Tilburg centre for Creative Computing
Tilburg University
http://www.uvt.nl/ticc

P.O. Box 90153
5000 LE Tilburg, The Netherlands
Email: ticc@uvt.nl

Copyright © Laurens van der Maaten and Eric Postma 2009.

February 5, 2009

TiCC TR 2009–001

Identifying the Real Van Gogh with Brushstroke Textons

Laurens van der Maaten **Eric Postma**

TiCC, Tilburg University

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. We propose the brushstroke textons method that builds a codebook of textons, i.e., representative patches, of a collection of paintings and represents paintings in terms of texton histograms. In addition, the method visualizes the similarities between the histogram representations of paintings by employing a recently proposed dimensionality reduction technique. The effectiveness of the brushstroke textons method is demonstrated on a collection of digitized high-resolution reproductions of paintings by Van Gogh and his contemporaries. The results show a clear separation of paintings created by Van Gogh and those created by other painters. In addition, the period of creation is faithfully reflected in the visualization. These promising results show that the texton brushstroke method, in particular when extended with color and global textural features, offers a new tool for art historians in support of their study of paintings.

Identifying the Real Van Gogh with Brushstroke Textons

Laurens van der Maaten **Eric Postma**

TiCC, Tilburg University

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, artist identification is 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.
Previous work on artist identification and painting analysis focused on the assessment of color use in (Van Gogh) paintings using filter-based approaches [1], on capturing statistical information from segmented or outlined brushstrokes [11], and on wavelet analysis of the paintings [4]. In this paper, we present a new method – the *brushstroke textons method* – that builds an effective representation of the brushwork of paintings and visualizes the representations by means of a novel dimensionality reduction technique. We show the feasibility of the approach by applying it to a collection of paintings by Van Gogh and his contemporaries. Our results underline the relevance of intelligent image analysis and machine learning techniques for the cultural heritage domain.

2 The Brushstroke Textons Method

Given that the segmentation of individual brushstrokes from a painting is unfeasible [4], from a image analysis perspective, brushstroke analysis corresponds to the analysis of the texture of the painting. The wildly overlapping brushstrokes form a textural cue of the painter's handwriting. 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, the filters in the bank are generally selected in such a way that they respond to spatial frequencies at various scales and orientations [9]. Such 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 [17]. One of the main arguments against the use of filter-based approaches is that they 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 brushstroke texture, such details may correspond to individual hairs in the brush used by the painter that contain valuable cues for artist identification. The smoothing problem

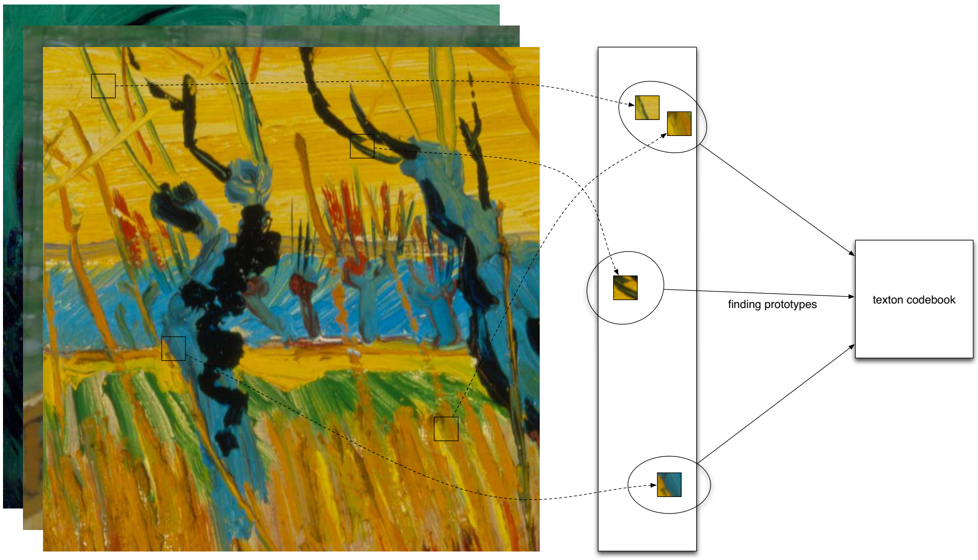

Figure 1: Illustration of the construction of a texton codebook. Patches are collected from random locations in each painting. Subsequently, prototypical brushstroke textons are identified by a clustering method and stored in a texton codebook.

may be resolved by employing filters with a small support such as in complex wavelet transforms [5], or by employing pixel-based texture representations such as texton histograms. In the brushstroke textons method, we combine the texton histogram approach of Varma and Zisserman [17] with a recently proposed visualization method based on dimensionality reduction [14]. In the following subsections, we separately discuss the texton histogram stage and the visualization stage of the brushstroke textons method.

2.1 Texton Histograms

In the first stage of the brushstroke textons method, paintings are represented in terms of texton histograms. A texton is the fundamental building block of texture (just like a grapheme is the fundamental building block of handwriting): we can view upon brushstroke texture as a superposition of brushstroke textons[1]. The texton histogram approach entails the use of normalized pixel-based texture representations and is, therefore, not hampered by the smoothing problem. In 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 that generates textons according to a painter-specific distribution over the textons in the codebook. For each painting, the texton distribution is measured by means of a texton histogram.

We construct the texton codebook as follows. We select $5,000$ square image patches from random locations in 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 [7], or the recently proposed affinity propagation [3]. 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.

Once the codebook is defined, each painting can be represented by means of a texton histogram. 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 correct for the

[1] Please 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.

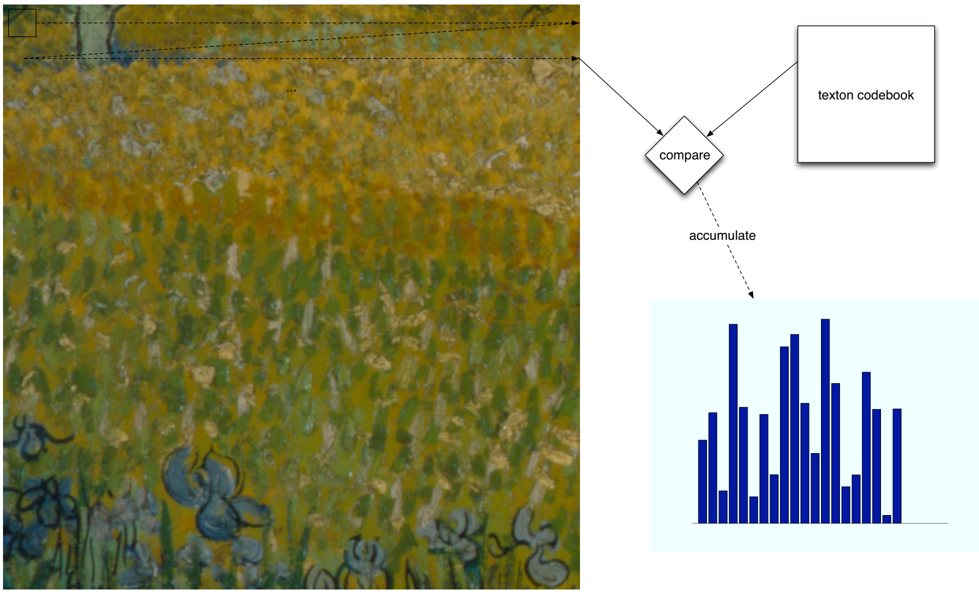

Figure 2: Illustration of the construction of a texton histogram. A window is slid over the texture image, and the histogram bin associated with the most similar codebook texton is incremented at each spatial location. After normalization, the texton histogram represents the relative number of times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

potential differences in size of the paintings. 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., texton sizes).

2.2 Visualization

The second stage of the brushstroke texton method is the visualization stage. Once relevant textural information is extracted from the paintings in the form of a texton histogram, a method is needed to visualize the similarity of textures or paintings to the art historian. The 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 in the data as possible, as a result of which similarities between texton histograms can be visualized in terms of distances between points in a scatterplot. The quality of the dimensionality-reduced visualization depends 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 datapoints may form a nonlinear manifold, and (2) PCA tends to ignore the local structure in the high-dimensional datapoints 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 datapoints that form a nonlinear manifold, e.g., Isomap [13], Local Linear Embedding [10], and diffusion maps [8]. We have reviewed these techniques [15] and found them to suffer from serious shortcomings that prohibit successful visualization of real-world data. Therefore, we opt

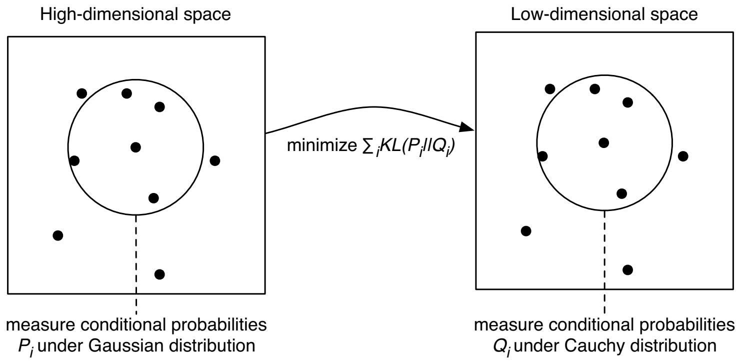

Figure 3: Illustration of the construction of a t-SNE map. The points in the low-dimensional map are arranged in such a way that the conditional probabilities in the map $q_{j|i}$ are similar to the conditional probabilities in the data $p_{j|i}$.

to use a novel dimensionality reduction technique, called t-Distributed Stochastic Neighbor Embedding (t-SNE) [14], 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 datapoints (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 datapoints x_i and x_j is given by the conditional probability

$$p_{j|i} = \frac{\exp\left(-\|x_i - x_j\|^2/2\sigma_i^2\right)}{\sum_{k\neq i}\exp\left(-\|x_i - x_k\|^2/2\sigma_i^2\right)}, \tag{1}$$

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 x_i and x_j by y_i and y_j, we define similar conditional probabilites between the points in the low-dimensional space. In order to address a problem known as the *crowding problem* [14], 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 datapoints y_i and y_j is defined to be

$$q_{j|i} = \frac{\left(1 + \|y_i - y_j\|^2\right)^{-1}}{\sum_{k\neq i}\left(1 + \|y_i - y_k\|^2\right)^{-1}}, \tag{2}$$

where again, we set $q_{i|i}$ to 0. The locations of the low-dimensional datapoints y_i are determined by minimizing the sum of the Kullback-Leibler divergences between the conditional probability distributions P_i and Q_i, i.e., by minimizing

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

with respect to the coordinates y_i of the points in the low-dimensional space using a gradient descent method[2]. t-SNE has been shown to perform very strong in the visualization of a variety of real-world datasets (see [14]). The construction of a t-SNE map is illustrated in Figure 3.

3 Evaluating the Brushstroke Textons Method

The texton histograms and the t-SNE dimensionality reduction algorithm constitute the brushstroke textons method which will be put to the test on a collection of paintings by Van Gogh and his contemporaries.

[2]Please note that minimizing this Kullback-Leibler divergence is identical to maximizing the log-likelihood of the data or to maximizing the cross-entropy of the distributions P_i and Q_i.

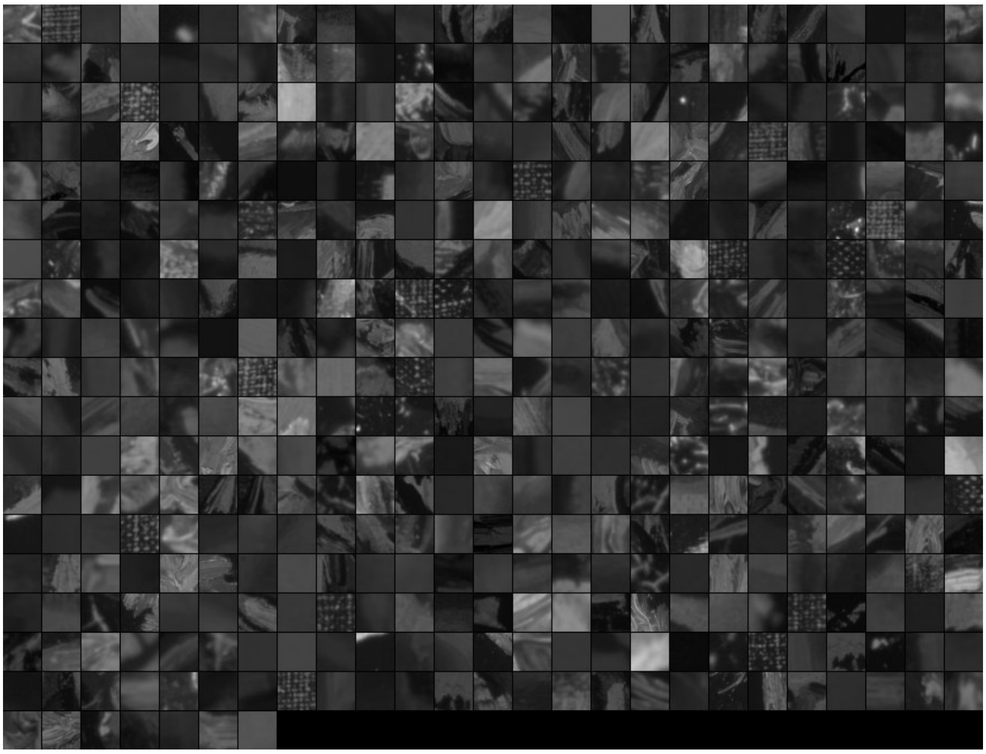

Figure 4: An example of one of the texton codebooks. This codebook was constructed using affinity propagation on textons of size $35\times$ 35 pixels.

3.1 Experimental Setup

We apply the approach on a recently released dataset of 117 high-resolution digital reproductions of paintings. 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 evaluate the brushstroke texton method, 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. Figure 4 shows one of the constructed texton codebooks. Altogether, the six texton histograms form feature vectors with approximately $3,000$ dimensions, which were first reduced to 50 dimensions using PCA. Subsequently, we 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 $\log_2(10)$).

3.2 Results

In Figure 5, we present one of the visualizations obtained with the brushstroke texton method. Each dot represents a single texton histogram (i.e., a single painting). The green dots represent Van Gogh's paintings, whereas the red dots those that are established non Van Gogh paintings. Paintings whose

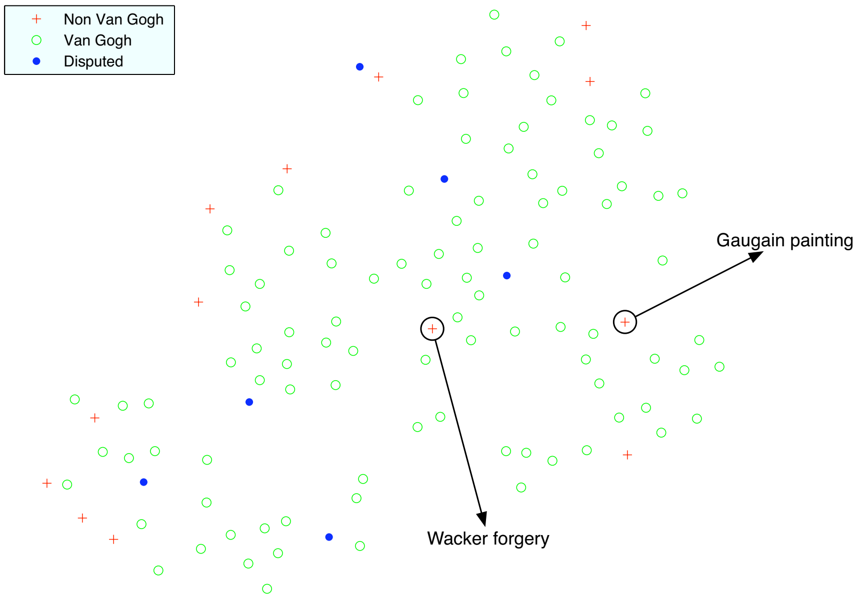

Figure 5: Visualization of the dataset of Van Gogh, non Van Gogh, and disputed Van Gogh paintings. The points in the scatter plot are labeled according to the authenticity of the paintings.

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 [4]: 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 the brushstroke textons method is completely data-driven. The visualizations obtained with the brushstroke texton method 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 the brushstroke textons method 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 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 brushstroke texton method to create visualisations of subsets of paintings to examine more subtle textural differences.

4 Discussion

The results presented above illustrate the potential of our method to support art historians in their analysis of paintings. Of course, the brushstroke textons method only offers an initial crude characterization

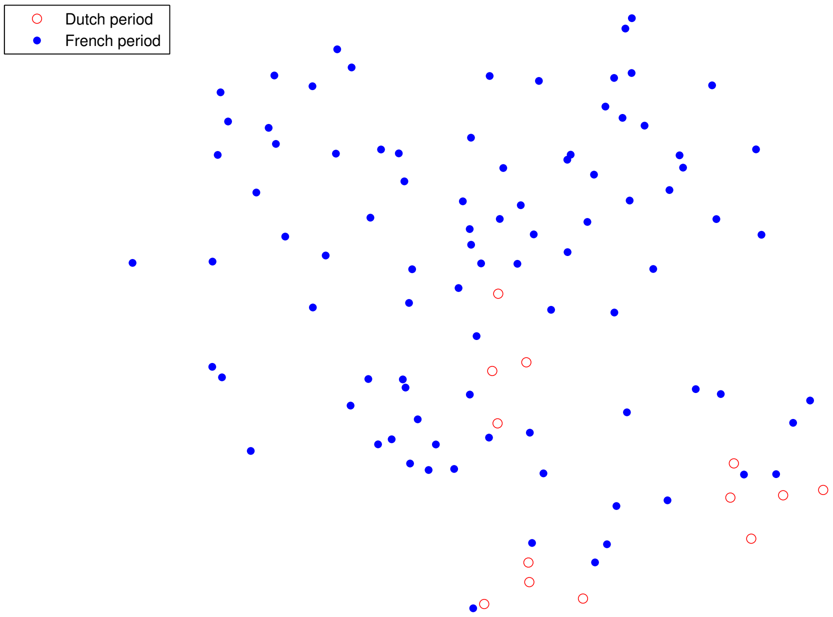

Figure 6: Visualization of the authentic Van Gogh paintings in the dataset. The points in the scatterplot are labeled according to the period in which they are painted.

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 features [4] or filter bank statistics such as those proposed in [9]. Moreover, the color use of Van Gogh (and more specifically, the use of complementary colors) should be captured in global painting features such as those proposed in [1]. Our texton-based approach may also be adpated to employ color information using one of the schemes proposed in [2].

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 brushstroke textons method 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 [6], paper analysis [16], analysis of lighting structure [12], and dendrochronological measurements).

In order to further improve the brushstroke textons method, more work is needed in order to present the results of the image analysis in more intuitive ways to the art expert. The true value for art historians is in the visualization and understanding of the visual characteristics (e.g., textons or configurations of textons) that give rise to the mapping. We envisage the future development of software that allows art experts to map and visualize subsets of paintings and selected regions of paintings. In that respect, the software incorporating the brushstroke texton method will become one of the many tools at the disposal of the art historian.

5 Concluding remarks

The brushstroke texton method, a novel computational method for artist identification, was presented that analyzes the brushstrokes in paintings by means of texton-based texture models, and visualizes the features obtained using t-SNE. The results indicate the sensitivity of the method to the textural cues underlying the authorship of a painting. An interactive version of the method will allow art experts to establish a quantitative assessment of the authenticity or date of creation of paintings. The authors hope to have made a good first step towards the development of systems that support art historians.

Acknowledgements

Laurens van der Maaten is supported by NWO-CATCH, project RICH (grant 640.002.401), and cooperates with RACM. 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.

References

[1] I.E. Berezhnoy, E.O. Postma, and H.J. van den Herik. Computer analysis of van Gogh's complementary colours. *Pattern Recognition Letters*, 28(6):703–709, 2007.

[2] G.J. Burghouts and J.M. Geusebroek. Color textons for texture recognition. In *Proceedings of the British Machine Vision Conference*, volume 3, pages 1099–1108, 2006.

[3] B.J. Frey and D. Dueck. Clustering by passing messages between data points. *Science*, 315:972–976, 2007.

[4] C.R. Johnson, E. Hendriks, I.E. Berezhnoy, E. Brevdo, S.M. Hughes, I. Daubechies, J. Li, E.O. Postma, and J.Z. Wang. Image processing for artist identification. *IEEE Signal Processing Magazine*, 25(4):37–48, 2008.

[5] N.G. Kingsbury. Complex wavelets for shift invariant analysis and filtering of signals. *Journal of Applied and Computational Harmonic Analysis*, 10(3):234–253, 2001.

[6] A.G. Klein, D.H. Johnson, W.A. Sethares, H. Lee, C.R. Johnson, and E. Hendriks. Algorithms for old master painting canvas thread counting from X-rays. In *Proceedings of the* 42^{nd} *Asilomar Conference on Signals, Systems, and Computers*, 2008.

[7] T. Kohonen. *Self-organization and associative memory:* 3^{rd} *edition*. Springer-Verlag New York, Inc., New York, NY, 1989.

[8] S. Lafon and A.B. Lee. 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.

[9] J. Portilla and E.P. Simoncelli. A parametric texture model based on joint statistics of complex wavelet coefficients. *International Journal of Computer Vision*, 40(1):49–71, 2000.

[10] S.T. Roweis and L.K. Saul. Nonlinear dimensionality reduction by Locally Linear Embedding. *Science*, 290(5500):2323–2326, 2000.

[11] R. Sablatnig, P. Kammerer, and E. Zolda. Hierarchical classification of paintings using face- and brush stroke models. In *Proceedings of the International Conference on Pattern Recognition*, pages 172–174, 1998.

[12] D.G. Stork. Optics and realism in Renaissance art. *Scientific American*, 291(6):76–84, 2004.

[13] J.B. Tenenbaum, V. de Silva, and J.C. Langford. A global geometric framework for nonlinear dimensionality reduction. *Science*, 290(5500):2319–2323, 2000.

[14] L.J.P. van der Maaten and G.E. Hinton. Visualizing data using t-SNE. Technical Report UTML TR 2008-001, University of Toronto, 2008.

[15] L.J.P. van der Maaten, E.O. Postma, and H.J. van den Herik. Dimensionality reduction: A comparative review. *Neurocomputing (submitted)*, 2008.

[16] M. van Staalduinen, J.C.A. van der Lubbe, F. Laurentius G. Dietz, and Th. Laurentius. Comparing X-ray and backlight imaging for paper structure visualization. In *Proceedings of the International Conference on Electronic Imaging and Visual Arts*, pages 108–113, 2006.

[17] M. Varma and A. Zisserman. Texture classification: Are filter banks necessary? In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, volume 2, pages 691–698, 2003.