Opinion

Statistics, vision, and the analysis of artistic style

Daniel J. Graham,[1] James M. Hughes,[2] Helmut Leder[1] and Daniel N. Rockmore[2,3*]

In the field of literature, there is an established set of techniques that have been successfully leveraged in the analysis of literary style, most often to answer questions of authenticity and attribution. With the digitization of huge troves of art images come significant opportunities for the development of statistical techniques for the analysis of artistic style. In this article, we suggest that the progress made and statistical techniques developed in understanding visual processing as it relates to natural scenes can serve as a useful model and inspiration for visual stylometric analysis. *WIREs Comp Stat* 2011 0 000–000

Keywords: stylometry; sparse coding; classification; logistic regression; image processing

AQ1

INTRODUCTION•

In 2011, a painting called 'Salvator Mundi,' purportedly by Leonardo Da Vinci, surfaced under mysterious circumstances in New York. A painting matching its description was known to be part of Leonardo's œuvre, though at least 20 copies of the original work by Leonardo's students and other imitators were also produced, some of which were in the past claimed by their owners to be the original Da Vinci. Without a secure provenance—or the record of the chain of ownership of works from their creation until the present—the 'determination' of a work's authenticity is more akin to trying a court case than administering a paternity test. While there is indeed a good amount of physical evidence (materials analysis of the pigments, canvas, and frame), at least as much weight comes from the opinions of experts, or 'connoisseurs,' who, assuming that the physical evidence agrees with their determination, act effectively as judge, jury, and even counsel, forming an opinion that is shaped by a lifetime of looking at and thinking about the works of the artist in question as well as the historical context in which these works were created.[1] Ultimately, they try to answer questions of the form, 'Is this work in the style of the artist?' or, 'Is the style of this work to be expected of the artist at this time in his or her career?' The stature of the connoisseur and the strength of his or her arguments make not only a significant historical impact, but a financial one as well.

These questions of style comparison are at least superficially ones that can easily be framed as questions of statistics. Given the known work of the artist and the period, how likely is it that the artist would create this kind of work at that specific time? At this stage, of course, nothing more has been done than to effectively translate the colloquial to the (vaguely) technical. The construction of a rigorous analysis requires objective measurements and a means of comparison performed in the service of articulating and understanding visual style. This is the goal of the nascent field of *visual stylometry*. With the increasing availability of large collections of high-resolution digital images of works of art, as well as new advances in image processing and machine learning and the understanding of the visual process, it is an area of research poised to make great progress.

The general thesis put forward here is that ultimately, style is a perceived phenomenon, and should be treated as such. In particular, our understanding of the human visual system and models of the way in which we 'see' natural scenes can serve as a useful means for visual stylometric analysis. In particular, the

AQ2

*Correspondence to: rockmore@cs.dartmouth.edu

[1]Faculty of Psychology, University• of Vienna

[2]Department of Computer Science, Dartmouth College

[3]Department of Mathematics, Dartmouth College

DOI: 10.1002/wics.197

two primary components of stylometry, selection and extraction of features and classification (or grouping) according to style, can benefit from this perspective. We show that going beyond mere metaphor and introducing statistical ideas from the understanding of human perception of natural scenes, while incorporating high-level perceptual information into models of style, proves effective in meaningfully organizing art images according to their style and in understanding some aspects of the perceptual basis of style.

VISUAL STYLOMETRY

The origins of visual stylometry can be traced back to the work of Giovanni Morelli (1816–1891), an Italian statesman who took it upon himself to both protect and rescue the reputations of his artist countrymen whom he felt were suffering from misattributions. Morelli was trained as both a paleontologist and a medical doctor and brought his scientific outlook and visual classification skills to bear on this difficult problem. The solution that he arrived at, called ‘scientific connoisseurship,’ was to compare specific details such as ears or hands in an unknown work to collection of exemplars (which he called a ‘schedule’), effectively asking whether the same kinds of details in the unknown work were consistent with the known exemplars. The choice of a less prominent feature like a hand or ear was a purposeful one, for Morelli believed that in these details the artist would be less driven by market and societal concerns and pressures.[2–4]

What Morelli did by eye and brain, researchers are now attempting to accomplish via statistical and image processing techniques. The problem remains a difficult one. Analogous efforts in the study of literature, where the notion of ‘stylometry’ was born, have been successful, and today there is a host of techniques that are used to answer questions of genre, authorship, and the dating of works—all of which are questions that we might hope to address in a quantifiable way in visual art as well.[5] Indeed, the following tasks are ones for which a combination of statistical and image processing tools might be able to provide a novel and compelling form of evidence that complements evidence acquired using more traditional techniques such as chemical and materials analysis, historical records, and, of course, human connoisseurship:

- Authentication: Given a set of ‘known’ works by a given artist, the task is to determine if an unknown work is or is not similar to the known works.
- Attribution: Given a work that could be by more than one artist, the task is to develop a model of each artist’s style to see which model best explains the work in question.
- Cultural evolution: Many questions in this area can be approached with image processing such as the influence of a given artist on contemporaries, students, and later artists, the influence of one artistic movement on a later movement, or the stylistic evolution of a particular type of work.
- Technical art history: Statistical research and simulation has produced important insights into artists’ methods for viewing, lighting, and capturing scenes on canvas.
- Conservation: Artwork from all eras—including the 20th century—shows degradation over time because of color fading, faulty previous attempts at conservation, air pollution, and other factors. Stylometric investigations can reveal important information used to conserve artworks and restore them to their original condition.

To date, progress in visual stylometry has been achieved in something of an *ad hoc* manner. In literature, or more generally, writing, we have the advantage of at least being able to identify a basic atom of relevance: the word. In visual art, we are less fortunate in that finding the basic elements of style is already a significant challenge. One might expect that standard image processing techniques such as SIFT[6] or color histogram statistics[7] would prove useful in the analysis of style in visual art, and indeed, many of these techniques have shown some success in sorting artworks with respect to general stylistic categories (e.g., broad stylistic groupings such as Impressionism, or sorting the works of famous artists by authorship). Nevertheless, many of these techniques often fail when confronted with the subtleties inherent in stylistic categorization.

For example, Gunsel and colleagues[8] work at the level of pixels and define six features derived from the pixel intensities in images of works of art. Using an •SVM classifier, the authors are able to separate Cubist, Classicist, and Impressionist works with good accuracy. AQ3 However, the addition of works in the Expressionist and Surrealist style causes performance to drop dramatically. While Cubist, Classicist, and Impressionist works may each be relatively distinct in terms of these low level features, it is clear that this approach does not scale up to finer grained distinctions. Other work[9] confirms the potential utility of these low level statistics for genre classification.

Some efforts attempt to aggregate pixel information at the level of the brushstroke. For example, the problem of brushstroke identification and extraction

AQ4

can be approached via filtering and subsequent modeling of lines, for example, using splines.[10,11] Sablatnig and colleagues[12] have developed sophisticated line 'skeleton' models for extracting brushstrokes from paintings. These kinds of models can also be used to extract strokes with respect to the order in which they were applied to the canvas. Subsequent processing and comparison can be accomplished by a variety of means including hidden Markov trees.[13–15] Unfortunately, it is often unclear how much the stylometric algorithms are measuring differences in the technical aspects of a painting (i.e., how the paint is applied to the canvas), and how much they are measuring differences in the visual 'space' represented in the painting (i.e., the depicted objects, scenes, and forms).

Several studies have examined artwork for regularities in the spatial statistical properties of the contours in art images. Perhaps the most well-known early example of work in the field of visual stylometry is that of Taylor,[16,17] who showed that the fractal dimension of works by Jackson Pollock, obtained using box-counting statistics, follows a very particular trend over the span of the artist's career. This work has not been without its critics.[18] Keren[19] suggested a model to describe local stylistic characteristics that uses the discrete cosine transform (DCT), along with a naïve Bayes classifier. The author's suggestion of a paradigm of 'recognition by type,' as opposed to recognition based on content, is a critical distinction in the application of image processing methods to the evaluation of artistic style, particularly in the context of image retrieval. Furthermore, Keren suggests that the obtained results are in accordance with human perception of style, and that his approach could be used to synthesize images of a particular style.[19]

In the work of Lyu et al.,[20] the authors applied spatial statistical techniques to the problem of distinguishing a set of authentic drawings by Pieter Bruegel the Elder from a set of imitation drawings using quadrature mirror filters to decompose the drawings, along with a hierarchical set of features derived from the filter coefficients, to demonstrate that the authentic drawings grouped tightly together in stylistic space. Work by Hughes et al.[21] demonstrated the technique of the Empirical Mode Decomposition (also known as the Hilbert–Huang transform)[22] to the same dataset. The authors showed that a newly created framework for a two-dimensional extension of EMD is capable of distinguishing between Bruegel and Bruegel-imitation drawings, as well as between a set of drawings by Rembrandt and his pupils. Other examples include the use of a combination of a wavelet decomposition and artificial neural networks capable of distinguishing between the works of Matisse and stylistic imitations of the artist's work.[23]

Other work focuses on nonspatially localized features or 'bags of features,' which are not ordered or arranged in any meaningful way. One example of this is the work of Berezhnoy et al.,[24] who analyzed the paintings of Van Gogh in terms of their use of complementary colors. They examined both the consistency of the artist's use of complementary colors, as well as the usefulness of this measure for organizing Van Gogh's works by the period in which they were painted. Another example is the work of Zujovic et al.[25] In this study, the authors used a collection of gray scale edge detection and Gabor wavelet pyramid sub-band statistics, along with color histogram statistics, to sort images into several stylistic categories. The idea of this study was to search for a solution that would be robust to a variety of digital image degradations and manipulations (e.g., downsampling, color space modification, compression, etc.). Thus, no attempt was made to normalize the images or to acquire high quality, distortion free images. While the authors show that this approach can be successful and may thus be of use in consumer-level applications involving variable-quality images, it remains unclear whether this approach can scale up to more than five categories or whether image quality biases themselves contributed to performance.

STATISTICAL MODELS OF VISION AND STYLE FEATURES

The use of wavelet and other multiscale techniques in stylometric analysis has been driven more by their utility in edge and orientation detection than any connection to the organization of the visual cortex.[26–28] On the other hand, our recent work has involved the development of new stylometric techniques that take as a starting point models of visual cortex. These models attempt to explain the organization of visual cortex and the encoding of visual information as an efficient, if not optimal, means of matching and coding the information contained in natural scenes. For example, the sparse coding model of Olshausen and Field seeks to explain the response properties of simple cells in primary visual cortex in terms of the statistical structure of natural images.[29] That is, the brain itself builds a model of the visual world that is an optimal representation of the statistical structure in natural scenes, where optimality is defined in terms of coding efficiency. According to their model, visual inputs should be represented using activations of as *few* model neurons as possible, such that the overall response characteristics of neurons is sparse, while at the same time preserving accuracy of neural representation by minimizing the error between inputs and their reconstructions.[29] Using a linear image model

$$I(x,y) = \sum_i \alpha_i \phi_i(x,y) + \epsilon, \qquad (1)$$

for some pixel location x, y in image patch I with coefficients α and Gaussian noise term ϵ, the goal is to learn the functions ϕ, assuming a sparse prior on α (e.g., a Cauchy distribution). The resultant functions ϕ, which are treated as a basis for representing inputs, possess a structure very similar to the response properties of cortical simple cells when trained on natural image inputs, in that they are spatially localized and have orientation and spatial frequency selectivity.[29,30] Figure 1a shows a set of sparse coding basis functions trained on the set of natural images used in Ref 29. Furthermore, although the sparse coding model builds a basis for the space $\mathbb{R}^n$, where n is the number of pixels in an input image patch, the basis functions are merely linearly independent and not guaranteed to be orthogonal. Such a representation results from the constraint that the patch coefficient distributions be sparse, that is, that only a few basis functions should have significantly nonzero weight for representing any particular patch. This stands in contrast to an orthogonal basis for the inputs such as a Fourier basis or a PCA basis created from the input patches, each of which would potentially utilize *all* basis functions in creating a minimum-error reconstruction for each patch.

The sparse coding model was first used for stylometry to capture the structure of a set of authentic drawings by Pieter Bruegel the Elder[31] where the basis functions were then used to distinguish authentic and imitation drawings based on how efficiently

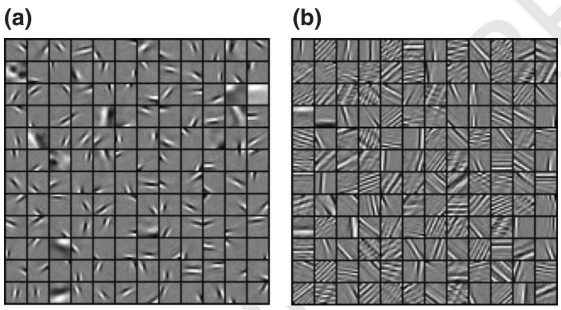

FIGURE 1 | Two set of basis functions trained using the sparse coding model. The set (a) was trained on the natural image set used in the original work of Olshausen and Field.[29] The set (b) corresponds to the art image (shown in Figure 2) that produced a set of basis functions that had minimal average histogram intersection using the distributions of spatial frequency and orientation bandwidth with respect to the basis functions (a).

they represented the statistical structure in a set of test images. A further study[32] explored the extent to which various statistics derived from the learned sparse coding basis functions were useful for classification according to style, with some promising initial results. In particular, the authors trained a set of basis functions to represent each work in a set of art images by several artists. The learned functions were then examined for features such as spatial frequency and orientation selectivity and spatial frequency and orientation bandwidth that were capable of grouping the images by artist.[32] Additionally, statistics derived from sparse coding basis functions have proven useful for organizing artwork according to its painterliness.[33]

The efficiency criterion of the sparse coding model suggests that comparing sets of learned filters—and the statistical characteristics of coefficient distributions for inputs—may highlight the underlying statistical differences in potentially different sets of inputs. Furthermore, because the sparse coding model captures higher order statistical characteristics as opposed to two-point correlations between spatial frequencies and orientations, a learned sparse coding model is a window into the important *structure* present in images, and in particular a window that allows insight into structure at a scale that is likely to be relevant for human perception.

Previous work has shown that artwork possesses a statistical structure similar to that of natural scene images, but with important distinctions. Work by Graham and Field[34] and Redies et al.[35] established that natural scenes and visual art share similar Fourier spatial frequency amplitude spectra, which are found to be roughly scale invariant (amplitude scales approximately as $1/f^k$ where $k \approx 1$, for frequency f). Spatial frequency and orientation information can be indicative of some of the higher order statistical properties that differentiate art from natural images and highlight the effectiveness of the sparse coding model at articulating these differences, albeit in a numerical fashion. In order to see this, we computed the following features on the set of basis functions shown in Figure 1a and on a set of 308 high-resolution images of works of art (the same set used in Ref 32):

1. Distribution of spatial frequency bandwidths: given the two-dimensional Fourier transform of each basis function, compute the bandwidth in octaves (measured by full width at half-maximum) of the function, averaged across all orientations, centered around its peak spatial frequency ω^*. This quantity measures how selective the basis functions are for their preferred spatial frequencies.

2. Distribution of orientation bandwidths: given the two-dimensional Fourier transform of each basis function, compute the bandwidth in octaves (measured by full width at half-maximum) of the function, averaged across all spatial frequencies, centered around its peak orientation θ^*. This quantity measures how selective the basis functions are for their preferred orientations.

Here, ω^* and θ^* are defined as follows: given the two-dimensional Fourier transform $F(\omega, \theta)$ of a basis function (viewed as a function of frequency ω and angle θ),

$$\omega^* = \arg\max_{\omega} \frac{1}{|\Theta|} \sum_{\theta \in \Theta} |F(\omega, \theta)|,$$

$$\theta^* = \arg\max_{\theta} \frac{1}{|\Omega|} \sum_{\omega \in \Omega} |F(\omega, \theta)|.$$

We can compare the distributions of these quantities between images using, for example, the symmetrized histogram intersection statistic,[36]

$$HI(H_1, H_2) = \frac{1}{2}\left[\frac{\sum_b \min(H_1[b], H_2[b])}{\sum_b H_1[b]} + \frac{\sum_b \min(H_1[b], H_2[b])}{\sum_b H_2[b]}\right]$$

for histograms H_1, H_2, and bins b. Note that histogram intersection is a *similarity* metric, rather than a dissimilarity metric, so that large values of *HI* imply greater similarity between images. For the experiments presented here, we considered both the spatial frequency bandwidth and orientation bandwidth histogram intersections.

Figure 1b shows a set of basis functions trained on an art image from our dataset; in particular, these basis functions correspond to the art image that produced the set of basis functions with *smallest* average histogram intersection for spatial frequency and orientation information. This set of basis functions was trained using the image in Figure 2a. This image, while depicting a 'natural scene,' that is, a village, is itself not particularly realistic. On the other hand, the image in Figure 2b, which is the image corresponding to the set of basis functions that had the *largest* average histogram intersection with the natural image bases, depicts a natural scene not unlike the ones used to train the natural image bases, albeit in a painterly style.

If on the other hand we consider only spatial frequency information, the result is somewhat different. The closest image, shown in Figure 2d, is indeed a 'natural scene' (of a village in a drawing by an imitator of Bruegel), thought it is quite different from the painterly one in Figure 2b. The furthest-away art image is an abstract rendering with little in common with natural scenes, and is shown in Figure 2c. Figure 3 shows the histograms of spatial frequency bandwidths for the basis functions corresponding to the images in Figure 2c and d, along with the spatial frequency bandwidth histogram for the natural image basis. Clearly, the bandwidth distributions for the natural image basis and the basis for the art image in Figure 2d and quite different from the basis trained on the image shown in Figure 2c, suggesting that this approach is capable of meaningfully relating the statistical structure of natural scenes and art.

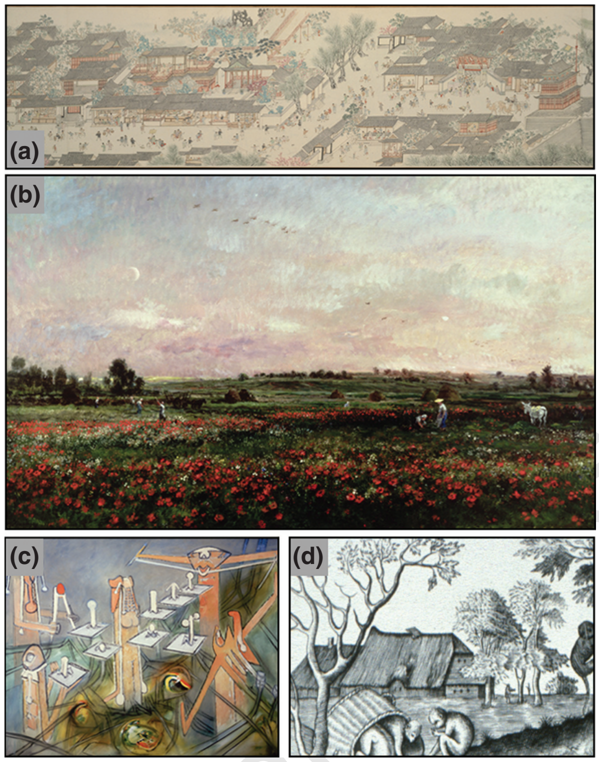

FIGURE 2 | Four art images from the database used in our analysis. Image (a) corresponds to the image that had the weakest similarity to the natural images, using the histogram intersection statistic on the spatial frequency and orientation bandwidth distributions. Image (b) is the art image that was maximally similar to the natural images. Note that it depicts a common natural scene (albeit in a painterly manner). Image (c) is the art image that produced the basis functions most different from the natural image basis functions, according to the histogram images between the spatial frequency bandwidth distributions *only*. Image (d) is a detail of the art image that was maximally similar under the same analysis (i.e., considering only spatial frequency information). Images (a)–(c) are courtesy of the Herbert F. Johnson Museum of Art, Cornell University, and image (d) is courtesy of the Metropolitan Museum of Art, New York.

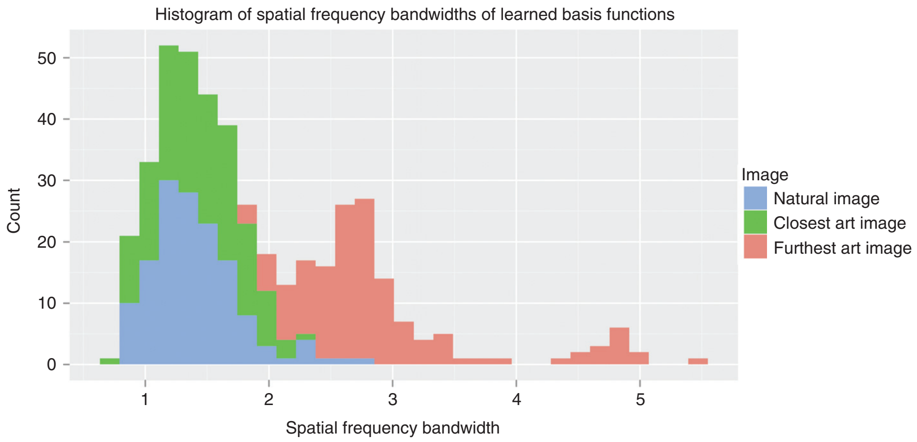

FIGURE 3 | Histogram of basis function spatial frequency bandwidths for the natural image basis (blue) and the images shown in Figure 2c and d (orange and green, respectively).

PERCEPTUALLY DRIVEN FEATURE SELECTION

Historically, psychologists and art historians have taken an interest in the perception of style in visual art, at least from a theoretical perspective.[37–41] More recent work by Leder and colleagues has elucidated a number of basic properties of style perception. In their work on empirical aesthetics, style as a property of art has been treated as a candidate for a special mode of visual processing. The model of aesthetic experiences by Leder and colleagues,[42] is built around the central assumption that processing of style—the manner of depiction as opposed to content of depictions—distinguishes relevant ways of approaching artworks and yields style-related processing as a distinctive, art-typical experience.

Beyond the process of generating features, it can be useful to take advantage of perceptual information for *classifying* art images according to their style.

Classification schemes can take advantage of perceptual information in several ways, but perhaps the most obvious is to consider perceptual similarities between works, derived using psychophysical experiments in which human subjects rate the similarity between two works on some predefined scale. Although such an approach is certainly laborious, it has the potential to yield a significant amount of fruitful data. Once these similarities are obtained, they can be used to train classification models in a supervised fashion or to 'discover' latent stylistic dimensions in an unsupervised learning model. As an example of the former, weighting statistical features based on perceptual information leads to automatic stylistic distinctions that agree well with human perceptual judgments.[32] In essence, stylistic perception is not random, and taking advantage of perceptual information can guide the use of image features in explaining variations in artistic style perception.

Another possible approach would be to use perceptual data to learn feature weights in the context of logistic regression.[43] Such a model is a natural candidate as it transforms an inner product that reflects similarity between two images into a scale (i.e., a probability in $[0, 1]$) that better reflects the intuitive notion of similarity. For example, if pairs of art images were labeled 'similar' and 'dissimilar,' these binary labels would serve as natural class labels in a logistic regression setting. For example, given two images I_1, I_2, a value L that equals 1 if the two images are similar and 0 if they are not, and some set of weights β, we let

$$P(L|I_1, I_2) = \sigma\left(\beta_0 + \sum_{j=1}^{M} \beta_j \kappa_j(\phi_j^{I_1}, \phi_j^{I_2})\right)^L \times \left[1 - \sigma\left(\beta_0 + \sum_{j=1}^{M} \beta_j \kappa_j(\phi_j^{I_1}, \phi_j^{I_2})\right)\right]^{1-L}.$$

In this case, κ_j is some similarity function between features $\phi_j^{I_1}$ and $\phi_j^{I_2}$, which could be scalar, vector-valued, etc., and σ is the logistic sigmoid function. More complex models could be utilized to take advantage of richer perceptual information, such as the similarity ratings mentioned above, rather than binary inputs.

Another possible avenue for taking advantage of perceptual information in classifying art images according to style is to utilize this information for feature elimination, for example, as in recursive feature elimination,[44] or by performing a biased dimensionality reduction that takes advantage of perceptual information. Some previous work has also been done to relate statistical features to the axes of embeddings of art images obtained using multidimensional scaling.[43,45]

CONCLUSION

Although a great deal of progress has been made in understanding and organizing art images according to their style, we argue that a significant majority of the research conducted in this field does not take advantage of the fundamental connection between visual perception and artistic style—and the ways in which style should be measured. Because style is ultimately perceived, it is important to concentrate on developing features and classification methods for organizing and understanding art images that utilize the information vision science and visual psychology can offer about human perception. Moreover, because our visual system is well adapted to the statistical structure of the natural world, we believe that the key to understanding the salient characteristics of artistic style is to quantify the intrinsic differences between art and natural images. Understanding this fundamental distinction will also allow us to refine our knowledge of human perception. In a sense, the way in which natural scene statistics are manipulated to create visual art will provide insight into the 'allowable deformations' of visual content that are important for human perception. Because we are dealing with digital representations of art, image processing techniques will remain critical in enabling us to quantify the style present in works of art. Nevertheless, these techniques should be shaped and utilized in accordance with our understanding of the fundamental processes in human vision. Such an approach will also allow us to extend the scientific reach of stylometric investigations from 'canned' problems such as authentication or attribution of works whose attribution is already known to subtler and more complex descriptions of art. With the growing availability of digital representations of many types of cultural heritage, the ability to meaningfully organize these objects is of ever-increasing importance.

ACKNOWLEDGMENTS

D.J.G., J.M.H., and D.N.R. gratefully acknowledge the support of the Kress Foundation for this work, under the grant 'The Workshop Practices of Botticelli Before Rome: Collaboration with Filippino Lippi in 'The Story Of Esther.'

REFERENCES

1. Spencer RD. *The Expert versus the Object: Judging Fakes and False Attributions in the Visual Arts*. Oxford University• Press; 2004.
2. Morelli G. *Italian Painters: Critical Studies of their Works*. J. Murray; 1892.
3. Wollheim R. *Giovanni Morelli and the Origins of Scientific Connoisseurship*. Allen Lane; 1973.
4. Ginzburg C. Morelli, Freud and Sherlock Holmes: clues and scientific method. *Hist Workshop J* 1980, 9:5–36.
5. Hockey S. The history of humanities computing. *A Companion to Digital Humanities*. Blackwell Publishing; 2004.
6. Lowe D. Object recognition from local scale-invariant features. Proc Seventh IEEE Conf Comput Vis 1999, 2, 1150–1157.
7. Rao A, Srihari R, Zhang Z. Spatial color histograms for content-based image retrieval. Proceedings of 11th IEEE International Conference on Tools with Artificial Intelligence, 1999, 183–186.
8. Gunsel B, Sariel S, Icoglu O. Content-based access to art paintings. IEEE International Conference on Image Processing, 2005. ICIP 2005, 2005, 2, 558–561.
9. Hughes JM, Graham DJ, Rockmore DN. Stylometrics of artwork: uses and limitations. Proceedings of SPIE: Computer Vision and Image Analysis of Art, 2010, 75310C1.
10. Li J, Yao L, Hendriks E, Wang JZ. Rhythmic brushstrokes distinguish Van Gogh from his contemporaries: findings via automated brushstroke extraction, 2011, submitted• for publication.
11. Yao L, Li J, Wang J. Characterizing elegance of curves computationally for distinguishing Morrisseau paintings and the imitations. 16th IEEE International Conference on Image Processing (ICIP), 2009, 73–76.
12. Sablatnig R, Kammerer P, Zolda E. Hierarchical classification of paintings using face- and brush stroke models. Proceedings of Fourteenth International Conference on Pattern Recognition 1998, 1, 172–174.
13. Johnson JCR, Hendriks E, Berezhnoy I, Brevdo E, Hughes S, Daubechies I, Li J, Postma E, Wang J. Image processing for artist• identification. *IEEE Signal Process Mag* 2008, 37.

14. van der Maaten L, Postma E. Identifying the real van Gogh with brushstroke textons, White paper, Tilburg University, February 2009.
15. Berezhnoy I, Postma E, van den Herik H. Authentic: computerized brushstroke analysis, 2005, 1586–1588.
16. Taylor RP, Micolich AP, Jonas D. Fractal analysis of Pollock's drip paintings. *Nature* 1999, 399
17. Taylor R, Guzman R, Martin T, Hall G, Micolich A, Jonas D, Scannell B, Fairbanks M, Marlow C. Authenticating Pollock paintings using fractal geometry. *Pattern Recognition Lett* 2007, 28:695–702.
18. Jones-Smith K, Matthur H. Fractal analysis: revisiting Pollock's drip paintings. *Nature* 2006, 444:E9–E10.
19. Keren D. Painter identification using local features and naïve Bayes. Proceedings of the 16th International Conference on Pattern Recognition, 2002, 2, 474–477.
20. Lyu S, Rockmore D, Farid H. A digital technique for art authentication. *Proc Natl Acad Sci* 2004, 101:17006–17010.
21. Hughes JM, Mao D, Rockmore DN, Wang Y, Wu Q. Empirical mode decomposition analysis for visual stylometry, 2011, submitted for publication.
22. Lin L, Wang Y, Zhou H. Iterative filtering as an alternative algorithm for empirical mode decomposition. *Adv Adapt Data Anal* 2009, 1:543–560.
23. Temel B, Kilic N, Ozgultekin B, Ucan ON. Separation of original paintings of matisse and his fakes using wavelet and artificial neural networks. *Istanbul Univ J Elect Electron Eng* 2009, 9.
24. Berezhnoy I, Postma E, van den Herik J. Computer analysis of van Gogh's complementary colours. *Pattern Recognition Lett* 2007, 28:703–709 Pattern Recognition in Cultural Heritage and Medical Applications.
25. Zujovic J, Gandy L, Friedman S, Pardo B, Pappas T. Classifying paintings by artistic genre: an analysis of features & classifiers. IEEE International Workshop on Multimedia Signal Processing, 2009. MMSP'09, 2009, 1–5.
26. Jones JP, Palmer LA. The two-dimensional spatial structure of simple receptive fields in cat striate cortex. *J Neurophysiol* 1987, 58:1187–1211.
27. Hubel DH, Wiesel TN. Receptive fields and functional architecture of monkey striate cortex. *J Physiol* 1968, 195:215–243.
28. Daugman JG. Uncertainty relation for resolution in space, spatial frequency, and orientation optimized by two-dimensional visual cortical filters. *J Opt Soc Am A* 1985, 2:1160–1169.
29. Olshausen B, Field D. Emergence of simple-cell receptive field properties by learning a sparse code for natural images. *Nature* 1996, 381:607–609.
30. Olshausen B, Field D. Sparse coding with an overcomplete basis set: a strategy employed by V1?. *Vision Res* 1997, 37:3311–3325.
31. Hughes JM, Graham DJ, Rockmore DN. Quantification of artistic style through sparse coding analysis in the drawings of Pieter Bruegel the Elder. *Proc Natl Acad Sci* 2010, 107:1279–1283.
32. Hughes JM, Graham DJ, Jacobsen CR, Rockmore DN. Comparing higher-order• spatial statistics and perceptual judgements in the stylometric analysis of art, 2011.
33. Foti NJ, Hughes JM, Rockmore DN. Nonparametric sparsification of complex multiscale• networks. *PLoS One* 2011.
34. Graham DJ. Statistical regularities of art images and natural scenes: spectra, sparseness and nonlinearities. *Spatial Vision* 2007, 21:149–164.
35. Redies C. Fractal-like image statistics in visual art: similarity to natural scenes. *Spatial Vision* 2007, 21: 137–148.
36. Barla A, Odone F, Verri A. Histogram intersection kernel for image classification. Proceedings of 2003 International Conference on Image Processing 2003, 3, 513–16.
37. Wölfflin H. *Principles of Art History: The Problem and Development of Style in Later Art*. Dover; 1950.
38. Gombrich E. *Art and Illusion: A Study in the Psychology of Pictorial Representation*. Phaidon Press; 2002.
39. Arnheim R. *Art and Visual Perception: A Psychology of the Creative Eye*. University of California Press; 1954.
40. Gibson JJ. The information available in pictures. *Leonardo* 1971, 4:27–35.
41. Goodman N. *Language• of Art*. Oxford; 1968.
42. Leder H, Belke B, Oeberst A, Augustin D. A model of aesthetic appreciation and aesthetic judgments. *Brit J Psychol* 2004, 95:489–508.
43. Bishop CM. *Pattern Recognition and Machine Learning*. 1st ed. Springer; 2007.
44. Guyon I, Weston J, Barnhill S, Vapnik V. Gene selection for cancer classification using support vector machines. *Machine Learn* 2002, 43:389–422.
45. Graham D, Friedenberg J, Rockmore D, Field D. Mapping the similarity space of paintings: image statistics and visual perception. *Visual Cogn* 2009.

FURTHER READING

Atick JJ, Redlich AN. What does the• retina know about natural scenes?

Neural Comput 1992, 4:196–210.

Graham DJ, Chandler DM, Field DJ. Can the theory of "whitening" explain the center-surround properties of retinal ganglion cell receptive fields? *Vision Res* 2006, 46:2901–2913.