftone matrices, the performance decreases significantly to 0.406. The importance of the role of the halftone matrix is therefore clearly substantial, however, since the objective of this metric is that of predicting grain from a digital halftone, parametrizing the metric with the halftone matrix is valid.

The grain metric has a number of parameters that can affect performance and whose values have previously been established based on the limited experimental data (and whose results are shown in Figure 5a). Three main aspects that have been explored to tune the method to obtain optimal results were:

- A key part of the algorithm is a blurring process performed over the halftoned Yule-Nielsen modified XYZs. This blurring window was set by default to be of size 3x3 as a uniform filter over which the average of all convoluted pixels is computed. A first improvement here is to change filtering from a uniform to a Gaussian filtering kernel instead, where its sigma (σ) value will be the controlled parameter, with a default (σ) of 1. This change makes the filtering more akin to that performed by the human visual system [20].
- Once the blurring is performed, the Yule-Nielsen modified XYZs are converted back to linear CIE XYZs from which the standard deviation is computed. When transforming into the Yule-Nielsen space a factor is needed which can vary depending on the level of optical dot gain, or as a result of data-fitting. The factor used previously in the grain metric was 4.
- The last variable explored is the size of the computed halftone matrix over which the entire computation is performed. Having the halftone bitmap from which grain is computed too small may leave it unrepresentative of larger patches, directly affecting accuracy; as size increases there is, however, a point of optimality, beyond which accuracy no longer improves and computation time increases significantly and unnecessarily. The default halftoning size is 64x64 pixels.

A full-factorial exploration was performed comparing the results of the computed metrics to the experimental dataset using Pearson correlation. In some cases, subsets of the dataset were used, for instance exploring the correlation with the patches that share the same halftone matrix. Furthermore, particular attention was paid to the behaviour of the model against the reported level of grain by the expert observers; due to its higher consistency and the consideration of being better able to distinguishing among different types of grain.

Figure 6 shows the performance over the variation of the blurring parameter (Gaussian kernel σ) using an already optimized Yule-Nielsen nonlinear factor of 5. The metric is compared against the ranking order given by the expert observer judgments of patches that use the same halftone matrix (solid line) and to the complete set of patches (dashed line). In the same of the patches that share the same halftone matrix, the maximum correlation is achieved at $\sigma = 0.9$ with a value of 0.7466; the correlation per patch of this set-up can be seen in Figure 5b. Interestingly, a second, albeit lower peak is found at $\sigma = 3.5$ with a correlation of 0.7275.

Table 2: Accuracy and computing time depending on the halftone patch size used. Time is shown in seconds taken to compute 80 patches.

Size	4	8	16	32	64	128	256
Corre	0.42	0.48	0.50	0.67	0.74	0.74	0.74
Time	0.09	0.1	0.15	0.31	0.99	3.42	14.19

When comparing the metric against all the patches from the dataset (mixing patches with different halftone matrices) a single global maximum is at $\sigma = 4.1$ with a only slightly lower correlation of 0.7322. Meanwhile when looking at ($\sigma = 0.9$), which was the best performing in the previous case, here, when using all the patches, the correlation drops significantly to 0.4675. Overall this reaffirms that having the halftoning strategy taken into account as a parameter is beneficial to the performance of the metric, although the metric can handle even using multiple halftone matrices, however with a significantly larger Gaussian kernel. In future work it will be useful to consider spatial frequency analysis here to understand the bimodal nature of data shown in Figure 5.

Similar behavior is seen when using the data of all the participants of the experiment (expert + non-expert), but with a lower overall correlation: the best performing case is found when using patches of the same halftone matrix at $\sigma = 0.9$ with a correlation of 0.7034.

The Yule-Nielsen non-linearity applied to the XYZ domain was found to give the best performance when set to a value of 5 in the maximum scoring case (experts looking at patches of the same halftone), but higher correlations can be found in other cases when changing this factor: when using the data reported from the experts on all the patches the best performing parameters were $\sigma = 4.1$ and Yule-Nielsen factor of 7 resulting in a correlation of 0.7412.

When it comes to the halftone patch size, the value that showed the best balance between accuracy and computational time was of 64x64 pixels (see Table 2). Sizes lower than 64x64 decreased performance in terms of correlation, and larger sizes show an exponential increase in performance time while not improving performance. This behavior has been found to be independent of any of the other factors.

After the study of the different results when tuning the variables of the method, we can conclude with a combination of parameters which bring optimal results to the dataset: $\sigma = 0.9$, YN factor $= 5$, patch size $= 64x64$.

Conclusions

The present paper brings three key contributions to the study and prediction of print smoothness, which are a scalable method for generating grain metamers, a rich ground-truth data set of psychovisual judgments made y 20 observers for a variety of grain levels and types, and an improved grain metric applicable to halftone patches before they are printed.

The grain metamer set generation technique used here re-

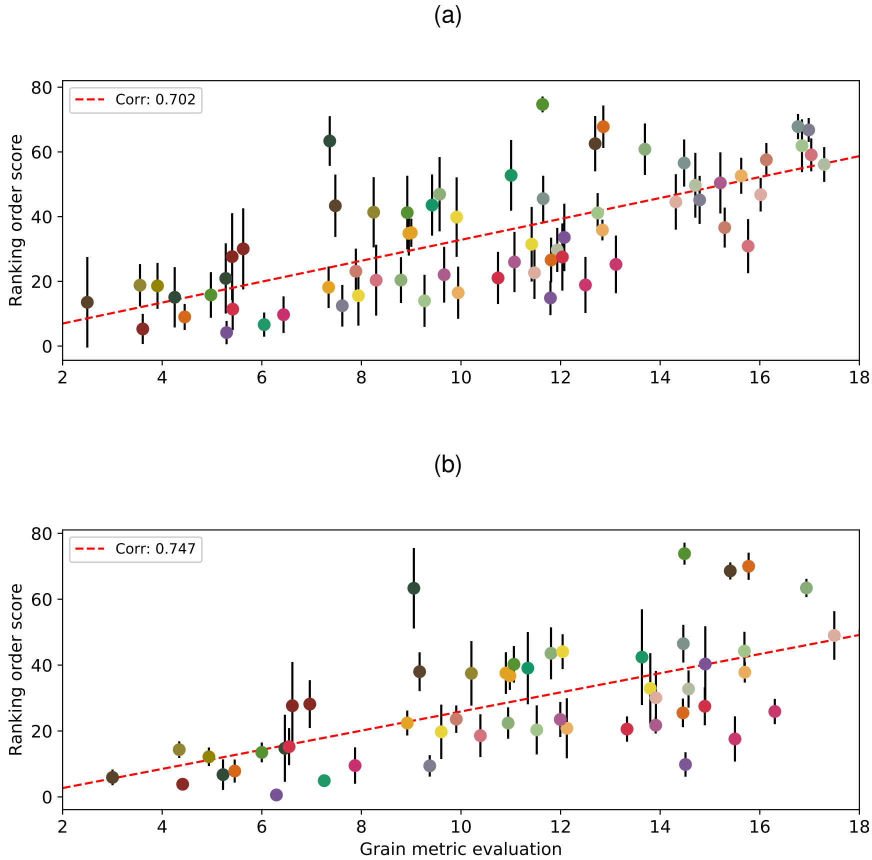

__Figure 5.__ Scores of the grain prediction metric with default parameters (a) and tuned ones (b) over the 66 patches that used the same halftone matrix against the ranking scores extrapolated from the psychophysical experiment from all the observers (a) as well as only the observers tagged as experts (b). Each circle shows a patch, with its color corresponding to the measured colorimetry of the patch. The error bars show the standard deviation between the participants and dashed line is the linear regression.

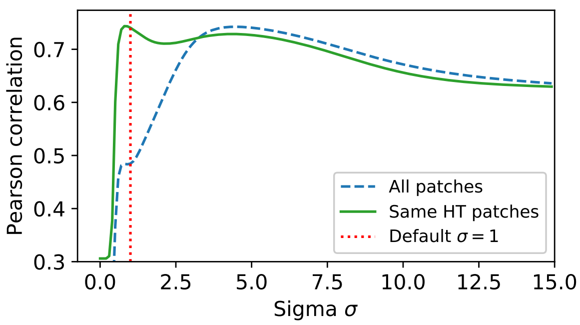

__Figure 6.__ Pearson correlation as a function of the Gaussian kernel sigma parameter. The results show the relationship for expert observers of the patches using the same halftone matrix (solid line) and all patches (dashed line) against the grain metric score; varying σ for a fixed Yule-Nielsen factor and patch size.

lied on the tight halftone-level control available via a HANS pipeline, while the characterization derived from the same HANS paradigm could in the future also be applied to halftone patterns generated using conventional colorant-channel based print imaging pipelines, since those halftones too can be characterized by substituting their NPs by NP colorimetries.

The newly-tuned grain metric has been shown to result in good correlation with psychovisual data, even when using the default parameters (0.7132). Focusing on the expert observers only, whose psychovisual data had a higher degree of consistency, as well as evolving and tuning the metric parameters, a somewhat higher correlation was achieved (0.7466). The metric also showed a good correlation even when applying it to patches that used different halftone matrices (0.7322).

The analysis and results shown in this paper also lead to further questions for investigation, such as the notion of grain preference. It is well known, for example in the digital photography world, that some grain (e.g. the grain of analogue film) is nicer than other grain (e.g. grain originating from digital noise). This could be one of the reasons why low and high grain are consistently identified, while mid-level grain shows greater variety. This would also be analogous to previous studies into gray neutrality preferences [21], where the balance implied by adaptation is often not what observers favor.

Acknowledgments

The authors would like to thank their colleagues at HP Inc. for their support, specially to Alessandro Beltrami, Alex Campa, Annarosa Multari, David Peinado, Hector Gomez, Javier Maestro, Jordi Bas, Jose Manuel Macrillante, Leyre Hernandez, Marc Casaldaliga, Michel Encrenaz, Montse Solano, Nigel Williams, Sonia Gallardo, Tanausu Ramirez, Utpal Sarkar, Victor Diego, Xavier Fariña, Xavier Quintero and Yuval Eckstein

References

[1] Ján Morovic, Peter Morovic, and Jordi Arnabat. Hans: Controlling ink-jet print attributes via neugebauer primary area coverages. *IEEE Transactions on Image Processing*, 21(2):688–696, 2012.

[2] Peter Morovič, Ján Morovič, Jay Gondek, and Robert Ulichney. Direct pattern control halftoning of neugebauer primaries. *IEEE Transactions on Image Processing*, 26(9):4404–4413, 2017.

[3] Ján Morovič and Peter Morovič. Hans print smoothness optimization and continuous control. In *Color and Imaging Conference*, volume 2017, pages 198–203. Society For Imaging Science & Technology, 2017.

[4] Ahmed Eid, Brian Cooper, and Ed Rippetoe. A unified framework for physical print quality. In *Image Quality and System Performance IV*, volume 6494, page 64940C. International Society for Optics and Photonics, 2007.

[5] Ahmed H Eid, Brian E Cooper, and Edward E Rippetoe. Characterization of mottle and low-frequency print defects. In *Image Quality and System Performance V*, volume 6808, page 680809. International Society for Optics and Photonics, 2008.

[6] ISO. Information technology office equipment measurement of image quality attributes for hardcopy output monochrome text and graphic images. ISO ISO/IEC 24790:2017, International Organization for Standardization, Geneva, Switzerland, 2017.

[7] Kevin D Donohue, Chengwu Cui, and M Vijay Venkatesh. Wavelet analysis of print defects. In *IS&Ts PICS Conference*, pages 42–47. Society For Imaging Science & Technology, 2002.

[8] Ki-Youn Lee, Yousun Bang, and Heui-Keun Choh. Characterization of'2d noise'print defect. In *Image Quality and System Performance VI*, volume 7242, page 72420M. International Society for Optics and Photonics, 2009.

[9] Paul J Kane, Theodore F Bouk, Peter D Burns, and Andrew D Thompson. Quantification of banding, streaking and grain in flat field images. In *IS&Ts PICS Conference*, pages 79–83. Society For Imaging Science & Technology, 2000.

[10] Bimal Mishra and Rene Rasmussen. Micro-uniformity: an image quality metric for measuring noise. In *IS&Ts PICS Conference*, pages 75–78. Society For Imaging Science & Technology, 2000.

[11] Dirk W Hertel. Exploring s-cielab as a scanner metric for print uniformity. In *Image Quality and System Performance II*, volume 5668, pages 51–61. International Society for Optics and Photonics, 2005.

[12] Tetsuya Itoh and Kazuomi Sakatani. Noise evaluation metric derived from digital am halftone image analysis. *Journal of Imaging Science and Technology*, 43(2):113–119, 1999.

[13] Peter Morovič, Ján Morovič, Xavier Fariña, Pere Gasparin, Michel Encrenaz, and Jordi Arnabat. Spectral and color prediction for arbitrary halftone patterns: a drop-by-drop, wysiwyg,ink on display print preview. In *Color and Imaging Conference*, volume 23, pages 2–6. Society For Imaging Science & Technology, 2015.

[14] Ahmed H Eid, Brian E Cooper, and Edward E Rippetoe. On the analysis of wavelet-based approaches for print grain artifacts. In *Image Quality and System Performance X*, volume 8653, page 86530K. International Society for Optics and Photonics, 2013.

[15] Tetsuya Itoh. Color noise of various printer outputs. In *Color Imaging: Device-Independent Color, Color Hardcopy, and Graphic Arts VI*, volume 4300, pages 314–326. International Society for Optics and Photonics, 2000.

[16] Xinwei Liu, Marius Pedersen, and Jon Yngve Hardeberg. Cid: Iq–a new image quality database. In *International Conference on Image and Signal Processing*, pages 193–202. Springer, 2014.

[17] Yuukou Horita, Keiji Shibata, Yoshikazu Kawayoke, and ZM Parvez Sazzad. Mict image quality evaluation database. *Online], http://mict. eng. u-toyama. ac. jp/mictdb. html*, 2011.

[18] Nikolay Ponomarenko and K Egiazarian. Tampere image database 2008 tid2008, 2009.

[19] HR Sheikh. Live image quality assessment database release 2. *http://live. ece. utexas. edu/research/quality*, 2005.

[20] L. G. Shapiro and G. C. Stockman. *Computer Vision*. Prentice Hall, 2001.

[21] Raja Bala. What is the chrominance of gray? In *Color and Imaging Conference*, volume 9, pages 102–107. Society For Imaging Science & Technology, 2001.

Author Biography

Sergio Etchebehere received his M.Sc. in color science from the University Jean Monnet (France), University of Granada (Spain) and University of Eastern Finland (Finland) in 2016, holding a B.Sc in computer science from Univerity Pompeu Fabra (Spain). Since 2017 he works as color and imaging scientist at HP Inc in Barcelona. His interests and major research work are in human vision, color theory, computer vision and printing and imaging devices.

Peter Morovič received his Ph.D. in computer science from the University of East Anglia (UK) in 2002 and holds a B.Sc. in theoretical computer science from Comenius University (Slovakia). He has been a senior color and imaging scientist at HP Inc. since 2007, has published 50+ scientific articles and has 150+ US patents filed (54 granted) to date. His interests include 2D/3D image processing, color vision, computational photography, computational geometry and his Erds number is 4.

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 150+ US patents (57 granted).