Wove Paper Analysis through Texture Similarities

P. Abry*, A.G. Klein†, P. Messier‡, S. Roux*, M.H. Ellis§, W.A. Sethares¶, D. Picard‖, Y. Zhai**,
D.L. Neuhoff**, H. Wendt††, S. Jaffard‡‡, and C.R. Johnson, Jr.[x]

*Univ Lyon, Ens de Lyon, Univ Claude Bernard, CNRS, Laboratoire de Physique, F-69342 Lyon, France
†Dept. of Engineering and Design, Western Washington University, Bellingham, WA 98225, USA
‡IPCH Lens Media Lab, Yale University, West Haven, CT 06516, USA
§Institute of Fine Arts, New York University, 14 East 78th Street, New York, NY 10075
¶Dept. of Electrical and Computer Engineering, Univ. of Wisconsin, Madison, WI 53706, USA
‖ETIS, UMR 8051 / ENSEA, Université de Cergy-Pontoise, CNRS, F-95000, Cergy, France
**Dept. of Electrical Engineering and Comp. Sci., Univ. of Michigan, Ann Arbor, MI 48109, USA
††IRIT, CNRS UMR 5505, University of Toulouse, France
‡‡Univ Paris Est, Lab. d'Analyse et de Mathématiques Appliquées, CNRS UMR 8050, UPEC, Créteil, France
[x]School of Electrical and Computer Engineering, Cornell University, Ithaca, NY 14850, USA

Abstract—**Wove paper, made on a papermaking screen or mold having a surface of smooth tightly woven wires, was the predominant paper type used for printing in the twentieth century. To aid in the study and classification of fine art prints on wove paper, the present work compares the results of five different image processing approaches for characterizing surface texture. Using a collection of popular and widely available wove papers, a reference dataset of raking light close-up images was assembled. Five research teams undertook their own processing strategies to detect outliers and affinities among the paper samples. Their success in identifying similarity groupings built into the dataset are reported.**

I. Introduction

The study of graphic art relies upon easily identifiable and describable characteristics of paper. One such marker has been the watermark, which designates the paper's manufacturer and provides clues for its dating, original dimensions, function, and country of origin. Watermarks have been present in paper for centuries. In addition, papers are identified by their color, thickness, structure or formation, sheen, surface texture or finish, and other visual and physical properties. These properties, however, cannot be used to confirm that the papers are from the same papermaking mill or belong to a particular brand or type from that manufacturer.

Until the widespread adoption of the papermaking machine in the early nineteenth century, paper was made by scooping up finely macerated pulp and water from a vat using a rectangular mold comprised of a porous screen surrounded by a removable wooden frame. Prior to 1750, the screen was fabricated from fine, densely spaced horizontal rows of laid wires lashed into position by thicker, more widely spaced vertical chain wires. When the mold was plunged into the vat and lifted out, the wires acted as a sieve, filtering out the pulp in thinner and thicker accumulations depending upon how much interference

the wires produced as the water drained through [1]. The grid-like pattern of crisscrossed chain and laid lines is thus replicated in the structure or formation of the final sheet of paper and is also replicated in its surface texture or finish. Paper having this formation is called laid paper, and the pattern left by the wire molds has been used recently to classify these papers [2]. After 1750, a smooth-surfaced paper was developed by using a woven screen also in a mold surrounded by a removable frame [3]. Reconfigured as an endless belt, the woven wire mesh was readily adopted for machine-made papers starting in the early 19th century [4]. This paper, called wove, eventually superseded laid paper. Its formation lacks the regular grid pattern characteristic of laid paper; the felt-like distribution of the paper pulp across the sheet is even and amorphous. The surface texture or finish of wove paper is likewise continuously smooth. The lack of unique and quantifiable chain line intervals and laid line density make the characterization of wove papers difficult.

Modern wove papers are identified by their proprietary watermarks. In many prints, however, a partial sheet was used or the sheet was trimmed down, paring off the watermark, which, by the 20th century, had been relegated to the edges of the sheet. Print connoisseurs, however, recognize that even the most nondescript wove papers display unique surface finishes. These textures vary not only from manufacturer to manufacturer and type to type, but also between both sides of the same sheet of paper - its front or felt side, here called the recto, and its back or wire side, here called the verso, as shown in Fig. 1. Subtle differences in patterns can be discerned and recorded using a raking light [5]. Due to the complexity of the topography, however, and the variable orientations of light, it is impossible to match the pattern by eye.

It was wondered if the application of computer-based, image processing tools to mark, measure, and compare the unique finishes of each wove paper, front and back, as recorded in raking light, could be used to identify papers from the
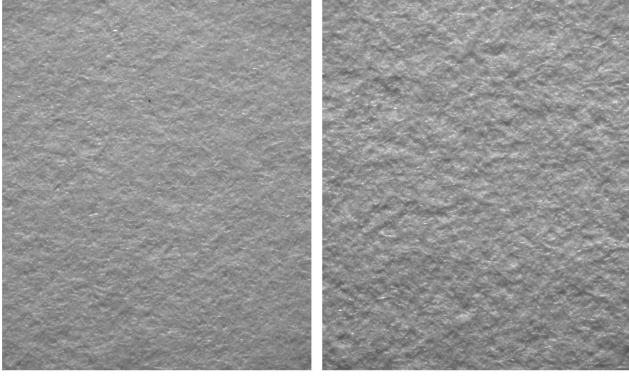
Fig. 1. Example raking light image of wove paper, showing recto (left) and verso (right)

same manufacturer. As part of the Historic Photographic Paper Classification Challenge [5], [6], a multitude of different approaches to texture similarity have been developed [5], [7]–[11]. These approaches were shown to yield encouraging results when used on silver gelatin and inkjet images [12]; here, we extend those prior results by reporting on the use of these texture similarity approaches on wove paper.

II. DATA SET

The paper samples selected for the data set are from *Specimens* [13], a 1953 publication of the Stevens-Nelson Paper Corporation. The samples are all of wove formation. Each sample was either hand-made, using an individually-dipped mold covered with a wire cloth, or mold-made, manufactured by machine using a small, mechanically driven, cylindrical mold. The surface texture or finish of each sample is not embossed or otherwise manipulated after manufacture and, thus, mimics the particular woven screen pattern favored by each manufacturer for each type of paper. The texture on the front and the back of each sample differs and are identified according to their presentation in the *Specimens* catalog.

Specimens, produced in an edition of over 5,000 copies, is an important reference today for graphic art curators, art historians, and paper conservators. Leading European and American artists, including Picasso, Matisse, Dubuffet, Lichtenstein, and Motherwell, or frequently their printers, selected printing papers from popular paper manufacturers, most of whom are represented in *Specimens*. The papers chosen to create printed works of art needed to have visual properties best suited to the printing technique at hand, i.e., silk screen, lithography, letter press, etc. Foremost among desirable characteristics was a smooth, continuous, and non-distracting surface. The samples from *Specimens* represent the wove papers often encountered in the study of mid-20th century graphics, a time of enormous changes in the genre of fine art prints.

Close up images of the surfaces of selected papers in *Specimens* were acquired while using a raking light, which is a linear light source at an oblique angle to the surface which enhances the highlights and shadows so that surface features are more clearly visible during image capture. Each sample

consists of 1×1.3 cm^2, scanned at a resolution corresponding to $6.51^2 = 42.4 \mu$m^2 per pixel.

The dataset consists of 180 close up images (90 recto, 90 verso) drawn from 36 different papers, and from 10 unique copies of *Specimens* to account for manufacturer variation. The dataset contains three levels of similarity: (1) samples from one same paper (3 subsets of 10 samples, labeled from 1 to 30, both recto and verso), (2) samples from identical sheets but different copies of *Specimens* (3 subsets of 10 samples, labeled from 31 to 60, both recto and verso); (3) and 30 papers (labeled from 61 to 90) of interest to paper conservators representing the diversity of wove papers (both recto and verso).

III. TEXTURE CHARACTERIZATION TOOLS

As they were fully described elsewhere [7]–[11], we only provide here a qualitative description of the five texture characterization image processing tools, emphasizing features and distances they rely on.

A. Anisotropic Multiscale Analysis (AMA)

Anisotropic multiscale analysis (AMA) [14] has been proposed in the context of the analysis of scale-free (or scale invariant) textures. It relies on the use of the Hyperbolic Wavelet Transform (HWT) [15]. The HWT consists of a variation of the 2D-Discrete Wavelet Transform (2D-DWT) [16], that explicitly takes into account the possible anisotropic nature of image textures. Indeed, instead of relying on a single dilation factor a used along both directions of the image (as is the case for the 2D-DWT), HWT relies on the use of two independent factors $a_1 = 2^{j_1}$ and $a_2 = 2^{j_2}$ along directions the horizontal (x_1) and vertical (x_2) directions. The HWT coefficients of imaged paper i are defined as inner products against wavelet templates, dilated with horizontal and vertical factors $a_1; a_2$ and translated at location k_1, k_2: $T_i((a_1, a_2), (k_1, k_2)) = \langle i(x_1, x_2), \frac{1}{\sqrt{a_1 a_2}}\psi(\frac{x_1-k_1}{a_1}, \frac{x_2-k_2}{a_2})\rangle$. Structure functions, consisting of space averages of the $T_i((a_1, a_2), (k_1, k_2))$ at scales a_1, a_2, are computed: $S_i((a_1, a_2), q) = \frac{1}{n_a}\sum_{\underline{k}}|T_i((a_1, a_2), (k_1, k_2))|^q$, with n_a the number of $T_i((a_1, a_2), (k_1, k_2))$ actually computed. To ensure that features do not depend on image intensity and that all scales contribute to texture characterization, the features consist of log-transformed normalized structure functions $\tilde{S}_i(a, q) = \ln \frac{S_i(a,q)}{\sum_{a'} S_i(a',q)}$. We use here $q = 2$ and a vector of seven dyadic scales $a = 2^l$, ranging from 2 pixels (6.51μm) to 2^7 (834μm), for a total of $7 \times 7 = 49$ features $\tilde{S}_i(a, q)$.

To measure proximity between two images i and j, a L^p norm cepstral-like distance is computed: $D(i, j) = \left(\sum_a |\tilde{S}_i(a, q) - \tilde{S}_j(a, q)|^p\right)^{\frac{1}{p}}$. We use here $p = 1$.

B. Pseudo-area-scale analysis (PASA) [8]

The PASA approach uses fractal analysis to decompose a surface into a patchwork of triangles of a given size. As the size of the triangles is increased, smaller surface features become less resolvable and the 'relative area' of the surface decreases. The topological similarity of two surfaces

is computed by comparing relative areas at various scales. Though wove samples do not provide a direct measure for height, light intensity is used as a proxy for height.

PASA first extracts a square $N \times N$ region from the center of the image (where N was chosen to be 1024), and normalizes the intensity of the resulting extracted image. The grid of N^2 equally spaced points (representing pixel locations) is decomposed into a patchwork of $2(\frac{N-1}{s})^2$ isosceles right triangles where s is a scale parameter representing the length of two legs of each triangle. The pixel values at each of the triangle vertices are then taken as the 'pseudo-height' of each of the vertices. The area of each triangle in 3-D space is then computed and the areas of all triangular regions are summed, resulting in the total relative area A_s at the chosen scale s, serving as features. a vector $\mathcal{S}$ of scales s ranging from 1 pixel to 34 pixels, (6.51 μm to 0.221 mm), for a total of 8 features. To assess the similarity of two images i and j, a χ^2 distance measure $d(i, j)$ is computed via

$$D(i, j) = \sum_{s \in \mathcal{S}} \frac{\left(A_s^{(i)} - A_s^{(j)}\right)^2}{A_s^{(i)} + A_s^{(j)}}.$$

C. Eigentexture Analysis (EGA) [9]

In the eigentexture approach (EGA), a collection of small patches are chosen from each samples. These patches are gathered into a large matrix and then simplified to retain only the most relevant eigendirections using a singular value decomposition (SVD) [17]. Features are extracted as follows: First, for each imaged paper j, N $p \times p$ pixel patches $X_{j,k} \in \mathbb{R}^{p \times p}$ are randomly picked (with $N = 2000$ and $p = 25$ in this case). The $X_{j,k}$ are lexicographically reordered into column vectors $a_{j,k} \in \mathbb{R}^{p^2}$ and stacked into matrices $A_j = [a_{j,1} \, a_{j,2} \, \cdots \, a_{j,N}]$. Second, to reduce dimensionality, SVD is applied $A_j = U_j \Sigma V_j$ for all j and the m columns of U_j corresponding to the m largest singular values, labeled $\boldsymbol{U}_j$, are retained (here, $m = 15$). The $M \times m = 3000$ features $\boldsymbol{U}_j$ are the representatives of the images and may be thought of as vectors pointing in the most-relevant directions. They concentrate on the analysis of patches of size p^2 pixels, corresponding to $162.75 \mu m^2$.

To compute similarity between images j and i, on first extract $Q = 2000$ $p \times p$ pixel patches Q_k from image i image and reorder them into vectors $q_k \in \mathbb{R}^{p^2}$. A distance from the kth patch to the jth image is computed as:

$$d(k, j) = ||q_k - \boldsymbol{U}_j(\boldsymbol{U}_j^\top q_k)||_2.$$

For each patch k of image i, one records the best match image $z_{k,i} = \arg\min_f d(k, f)$. Similarity between image i and j is then computed by a majority vote as the percentage of best matches from image i that belong to image j: $Dm(i, j) = (\#k \in \{1, \ldots, Q\}, |z_j(k) = i|)/Q$ and $D(i, j) = (Dm(i, j) + Dm(j, i))/2$.

D. Tensor aggregation of deep convolutional neural networks (CNN-VLAT) [10]

This method can be considered as a computer vision baseline. First, local features are extracted from image i as the output of a deep convolutional neural network (CNN) trained on natural images. The third convolutional layer of AlexNet [18] was used here, resulting into 13×13 localized features of size 384 each. PCA is further used to reduce dimensionality from 384 to 56, each component being potentially interpretable as a combination of non-linear filters detecting basic patterns such as angles, crosses or combinations of blob like patterns.

Second, we aggregate the local features into a single representation using the VLAT approach developed in [19], which has been shown to perform very well on similarity search tasks. The main idea is to cluster the local descriptors space and then to compute second order statistics for each cluster. Let $\{\mu_k\}_k$ be the set of cluster centers and $\{\mathcal{C}_k = \{x_i | k = argmin_j \|\mu_j - x_i\|\}\}_k$ the associated sets of local descriptors x_i belonging to each cluster. For each cluster, the VLAT representation is then the following matrix:

$$T_k = \sum_{x_i \in \mathcal{C}_k} (x_i - \mu_k)(x_i - \mu_k)^\top$$

Then, the matrices T_k for all k are flatten into vectors and concatenated into a vector $\mathbf{v}$. The final features $\mathbf{r}$ are obtained by a non linear transformation ensuring an ℓ_2 normalization:

$$\mathbf{r} = \frac{sign(\mathbf{v})\sqrt{|\mathbf{v}|}}{\|\mathbf{v}\|}$$

We used here 16 clusters, for a resulting dimension of $16 \times 56 \times 56 = 50176$. Distances between images i and j are computed by inner product between the features of each images: $D(i, j) = 2 - \langle \mathbf{r}_i, \mathbf{r}_j \rangle$.

E. Local radius index (LRI) [11]

A texture contains repetitive smooth regions, called *texture elements*, and transitions between them, i.e., *edges*. The Local Radius Index (LRI) [11], [20] aims to capture the sizes and shapes of the texture elements by considering the distances, at various angles, between the edges that surround them. More precisely, one version of LRI (LRI-A) is designed to capture inter-edge distance distributions in different directions by focusing on pixels adjacent to edges. The LRI-A feature consists of eight histograms, one for each direction corresponding to an adjacent pixel, where given a pre-specified threshold $T > 0$, the value m (or $-m$) of the n-th histogram is the frequency with which m successive pixels in direction n have intensity at least T larger (or smaller) than the current pixel. The value 0 of the n-th histogram equals the frequency with which the adjacent pixel in direction n differs by less than T. To limit histogram sizes, the frequencies for all values m greater than some upper limit K are lumped into the frequency for $m = K$, and similarly, for $m < -K$. Typically, $K = 4$ works well, and the threshold T is taken to be one half an image's standard deviation. Histograms of all 8 directions are concatenated into a feature vector h, hence of size $8 \times (2 \times 4 + 1) = 72$, analyzing essentially the fine scales of the image, from 1 to 9 pixels (i.e., from 6.51 to $58.59 \mu m$.

A texture similarity metric is obtained as the Kullback-Leibler divergence between the histogram feature vectors of images i and j: $D(i, j) = h_i \ln(h_i/h_j) + h_j \ln(h_j/h_i)$.

In the results reported in the next section, the LRI-A based metric is combined with metrics based on complementary features [11], [20], such as LBP [21] and measures of intensity in various frequency bands. Further, the LRI-A feature used here is the version in [11], previously found to be effective in photographic paper classification. This version is sensitive to rotations and is therefore not rotationally invariant. The principal improvement of this version over conventional LRI-A is that instead of keying the threshold T to an image's standard deviation, T is keyed to the empirical cumulative distribution function of adjacent pixel differences.

IV. Wove paper automated characterization

[●] **Context.** In the context of wove paper characterization, the goal is usually not to perform supervised learning classification as datasets often consists of unlabelled samples, that should be grouped in a priori unknown numbers of classes. The proposed image processing texture characterization techniques are thus rather essentially extracting features from wove paper samples and computing distances between features of two different samples, which could served as inputs for classical unsupervised clustering strategies, e.g., spectral clustering. Each of the 5 teams was provided a version of the dataset with scrambled labels, and no expert information, to preclude any form of supervised learning. Results were reordered a posteriori to ease comparisons and interpretation, as in Fig. 2. Further, the goal of the present contribution is not to compare performance achieved by each method but rather to assess whether or not different image processing texture characterization tools, relying on features very different in nature, and on distances based on metrics different in spirit, are able or not to successfully assess similarities amongst wove paper samples.

Expert assessment was constructed from experts involved in this study by in visual inspection of paper surfaces and using six qualitative levels of match from *perfect* to *very poor*. Expert assessment are used to evaluate both qualitatively and quantitatively the automated quantifications of distances between all pairs of images.

[●] **Qualitative description of performance.** Fig. 2 reports, in the form of matrices, distances computed between each pair of wove paper samples, independently for the recto and verso, and compares them to the expert assessment. Fig. 2 thus permits either visual qualitative comparisons of distances yield by different image processing texture characterization tools or comparisons against the expert assessment. Fig. 2 shows that most computed distances reproduce the 10×10 black squares along the main diagonal, clearly visible on expert assessment matrices, materializing the very low distances amongst the different samples from the same sheet, or between samples from different copies of the same paper. This holds for the analysis of both recto and verso. Further, intermediate gray levels quantifying proximity between samples 1-10 and 11-20, or 1-10 and 21-30, or 11-20 and 21-30, for the recto expert assessment matrix are also satisfactorily reproduced by several image processing tools. This also holds for proximity between samples 1-10 and 11-20 for the recto expert assessment matrix,

but much less clearly for proximity between samples 21-30 and 31-40.

[●] **Quantitative description of performance.** To further quantify similarity assessment performance, we adopt metrics used in the informational retrieval community to assess the performance of each approach. Such metrics are based not on the distances or affinities themselves, but on the *rank* of true matches when, for a given query image, all other images in the dataset are ordered by increasing distance (or decreasing affinity) to the query image. In particular, we consider three performance metrics: (i) *precision at one* (P@1) which is the mean fraction of time that the top ranked match (having smallest distance to the query image) is a true match, (ii) *mean reciprocal rank* (MRR) which measures the mean inverse rank of the first true match [22], and (iii) *mean average precision* (MAP) [23]. The MAP is calculated as follows: for each query image and positive integer n less than or equal to the size of the data set, compute the fraction of the n highest ranked images that are true matches, and then average these fractions over all values of n for which the nth highest ranked image was actually a true match; then, average these values across all images. The compared performance metrics are reported in Table I. Despite differences, Table I shows that most image

TABLE I
PERFORMANCE METRICS FOR EACH APPROACH

	recto			verso		
	P@1	MRR	MAP	P@1	MRR	MAP
AMA	98.3%	99.2%	95.1%	100.0%	100.0%	97.7%
PASA	95.0%	96.9%	73.6%	88.3%	92.2%	65.7%
EGA	43.3%	59.8%	44.7%	48.3%	62.4%	50.3%
CNN-VLT	68.3%	81.5%	64.3%	88.3%	92.9%	72.4%
LRI	95.0%	97.5%	94.4%	100.0%	100.0%	96.8%

processing approaches achieve satisfactory performance in comparing wove paper performance. In addition, the verso side appears easier to characterize as the performance is generally better compared to that achieved on the recto.

V. Conclusions and Perspectives

Reporting satisfactory performance, when compared to expert assessment, in the quantification of wove paper similarities for several different image processing texture characterization tools, this contribution demonstrates, as a proof-of-concept, that the computerized and automated assessment of similarities between wove papers can be achieved.

We believe that twentieth century wove paper popular for fine art prints display a variety of complex surface textures that can be characterized in several ways. In fact, the experts themselves admit to the difficulty in producing consistent ground-truth observations. Therefore, rather than comparing the performance achieved by each approach, we emphasize here that this automated assessment of similarity between wove papers can be achieved from tools very different in principle: they rely on features that are different in size, nature, and physical scale, as well as on different distances (L1-norm, L2-norm, Chi-square, Kullback Leibler). With such a diversity
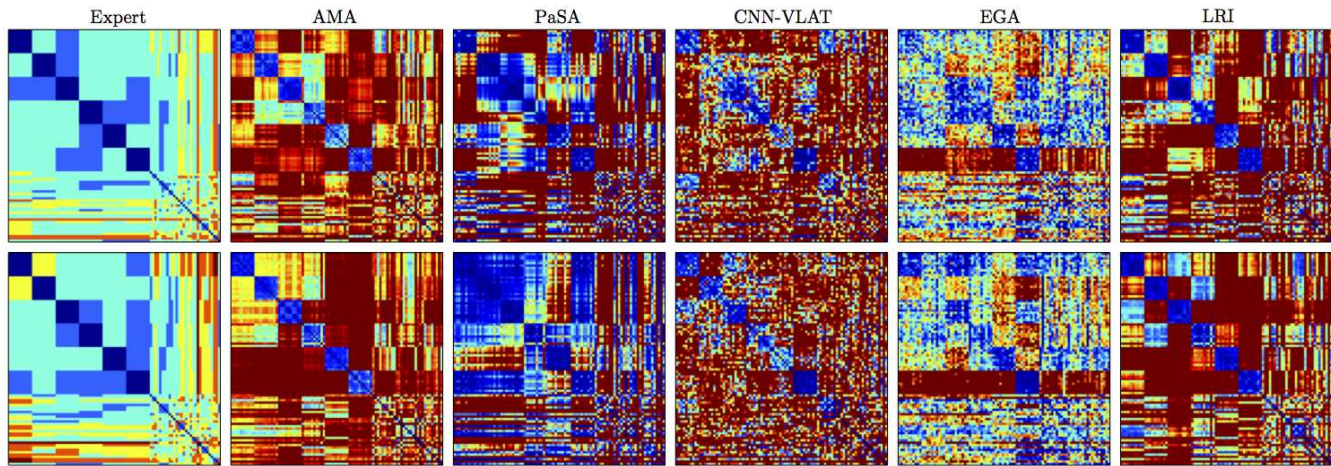
Fig. 2. Expert assessment (most left) compared to the distance matrices computed from the five different image processing texture characterization techniques. Distances range from dark blue (low) to red (high). Top: Recto; Bottom: Verso.

of differences, these different techniques can be used jointly to induce robustness in automated assessment. Indeed, several approaches consistently indicate similarities between samples that were not tagged as close in the expert assessment matrices. This may provide feedback to experts on the procedure to revise ground-truth, an involved issue from the expert's own experience. The ability to identify and differentiate wove paper in an automated manner based upon its surface texture will add immeasurably to print connoisseurship by objectively documenting an artist's preference and intent, or by revealing attributes of papers having specific functional and aesthetic aims. Further, the automated calculations of distances between all pairs of samples pave the road towards automated unsupervised clustering, e.g., using spectral clustering, our next goal. This will require addressing issues such as the status of recto and verso, so far analyzed independently, or, more classically, that of the optimal number of clusters. Interestingly, this latter issue can not be considered without interaction with wove paper experts, detailing explicitly the final goal of the clustering.

REFERENCES

[1] R. L. Hills, *Papermaking in Britain 1488-1988: A short history*. Athlone Press, 1988.

[2] C. R. Johnson, W. A. Sethares, M. H. Ellis, and S. Haqqi, "Hunting for paper moldmates among Rembrandt's prints: Chain-line pattern matching," *IEEE Signal Processing Magazine*, vol. 32, no. 4, pp. 28–37, July 2015.

[3] D. Hunter, *Papermaking: The history and technique of an ancient craft*. Dover publications, 1978.

[4] E. Sutermeister, *The Story of Papermaking*. SD Warren Co., 1954.

[5] C. R. Johnson, P. Messier, W. A. Sethares, A. G. Klein, C. Brown, A. H. Do, P. Klausmeyer, P. Abry, S. Jaffard, H. Wendt *et al.*, "Pursuing automated classification of historic photographic papers from raking light images," *Journal of the American Institute for Conservation*, vol. 53, no. 3, pp. 159–170, 2014.

[6] P. Messier. (2013) Paper Texture ID Challenge. [Online]. Available: http://www.papertextureid.org/about.html

[7] P. Abry, S. G. Roux, H. Wendt, P. Messier, A. G. Klein, N. Tremblay, P. Borgnat, S. Jaffard, B. Vedel, J. Coddington *et al.*, "Multiscale anisotropic texture analysis and classification of photographic prints: Art scholarship meets image processing algorithms," *IEEE Signal Processing Magazine*, vol. 32, no. 4, pp. 18–27, 2015.

[8] A. G. Klein, A. H. Do, C. A. Brown, and P. Klausmeyer, "Texture classification via area-scale analysis of raking light images," in *Proc. Asilomar Conf. on Signals, Systems and Computers*, Nov. 2014, pp. 1114–1118.

[9] W. A. Sethares, A. Ingle, T. Krč, and S. Wood, "Eigentextures: An SVD approach to automated paper classification," in *2014 48th Asilomar Conference on Signals, Systems and Computers*, Nov 2014, pp. 1109–1113.

[10] D. Picard and I. Fijalkow, "Second order model deviations of local Gabor features for texture classification," in *2014 48th Asilomar Conference on Signals, Systems and Computers*, Nov 2014, pp. 917–920.

[11] Y. Zhai and D. L. Neuhoff, "Photographic paper classification via local radius index metric," in *IEEE Intl. Conf. on Image Processing (ICIP)*, Sep. 2015, pp. 1439–1443.

[12] P. Messier, C. R. Johnson, H. Wilhelm, W. A. Sethares, A. G. Klein, P. Abry *et al.*, "Automated surface texture classification of inkjet and photographic media," in *Proc. Intl. Conf. on Digital Printing Technologies (NIP 29)*, Sep. 2013.

[13] *Specimens: A Stevens-Nelson Paper Catalogue*. New York: Stevens-Nelson Paper Corporation, 1953.

[14] S. G. Roux, M. Clausel, B. Vedel, S. Jaffard, and P. Abry, "Self-similar anisotropic texture analysis: The hyperbolic wavelet transform contribution," *IEEE Trans. Image Process.*, vol. 22, no. 11, pp. 4353–4363, 2013.

[15] R. DeVore, S. Konyagin, and V. Temlyakov, "Hyperbolic wavelet approximation," *Constructive Approximation*, vol. 14, pp. 1–26, 1998.

[16] S. Mallat, *A wavelet tour of signal processing*. Academic press, 1999.

[17] T. K. Moon and W. C. Stirling, *Mathematical Methods and Algorithms for Signal Processing*. Prentice Hall, 2000, pp. 138–141, 369–395.

[18] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "Imagenet classification with deep convolutional neural networks," in *Advances in neural information processing systems*, 2012, pp. 1097–1105.

[19] D. Picard and P.-H. Gosselin, "Efficient image signatures and similarities using tensor products of local descriptors," *Computer Vision and Image Understanding*, vol. 117, no. 6, pp. 680–687, 2013.

[20] Y. Zhai, D. L. Neuhoff, and T. N. Pappas, "Local radius index - a new texture similarity feature," in *IEEE Intl. Conf. on Acoustics, Speech, and Signal Processing (ICASSP)*, May 2013, pp. 1434–1438.

[21] T. Ojala, M. Pietikainen, and T. Maenpaa, "Multiresolution gray-scale and rotation invariant texture classification with local binary patterns," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 24, pp. 971–987, Jul. 2002.

[22] E. M. Voorhees, "The TREC-8 question answering track report," in *Proc. TREC-8*, 1999.

[23] ——, "Variations in relevance judgments and the measurement of retrieval effectiveness," *Inf. Process. Manage.*, vol. 36, no. 5, pp. 697–716, Sep. 2000.