

Texton-Based Analysis of Paintings

Laurens J.P. van der Maaten^{a,b} and Eric O. Postma^c

^aUniversity of California, San Diego, 9500 Gilman Drive, La Jolla CA 92093, USA

^bDelft University of Technology, Mekelweg 4, 2628 CD Delft, The Netherlands

^cTilburg University, P.O. Box 90153 5000 LE Tilburg, The Netherlands

ABSTRACT

The visual examination of paintings is traditionally performed by skilled art historians using their eyes. Recent advances in intelligent systems may support art historians in determining the authenticity or date of creation of paintings. In this paper, we propose a technique for the examination of brushstroke structure that views the wildly overlapping brushstrokes as texture. The analysis of the painting texture is performed with the help of a texton codebook, i.e., a codebook of small prototypical textural patches. The texton codebook can be learned from a collection of paintings. Our textural analysis technique represents paintings in terms of histograms that measure the frequency by which the textons in the codebook occur in the painting (so-called *texton histograms*). We present experiments that show the validity and effectiveness of our technique for textural analysis on a collection of digitized high-resolution reproductions of paintings by Van Gogh and his contemporaries.

As texton histograms cannot be easily be interpreted by art experts, the paper proposes two approaches to visualize the results on the textural analysis. The first approach visualizes the similarities between the histogram representations of paintings by employing a recently proposed dimensionality reduction technique, called t-SNE. We show that t-SNE (applied on texton histograms) separates paintings created by Van Gogh from those created by other painters. In addition, the period of creation is faithfully reflected in the t-SNE visualizations. The second approach visualizes the similarities and differences between paintings by highlighting regions in a painting in which the textural structure of the painting is unusual. We illustrate the validity of this approach by means of an experiment in which we highlight regions in a painting by Monet that are not very “Van Gogh-like”. Taken together, we believe the tools developed in this study are capable of assisting for art historians in support of their study of paintings.

Keywords: Art analysis, texture features, non-linear dimensionality reduction.

1. INTRODUCTION

Digital analysis techniques are entering the domain of art historians. The development of advanced computer vision algorithms and the availability of powerful computers to run these algorithms in reasonable time have opened up a wide range of new possibilities for the computer-supported examination of paintings. Currently, painting analysis and artist identification are performed by art experts that have considerable experience with the works of one or more painters. Although experts have a variety of techniques at their disposal – such as canvas weave count, dendrochronological analysis of the wood of the frame, chemical analysis of the pigments, and x-radiography of the painting or support – visual assessment of the painting is still one of their most important tools. An important cue in identifying the artist of a painting is the “handwriting” of the painter: the brushstrokes and brushstroke configurations that reveal the painter’s style. Despite the variations in form and appearance of a painter’s brushstrokes, the artist’s handwriting can be recognized by skilled art experts, although it is hard to explicate to laymen what the characteristic elements of a painter’s handwriting are. Intelligent image analysis and machine learning techniques that are sensitive to the brushstroke texture may support the art expert in detecting and visualizing painter-specific brushstrokes, and provide objective evidence for the attribution of a painting to an artist.

Further author information: (Send correspondence to L.v.d.M.)
L.v.d.M.: E-mail: lvdmaaten@gmail.com, Telephone: 1 858 534 8647
E.P.: E-mail: e.o.postma@uvt.nl, Telephone: +31 (0)13 466 2655

A comprehensive overview of artistic image analysis is given by Hurtut.¹ Previous work on digital techniques for painting analysis and artist identification focused on the assessment of, e.g., color use in (Van Gogh) paintings using filter-based approaches,² on capturing statistical information from segmented or outlined brushstrokes,³ on wavelet analysis of the paintings,^{4,5} analysis of lighting structure,^{6,7} analysis of illuminants and perspective,^{8–11} and analysis of shapes depicted in the paintings.¹² Though the techniques presented in these papers are very useful and have advanced the state-of-the-art in, for instance, artist identification, the results of the techniques are typically hard to interpret. In particular, most statistical techniques represent a painting as a point in a high-dimensional “feature space”.^{2–5,8} As art experts are only familiar with the “painting space”, it is hard for them to interpret the statistical measurements performed by modern digital techniques.

In this paper, we present a technique that evaluates the brushstroke configuration by considering it to be built up from small prototypical textural patches called *textons*. The technique learns a texton codebook from a collection of paintings, and subsequently, represents paintings by means of a histogram that measures how frequent a particular texton appears in the painting at hand. Our motivation for using such a texton-based technique for brushstroke analysis is two-fold: (1) texton-based texture analysis has been recently shown to perform very well in texture analysis and synthesis tasks^{13–15} and (2) the texton-based framework allows for natural ways to visualize the results of the texture analysis in representations that can be interpreted by art experts. In particular, we present two approaches to the visualization of the results of the brushstroke analysis. The first approach visualizes the similarities between the texton histograms computed from paintings using a recently proposed dimensionality reduction technique, called t-Distributed Stochastic Neighbor Embedding (t-SNE).¹⁶ Combined with the texton-based texture analysis, t-SNE reveals a clear separation of paintings created by Van Gogh and those created by other painters. The second approach visualizes the results of the brushstroke analysis by “projecting back” into the painting space: the approach highlights regions in the painting in which the textural structure of the painting is unusual.

The outline of the remainder of this paper is as follows. In Section 2, we present our approach to brushstroke analysis using texton-based texture analysis. Section 3 outlines how the results of the brushstroke analysis can be visualized using t-SNE in order to, for instance, identify differences between contemporary painters. Section 4 presents the results of experiments in which we combine the technique for brushstroke analysis with t-SNE, and apply the resulting combination on a collection of paintings by Van Gogh and his contemporaries. In Section 5, we discuss some of the shortcomings of the combination of techniques considered in Section 4, and we propose the approach that projects the results of the brushstroke analysis back into the painting space. Section 6 provides a more detailed discussion of the results presented in this study. Section 7 presents our conclusions, as well as possible directions for future work.

2. TEXTURE ANALYSIS

Given that the segmentation of individual brushstrokes from a painting is unfeasible,⁵ from an image analysis perspective, brushstroke analysis can be considered as the analysis of the texture of the painting. The wildly overlapping brushstrokes form a textural cue of the painter’s handwriting. We note that using texture analysis is not the only way to address the analysis of brushstrokes: one may also try to “segment” individual brushstrokes in the painting and compute characteristics from the segmented brushstrokes.^{2,17} Since brushstroke segmentation is tedious due to the lack of contrast between individual brushstrokes, we do not opt for such an approach here.

Typically, texture analysis is performed by applying a bank of filters that respond to intensity transitions in the input image. Motivated by the characteristics of the primate visual system,¹⁸ the filters in the bank are generally selected in such a way that they respond to spatial frequencies at various scales and orientations,^{19–21} such as a Gabor filter bank.²² Filters with very similar receptive fields as Gabor filters also emerge when the independent components of natural images are estimated,²³ when natural images are analyzed using sparse coding,²⁴ or when certain models based on Markov Random Fields are trained on natural images.²⁵ These observations, together with cell recordings from cells in the primate visual cortex V1,¹⁸ suggest Gabor-like filters form the basis of natural images.

Filter-based approaches have long been considered to be the supreme method for texture analysis. Recently, however, the supremacy of filter-based approaches in texture analysis has been questioned, with the presentation of a pixel-based approach to texture analysis that performs on par with computationally more expensive



Figure 1. Illustration of the construction of a texton codebook. Patches are collected from random locations in each painting. Subsequently, prototypical brushstroke textons are identified by a clustering method and stored in a texton codebook.

approaches based on filters.¹⁴ One of the main arguments against the use of filter-based approaches is that they typically employ filters with a Gaussian envelope, as a result of which the filters *smooth* the image before they measure the presence of (oriented) high spatial frequencies. The smoothing is necessary to remove noise that masks image gradients. Unfortunately, the smoothing may distort or remove pivotal information of relevance to the texture analysis task. In the case of the brushstroke texture of paintings, such details may correspond to individual hairs in the brush used by the painter that contain valuable cues for painting analysis and artist identification. The smoothing problem may be resolved by employing filters with a small support such as in complex wavelet transforms,²⁶ or by employing pixel-based texture representations such as the texton-based texture representation we introduce below.

Textons are considered to be the fundamental building blocks of texture,²⁷ just like grapheme are the fundamental building blocks of handwriting²⁸ and phonemes are the fundamental building blocks of speech.²⁹ We can view brushstroke texture as a superposition of brushstroke textons. Note that although a texton is a fundamental building block of texture, it is not identical to a brushstroke: the fundamental building block of paintings*. Texton-based approaches entail the use of normalized pixel-based texture representations and are, therefore, not hampered by the smoothing problem. In our texton-based modeling of paintings, we first identify the prototypical brushstroke textons of the complete collection of digital reproductions of paintings, and store those in a texton codebook. The artist can be viewed upon as a probabilistic generator of textons, who generates textons according to a painter-specific distribution over the textons in the codebook[†]. For each painting, the texton distribution is approximated by means of a texton histogram. We discuss the two main stages of the texton-based approach, viz. the texton codebook construction and the texton histogram computation, separately below.

We construct the texton codebook as follows. We select 5,000 square image patches from random locations in

* Approaches that try to use brushstrokes as the fundamental blocks have also been proposed,^{2,17} but are difficult to apply in practice because segmentation of brushstrokes or handwriting remains an unsolved problem.³⁰

[†]Note that the distribution by which a painter generates textons may be different for different periods in his artistic life.




Figure 2. Illustration of the construction of a texton histogram. A window is slid over the texture image, and the histogram bin associated with the most similar codebook texton is incremented at each spatial location. After normalization, the texton histogram represents the relative number of times that a codebook texton appears in the painting.

each painting. In the collection of patches that were selected from all paintings, prototypical patches are identified by applying a clustering algorithm such as k -means clustering, Kohonen maps,³¹ spectral clustering,^{32,33} or affinity propagation.³⁴ The cluster centers or exemplars obtained from the clustering algorithm form the texton codebook. The construction of the texton codebook is illustrated in Figure 1. An example of a codebook learned from a collection of Van Gogh paintings using affinity propagation is shown in Figure 4.

Once the texton codebook is learned, each painting can be represented by means of a texton histogram that estimates the distribution by which the painter generated textons in the painting process. The height of each bin in the texton histogram for a painting represents the frequency of occurrence of the associated codebook texton in the painting. The histogram is created using a sliding window that moves over all locations in the painting. At each location, the contents of the sliding window are compared to all codebook textons. The most similar codebook texton (e.g., in the Euclidean sense) is defined as the matching texton. The texton histograms are normalized to sum up to 1 in order to estimate the painter's texton distribution. The extraction of texton histograms from a painting is illustrated in Figure 2.

In art experts' visual analysis of a painting, a variety of different scales is usually examined by the experts. The examined scales range from a coarse level (such as the depicted scene, the use of complementary colors, etc.) to a very fine level (such as the size of employed brushes, the thickness of brush hairs, etc.). It thus seems likely that a computer analysis of paintings should examine a range of scales as well. Texton-based texture modeling readily facilitates such a multi-scale approach: the computation of texton histograms can be performed using textons of different sizes, and the resulting histograms can be combined into a multi-dimensional feature vector. In our experiments with the brushstroke texture method, we extract texton histograms for six different scales (i.e., for six different sizes of the sliding window), and we combine the histograms computed on each scale into a single painting representation.

3. VISUALIZATION USING T-SNE

Although the texton histograms developed in the previous section have been shown to be good representations for textured surfaces such as paintings,^{14,15} they are often hard to interpret by art experts. Indeed, visualizations of the texton codebooks may reveal some information. For instance, in the case of Van Gogh paintings, the texton codebooks may reveal many textons with large sudden intensity transitions, reflecting Van Gogh's typical painting style of using pronounced brushstrokes. However, the information that can be obtained by interpreting texton codebooks is certainly limited, and typically only shows information that was already known by art experts beforehand. Moreover, visualization of the texton codebooks does not provide insight into the differences and similarities between paintings.

In order to provide art experts insight into the amount of textural differences between paintings (as measured by the texton-based approach), a method is needed that can lay out the paintings in a "painting map" that faithfully represents the differences between the texton histograms. The texton histograms constitute vectors of a high dimensionality: the dimensionality is equal to the number of textons in the codebook. Dimensionality reduction techniques transform the high-dimensional texton histograms into points in a two or three-dimensional space while preserving as much of the structure of the data as possible. Hence, they may be used to visualize the similarities between texton histograms in terms of distances between points in a two-dimensional scatterplot. The quality of the dimensionality-reduced visualization depends on the quality of the features and on the quality of the dimensionality reduction technique. Dimensionality reduction is traditionally often performed using Principal Components Analysis (PCA), however, this technique is not very suitable for our purposes for two main reasons: (1) PCA is hampered by its linear nature, whereas the high-dimensional data points may form a nonlinear manifold, and (2) PCA tends to ignore the local structure in the high-dimensional data points by focusing on the preservation of the global structure of the data. Recently, several dimensionality reduction techniques have been developed that focus on retaining the local structure of high-dimensional data points that form a nonlinear manifold, e.g., Isomap,³⁵ Local Linear Embedding,³⁶ and diffusion maps.³⁷ We have reviewed these techniques³⁸ and found them to suffer from serious shortcomings that prohibit successful visualization of real-world data. Therefore, we opt to use a novel dimensionality reduction technique, called t-Distributed Stochastic Neighbor Embedding (t-SNE),¹⁶ that is much better at preserving the local structure of the high-dimensional data in the visualization.

t-SNE starts by converting the pairwise distances between the high-dimensional data points (i.e., the texton histograms) to conditional probabilities by centering a Gaussian over each point and measuring the density of all other points under the Gaussian. Mathematically, the pairwise similarity $p_{j|i}$ between the high-dimensional data points \mathbf{x}_i and \mathbf{x}_j is given by the conditional probability:

$$p_{j|i} = \frac{\exp(-\|\mathbf{x}_i - \mathbf{x}_j\|^2/2\sigma_i^2)}{\sum_{k \neq i} \exp(-\|\mathbf{x}_i - \mathbf{x}_k\|^2/2\sigma_i^2)}, \quad (1)$$

where σ_i is set in such a way as to give each Gaussian the same entropy or perplexity, and $p_{i|i}$ is set to 0. Denoting the low-dimensional counterparts of \mathbf{x}_i and \mathbf{x}_j by \mathbf{y}_i and \mathbf{y}_j , we define similar conditional probabilities between the points in the low-dimensional space. In order to address a problem known as the *crowding problem*[‡], t-SNE does not employ a Gaussian in the low-dimensional space, but a Cauchy distribution (which is a special case of a Student-t distribution, hence the name t-SNE). Mathematically, the pairwise similarity $q_{j|i}$ between the low-dimensional data points \mathbf{y}_i and \mathbf{y}_j is defined as:

$$q_{j|i} = \frac{(1 + \|\mathbf{y}_i - \mathbf{y}_j\|^2)^{-1}}{\sum_{k \neq i} (1 + \|\mathbf{y}_i - \mathbf{y}_k\|^2)^{-1}}, \quad (2)$$

where again, we set $q_{i|i}$ to 0. The locations of the low-dimensional data points \mathbf{y}_i are determined by minimizing the sum of the Kullback-Leibler divergences between the conditional probability distributions P_i and Q_i . Hence,

[‡]The crowding problem is a fundamental problem that occurs when one tries to model intrinsically high-dimensional data in a low-dimensional space in such a way, that local similarities are preserved as good as possible.




Figure 3. Illustration of the construction of a t-SNE map. The points in the low-dimensional map are arranged in such a way that the conditional probabilities in the map $q_{j|i}$ are similar to the conditional probabilities in the data $p_{j|i}$.

t-SNE minimizes the cost function:

$$C = \sum_i KL(P_i||Q_i) = \sum_i \sum_{j \neq i} p_{j|i} \log \frac{p_{j|i}}{q_{j|i}}, \quad (3)$$

with respect to the coordinates \mathbf{y}_i of the points in the low-dimensional space using a gradient descent method (although alternative optimization techniques have also been investigated³⁹). t-SNE has been shown to perform very strong in the visualization of a variety of real-world datasets.¹⁶ The construction of a t-SNE map is illustrated in Figure 3.

4. EXPERIMENTS

In this section, we present the results of our experiments with texton-based texture analysis and t-SNE on a collection of paintings by Van Gogh and his contemporaries. The setup of these experiments is described in Section 4.1. Section 4.2 presents the results of our experiments

4.1 Experimental Setup

We apply our techniques on a recently released dataset of 117 high-resolution digital reproductions of paintings.⁴⁰ The dataset contains high-resolution digital 48-bit color reproductions of 117 paintings attributed to Van Gogh and related painters, which we transformed to 8-bit grayscale images for our experiments. The reproductions were created using ektachromes made available by the Van Gogh Museum and the Kröller-Müller Museum (both in The Netherlands). The paintings were normalized in such a way that a square inch of the painting is represented by 196.3×196.3 pixels. Of the 117 paintings, 13 are known not to be painted by Van Gogh and 6 are of disputed authorship. The remaining 98 paintings are generally accepted as authentic Van Gogh paintings. Each painting is labeled with its authenticity (Van Gogh, non Van Gogh, or disputed), and all authentic Van Gogh paintings are labeled by their creation date (ranging from 1884 to 1890) and creation place.

In order to perform the texton-based texture analysis, we extracted texton histograms for textons of six different sizes: 25×25 , 35×35 , 45×45 , 55×55 , 65×65 , and 75×75 pixels. The six texton codebooks employed in the experiment were constructed using affinity propagation, and contained approximately 500 textons each. The preference value of affinity propagation was set to the median of the negative pairwise Euclidean distances between the extracted textons. Figure 4 shows one of the constructed texton codebooks. Altogether, the six texton histograms form feature vectors with approximately 3,000 dimensions.

In order to visualize the results of the texture analysis in a two-dimensional map, the 3,000-dimensional texton histograms were first reduced to 50 dimensions using Principal Components Analysis (PCA). Subsequently, we




Figure 4. An example of one of the texton codebooks. This codebook was constructed using affinity propagation on textons of size 35×35 pixels.

use t-SNE to reduce the dimensionality of the resulting feature vectors to 2 dimensions. In the high-dimensional space, we set the variance parameters σ_i in such a way that the conditional distributions P_i had a perplexity of 10 (or equivalently, a Shannon entropy of $\log_2(10)$).

4.2 Results

In Figure 5, we present one of the visualizations obtained by performing t-SNE on texton histograms extracted from the paintings in our collection. Each dot represents a single texton histogram (i.e., a single painting). The green dots represent Van Gogh's paintings, whereas the red dots represent established non Van Gogh paintings. Paintings whose attribution is disputed are indicated by blue dots. The visualization reveals that all-but-two non Van Gogh paintings are depicted in the periphery of the visualization. Apparently, the brushstroke texture in these paintings is appropriately captured by the texture histograms and offers an effective, albeit crude, indication of textural differences and similarities. The two paintings that do not stand out in the visualization are the so-called Wacker forgery and a painting by Gauguin. The Wacker forgery is one of a series of forgeries, which fooled renowned Van Gogh experts for years. The Wacker forgery in our collection is quite easy to discriminate from the genuine Van Gogh paintings using global texture analysis:⁵ the Wacker forgery contains more high spatial frequencies than the genuine Van Gogh's. Presumably, the local textons do not capture these global statistics. The same may apply to the painting created by Gauguin. Further analysis is needed to establish this. Despite these two anomalies, the visualization places 11 of the 13 non-Van Gogh paintings in the periphery of the visualization. This is quite a remarkable result, given that our approach is completely data-driven. The visualizations obtained with our approach also suggest attributions of the disputed paintings: some are located in the middle of a cluster of genuine Van Gogh paintings, whereas others are located close to the non Van Gogh paintings in the periphery.

Figure 6 shows a visualization obtained by applying t-SNE on texton histograms that correspond to established Van Gogh paintings only. The dots are colored according to the two main periods in Van Gogh's oeuvre: the Dutch period (1883-1886; red dots) and the French period (1886-1890; blue dots). The visualization shows




Figure 5. Visualization of the dataset of Van Gogh, non Van Gogh, and disputed Van Gogh paintings. The points in the scatter plot are labeled according to the authenticity of the paintings.

a separation between the paintings from both periods (all Dutch paintings are captured in one of three small clusters), and thus captures diagnostic textural elements of the development of Van Gogh's paintings style from his originally sober style to his later (and more famous) lively impressionistic paintings. Art historians may use the our techniques to create visualizations of subsets of paintings to examine more subtle textural differences.

5. PROJECTION INTO PAINTING SPACE

Although scatter plots such as those presented in Figure 5 and 6 provide insight into which paintings (or parts of paintings) are similar and dissimilar, in this case, in terms of their textural content, they also lead to a variety of new questions. Most importantly, the scatter plots raise questions such as the following. *Why are two paintings that are visualized far apart dissimilar? Which features lead to this dissimilarity? Where are these features located in the painting? What are the main textural elements that cause the difference between Van Gogh paintings and paintings by his contemporaries?* Undoubtedly, the reader is able to come up with a variety of related questions.

A prominent approach to try and answer such questions is by “projecting back” the results obtained from the statistical analyses in the “painting space”. A prominent example of such back-projection is the inspection of the receptive fields that are learned by, for instance, Principal Components Analysis (e.g., the well-known *eigenfaces*⁴¹), Nonnegative Matrix Factorization (NMF),⁴² Independent Component Analysis (ICA),⁴³ Restricted Boltzmann Machines (RBMs),⁴⁴ or one of several types of feedforward neural networks. In particular, receptive fields learnt by NMF are very appropriate to this end, as they form a can be thought of as a statistical variant of the recognition-by-parts theory.⁴⁵ Specifically, NMF tends to learn a single “clean” feature (a part) per receptive field that, as a result of which the features can be interpreted quite well. In contrast, more powerful techniques such as PCA tend to learn features that are much more complicated, because they correspond to complex combinations of the features that are present in the data. When presented with our dataset of Van Gogh paintings, NMF presumably learns features that correspond to various stroke shapes (or to various parts of a stroke).





Figure 6. Visualization of the authentic Van Gogh paintings in the dataset. The points in the scatterplot are labeled according to the period in which they are painted.


Unfortunately, the interpretable features that are learnt by NMF come at a price, as these features could only be obtained by severely restricting the flexibility of the model compared to, e.g. PCA. In practice, this results in inferior representations of the paintings when they are projected onto the basis learnt by NMF. In other words, the locations of the points in a scatter plot (such as those in Figure 5 and 6) produced by NMF hardly reflect any information on the intrinsic similarity between the paintings. The reverse is also true: the receptive fields learned by very powerful methods such as t-SNE (if we would explicitly learn these receptive fields⁴⁶) are very hard to interpret, as they are comprised of extremely complex combinations of the features (parts) that are present in the paintings.

As we opted for using a powerful method such as t-SNE in our analysis, we thus need an alternative means to project back the results of the analysis into the “painting space”. Texton-based texture features facilitate such a back-projection by computing the difference between the texton histograms of two paintings. Subsequently, one can employ the texton histogram difference to highlight regions in one of the paintings that occur often in that painting, but not in the other texture image. The highlighting of the painting can be performed using a sliding scale, where the amount of highlighting is proportional to the texton difference between the two paintings.


In Figure 7, we demonstrate this way of back-projecting texton-based texture features into the “painting space”. In the example, we compare a Van Gogh painting with a painting by one of his contemporaries (viz., a painting by Claude Monet), and highlight the regions in the contemporary painting that are not very much “Van Gogh-like”. The resulting visualization indicates that the sky of the contemporary painting is very different from the Van Gogh painting in terms of texton measurements, a result that was acknowledged by art experts. Although in our illustration in Figure 7, we only compared the Monet painting with a single painting by Van Gogh, the approach can readily be extended to facilitate comparisons with the complete oeuvre of Van Gogh.



(a) Van Gogh painting.



(b) Monet painting.



(c) Monet painting with highlights.

Figure 7. Illustration of the back-projection of texton-based texture features onto the painting space. The bottom image is a highlighted version of the Monet painting in the top-right of the figure, in which only textons are highlighted that do not occur frequently in the Van Gogh painting shown in the top-left of the figure. Part of the paintings in (a) and (b) is blocked out to prevent unauthorized digital reproduction of the paintings.

6. DISCUSSION

The results presented above illustrate the potential of digital image analysis and machine learning to support art historians in their analysis of paintings. Of course, the combination of texton-based texture analysis and t-SNE only offers an initial crude characterization of paintings. A complete approach to computer-assisted artist identification should integrate more information than just the local texture characteristics that our texton histograms capture. For instance, the interaction between brushstrokes should be captured, prompting the use of textural features that are less local than our texton features, such as wavelet features⁵ or filter bank statistics.²⁰ Moreover, the color use of Van Gogh (and more specifically, the use of complementary colors) should be captured in global painting features.² Our texton-based approach may also be adapted to employ color information.⁴⁷

The development of an approach that combines a variety of features is the most viable way to obtain clear separation between paintings that are created by different artists based on the visual assessment of the paintings. Typically, the numerical results obtained with applying the texton-based texture analysis and other image analysis methods on digital reproductions of paintings will complement the results obtained with other types of analysis, such as provenance analysis, canvas weave count,⁴⁸ paper analysis,⁴⁹ analysis of lighting structure,^{6,7} analysis of illuminants and perspective,⁸⁻¹¹ analysis of shape,¹² color analysis,² dendrochronological measurements on painting frames and panels,⁵⁰ and radiocarbon dating of canvas.⁵¹ These features can be combined before t-SNE is performed using, e.g., multiple kernel learning (MKL).⁵²

Admittedly, we only presented techniques that can be used as the basis for an interactive intelligent system that can assist art experts in their analysis and contribution of paintings. For instance, a production version of our software should allow art experts to interactively map and visualize subsets of paintings and selected regions of paintings. A production version of our software should also seamlessly be integrated in software that is already widely used by art experts, such as Adobe Photoshop.

7. CONCLUDING REMARKS

The paper presented new techniques that are designed to assist art experts in their assessment and analysis of paintings. We proposed texton-based texture analysis as a technique to analyze the structure of the brushstrokes in a painting. We visualized the results of the brushstroke analysis using a novel dimensionality reduction technique (t-SNE) that constructs a two-dimensional map of the paintings in such a way, that the map reveals the difference between the paintings in terms of brushstroke structure. The results of our experiments with a combination of the texton-based texture analysis and t-SNE indicate the sensitivity of the method to the textural cues underlying the authorship of a painting. An interactive version of this combination of techniques will allow art experts to establish a quantitative assessment of the authenticity or date of creation of paintings. Moreover, we presented an approach that highlights regions in a painting in which the brushstroke texture is different from the typical brushstroke structure in other paintings in a collection. The highlighting technique allows art experts to use the power of computers to examine every detail in a painting. Taken together, the authors hope to have made a good first step towards the development of interactive systems that support art historians in their analyses and attributions of paintings.

Future work aims at aggregating our approach for brushstroke analysis with approaches that use wavelet analysis and filter response statistics, as well as with techniques for, e.g., weave counting based on x-rays of a painting. Combining these approaches into a single system also requires the development of techniques for the automatic alignment of painting photographs and x-rays (and possibly, also other image types such as hyperspectral images), and it requires the combination of different types of features using, e.g., MKL.

ACKNOWLEDGMENTS

Laurens van der Maaten acknowledges support by the Netherlands Organisation for Scientific Research (NWO; grant 640.002.401 and 680.50.0908), and by the EU-FP7 programme on Social Signal Processing (SSPNet). The authors thank the Van Gogh Museum and the Kröller-Müller Museum for their cooperation and for their generous release of the Van Gogh dataset that was employed in this study. The authors thank the Dutch State Service for Cultural Heritage (RCE) for their cooperation. The authors thank Rick Johnson, Ingrid Daubechies, David Stork, Shannon Hughes, Eugene Brevdo, Igor Bereznoy, Jia Li, James Wang, Jan van der Lubbe, Ella Hendriks, and Louis van Tilborgh for many helpful discussions on digital and traditional painting analysis.

REFERENCES

- [1] Hurtut, T., “2D Artistic images analysis, a content-based survey,” (2010).
- [2] Berezhnoy, I., Postma, E., and van den Herik, H., “Computer analysis of van Gogh’s complementary colours,” *Pattern Recognition Letters* **28**(6), 703–709 (2007).
- [3] Sablatnig, R., Kammerer, P., and Zolda, E., “Hierarchical classification of paintings using face- and brush stroke models,” in [*Proceedings of the International Conference on Pattern Recognition*], 172–174 (1998).
- [4] Lyu, S., Rockmore, D., and Farid, H., “A digital technique for art authentication,” *Proceedings of the National Academy of Sciences* **101**(49), 17006–17010 (2004).
- [5] Johnson, C., Hendriks, E., Berezhnoy, I., Brevdo, E., Hughes, S., Daubechies, I., Li, J., Postma, E., and Wang, J., “Image processing for artist identification,” *IEEE Signal Processing Magazine* **25**(4), 37–48 (2008).
- [6] Stork, D., “Optics and realism in Renaissance art,” *Scientific American* **291**(6), 76–84 (2004).
- [7] Johnson, K. and Stork, D., “Estimating the location of illuminants in realist master paintings: Computer image analysis addresses a debate in art history of the baroque,” in [*Proceedings of the International Conference on Pattern Recognition*], **1**, 255–258 (2006).
- [8] Criminisi, A., Kemp, M., and Zisserman, A., “Bringing pictorial space to life: Computer techniques for the analysis of paintings,” in [*Digital art history: A subject in transition*], Bentkowska-Kafel, A., Cashen, T., and Gardner, H., eds., 77–100, Intellect Books, Bristol, UK (2005).
- [9] Stork, D., “Did early renaissance painters trace optical projections? evidence pro and con,” in [*Electronic Imaging, Vision Geometry XIII*], Latecki, L., Mount, D., and Wu, A., eds., 25–31 (2005).
- [10] Stork, D., “Did Georges de la Tour use optical projections while painting Christ in the carpenter’s studio?,” in [*Proceedings of SPIE Electronic Imaging*], Said, A. and Apostolopoulos, J., eds., **5685**, 214–219 (2005).
- [11] Stork, D., “Computer vision, image analysis and master art, Part I: Perspective and form,” *IEEE Multimedia* **13**(3), 16–20 (2006).
- [12] Stork, D. and Duarte, M., “Computer vision, image analysis and master art, Part III: Quantifying shape in realist art,” *IEEE Multimedia* **14**(1), 14–18 (2007).
- [13] Efros, A. and Freeman, W., “Image quilting for texture synthesis and transfer,” in [*Proceedings of the 28th Annual Conference on Computer Graphics and Interactive Techniques*], 341–346 (2001).
- [14] Varma, M. and Zisserman, A., “Texture classification: Are filter banks necessary?,” in [*Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*], **2**, 691–698 (2003).
- [15] van der Maaten, L. and Postma, E., “Texton-based texture analysis,” in [*Proceedings of the 19th Belgian-Dutch Conference on Artificial Intelligence*], Dastani, M. and de Jong, E., eds., 213–220 (2007).
- [16] van der Maaten, L. and Hinton, G., “Visualizing data using t-SNE,” *Journal of Machine Learning Research* **9**(Nov), 2431–2456 (2008).
- [17] Shahram, M., Stork, D., and Donoho, D., “Recovering layers of brush strokes through statistical analysis of color and shape: An application to van Goghs Self portrait with grey felt hat,” in [*Computer Image Analysis in the Study of Art, Proceedings of SPIE*], Stork, D. and Coddington, J., eds., **6810**, 68100D–1–8, SPIE (2008).
- [18] Jones, J. and Palmer, L., “An evaluation of the two-dimensional Gabor filter model of simple receptive fields in cat striate cortex,” *Journal of Neurophysiology* **58**(6), 1233–1258 (1987).
- [19] Simoncelli, E. and Freeman, W., “The steerable pyramid: A flexible architecture for multi-scale derivative computation,” in [*Proceedings of the IEEE 2nd International Conference on Image Processing*], 444–447 (1995).
- [20] Portilla, J. and Simoncelli, E., “A parametric texture model based on joint statistics of complex wavelet coefficients,” *International Journal of Computer Vision* **40**(1), 49–71 (2000).
- [21] Leung, T. and Malik, J., “Representing and recognizing the visual appearance of materials using three-dimensional textons,” *International Journal of Computer Vision* **43**(1), 29–44 (2001).
- [22] Daugman, G., “Uncertainty relation for resolution in space, spatial frequency, and orientation optimized by two-dimensional visual cortical filters,” *Journal of the Optical Society of America A* **2**(7), 1160–1169 (1985).
- [23] Hyvärinen, A., Hurri, J., and Hoyer, P., [*Natural Image Statistics: A probabilistic approach to computational early vision*], Springer, New York, NY (2008).

- [24] Olshausen, B. and Field, D., “Emergence of simple-cell receptive field properties by learning a sparse code for natural images,” *Nature* **381**, 607–609 (1996).
- [25] Osindero, S., Welling, M., and Hinton, G., “Topographic product models applied to natural scene statistics,” *Neural Computation* **18**(2), 381–414 (2006).
- [26] Kingsbury, N., “Complex wavelets for shift invariant analysis and filtering of signals,” *Journal of Applied and Computational Harmonic Analysis* **10**(3), 234–253 (2001).
- [27] Julesz, B., “Textons, the element of texture perception and their interactions,” *Nature* **290**, 91–97 (1981).
- [28] Schomaker, L., Franke, K., and Bulacu, M., “Using codebooks of fragmented connected-component contours in forensic and historic writer identification,” *Pattern Recognition Letters* **28**(6), 719–727 (2007).
- [29] Cheour, M., Ceponiene, R., Lehtokoski, A., Luuk, A., Allik, J., Alho, K., and Näätänen, R., “Development of language-specific phoneme representations in the infant brain,” *Nature Neuroscience* **1**(5), 315–353 (1998).
- [30] Sayre, K., “Machine recognition of handwritten words: a project report,” *Pattern Recognition* **5**(3), 213–228 (1973).
- [31] Kohonen, T., [*Self-organization and associative memory: 3rd edition*], Springer-Verlag New York, Inc., New York, NY (1989).
- [32] Shi, J. and Malik, J., “Normalized cuts and image segmentation,” *IEEE Transactions on Pattern Analysis and Machine Intelligence* **22**(8), 888–905 (2000).
- [33] Ng, A., Jordan, M., and Weiss, Y., “On spectral clustering: Analysis and an algorithm,” in [*Advances in Neural Information Processing Systems*], **14**, 849–856, The MIT Press, Cambridge, MA (2001).
- [34] Frey, B. and Dueck, D., “Clustering by passing messages between data points,” *Science* **315**, 972–976 (2007).
- [35] Tenenbaum, J., de Silva, V., and Langford, J., “A global geometric framework for nonlinear dimensionality reduction,” *Science* **290**(5500), 2319–2323 (2000).
- [36] Roweis, S. and Saul, L., “Nonlinear dimensionality reduction by Locally Linear Embedding,” *Science* **290**(5500), 2323–2326 (2000).
- [37] Lafon, S. and Lee, A., “Diffusion maps and coarse-graining: A unified framework for dimensionality reduction, graph partitioning, and data set parameterization,” *IEEE Transactions on Pattern Analysis and Machine Intelligence* **28**(9), 1393–1403 (2006).
- [38] van der Maaten, L., Postma, E., and van den Herik, H., “Dimensionality reduction: A comparative review,” Tech. Rep. TiCC-TR 2009-005, Tilburg University (2009).
- [39] Carreira-Perpiñán, M., “The elastic embedding algorithm for dimensionality reduction,” in [*Proceedings of the 27th International Conference on Machine Learning*], (2010).
- [40] Johnson, C., ed., [*Proceedings of the First International Workshop on Image Processing for Artist Identification*] (2007).
- [41] Turk, M. and Pentland, A., “Face recognition using eigenfaces,” in [*Proceedings of the Computer Vision and Pattern Recognition 1991*], 586–591 (1991).
- [42] Lee, D. and Seung, H., “Learning the parts of objects by non-negative matrix factorization,” *Nature* **401**(6755), 788–791 (1999).
- [43] Hyvärinen, A. and Oja, E., “Independent component analysis: algorithms and applications,” *Neural Networks* **13**, 411–430 (2000).
- [44] Hinton, G., “Training products of experts by minimizing contrastive divergence,” *Neural Computation* **14**(8), 1771–1800 (2002).
- [45] Biederman, I., “Recognition-by-components: A theory of human image understanding,” *Psychological Review* **94**(2), 115–147 (1987).
- [46] van der Maaten, L., “Learning a parametric embedding by preserving local structure,” in [*Proceedings of the International Conference on Artificial Intelligence and Statistics (AI-STATS)*], 384–391 (2009).
- [47] Burghouts, G. and Geusebroek, J., “Color textons for texture recognition,” in [*Proceedings of the British Machine Vision Conference*], **3**, 1099–1108 (2006).
- [48] Klein, A., Johnson, D., Sethares, W., Lee, H., Johnson, C., and Hendriks, E., “Algorithms for old master painting canvas thread counting from X-rays,” in [*Proceedings of the 42nd Asilomar Conference on Signals, Systems, and Computers*], (2008).

- [49] van Staalduinen, M., van der Lubbe, J., G. Dietz, F. L., and Laurentius, T., “Comparing X-ray and backlight imaging for paper structure visualization,” in [*Proceedings of the International Conference on Electronic Imaging and Visual Arts*], 108–113 (2006).
- [50] Bauch, J. and Eckstein, D., “Dendrochronological dating of oak panels of Dutch seventeenth-century paintings,” *Studies in Conservation* **15**(1), 45–50 (1970).
- [51] van Strydonck, M., Masschelein-Kleiner, L., Alderliesten, C., and de Jong, A., “Radiocarbon dating of canvas paintings: Two case studies,” *Studies in Conservation* **43**(4), 209–214 (1998).
- [52] Bach, F., Lanckriet, G., and Jordan, M., “Multiple kernel learning, conic duality, and the SMO algorithm,” in [*Proceedings of the 21st International Conference on Machine Learning*], (2004).