

UNCOVERING ELEMENTS OF STYLE

Josephine Wolff, Maximiliaan Martens, Sina Jafarpour, Ingrid Daubechies, and Robert Calderbank

ABSTRACT

This paper relates the style of 16th century Flemish paintings by Goossen van der Weyden (GvdW) to the style of preliminary sketches or underpaintings made prior to executing the painting. Van der Weyden made underpaintings in markedly different styles for reasons as yet not understood by art historians. The analysis presented here starts from a classification of the underpaintings into four distinct styles by experts in art history. Analysis of the painted surfaces by a combination of wavelet analysis, hidden Markov trees and boosting algorithms can distinguish the four underpainting styles with greater than 90% cross-validation accuracy. On a subsequent blind test this classifier provided insight into the hypothesis by art historians that different patches of the finished painting were executed by different hands.

Index Terms— Wavelet Transforms, Hidden Markov Trees, Image Analysis, Machine Learning, Artwork Classification.

1. INTRODUCTION

In recent decades, mathematical methods of statistical analysis have contributed greatly to detecting artistic forgeries, denoising digital images, and uncovering underpaintings or other images beneath the surface of famous works of art. A very powerful mathematical tool for image processing has been wavelet analysis, in which special functions called wavelets are used to analyze images at different scales and extract important or identifying information from them. In 2008, a research team at Dartmouth used wavelet decompositions to authenticate a series of drawings by Netherlandish artist Pieter Bruegel the Elder [7]. Subsequent work by the same research team also showed that Pietro Perugino’s painting *Virgin and Child with Saints* had likely been completed by several different artists. Mathematical modeling suggested

that each painter brought a distinctive style or “signature” to the work. These same principles of wavelet analysis have also been applied to a number of other image processing problems, including detecting handwriting forgeries and recovering underpaintings by famous artists [6].

We follow prior stylistic analysis [4] in applying wavelet analysis, hidden Markov models and machine learning algorithms to a set of painting details attributed to Goossen van der Weyden. However our primary objective is different; it is to apply wavelet analysis broadly defined to develop a correspondence between the van der Weyden’s painting details and the style of the underpaintings or sketches beneath the paint.

This paper is organized as follows. Section 2 briefly describes the application of wavelets to brushstroke analysis, Section 3 describes how features are represented by hidden Markov models, and Section 4 describes the machine learning algorithms used for classification. Section 5 describes the data set in greater detail and Section 6 contains our results and conclusions.

2. WAVELET TRANSFORMS

Wavelet transforms allow us to split up a signal — or, in this case, an image — into several signals all corresponding to different frequency bands or scales. A wavelet transform divides up data into frequency components and looks at each of those components with a resolution that corresponds to its scale. Five primary properties of wavelet transforms make them especially useful for performing image analysis: locality, multi-resolution representation, edge detection, compression, and decorrelation [3].

Dual-tree complex wavelet transforms [8] consist of two real discrete wavelet transforms (implemented via wavelet filter bank trees) that produce the real and imaginary parts of the transform, respectively. The filter banks use combinations of band-pass and low-pass filters to separate the signal into several components characterized by specific frequency magnitude ranges; for each such range, six angular subbands (see Figure 1) provide better orientation selection than standard wavelet bases (see [4] for a more detailed explanation). When studying paintings, this directional orientation allows for a more detailed analysis of brushstroke direction, color patterns and local differences.

J. W. is with the Computer Science and Artificial Intelligence Laboratory, Massachusetts Institute of Technology, email: jwolff@mit.edu. M.M. is with Afdeling Beeldende Kunst, Ghent University, email: maximiliaan.martens@ugent.be. S.J. is with the Department of Computer Science, Princeton University, email: sina@cs.princeton.edu. I.D. is with the Department of Mathematics, Princeton University, email: ingrid@math.Princeton.EDU. R.C. is with the Department of Computer Science, Duke University, email: calderbk@cs.duke.edu. The work of R. Calderbank, I. Daubechies, and S. Jafarpour is supported in part by NSF under grant DMS 0701226, by ONR under grant N00173-06-1-G006, and by AFOSR under grant FA9550-09-1-0551.




Fig. 1. Left: Wavelets for the six orientation subbands in the 2-D dualtree complex wavelet transform, shown for two successive scales. Right: Quad-tree organization of wavelet coefficients in one orientation direction of the wavelet transform. For each spatial localization of a “parent coefficient” the next finer-scale subband has four “children wavelet coefficients”.

3. HIDDEN MARKOV TREES

In the wavelet transform of an image, most coefficients are very likely to be similar in size to their neighbors. This results in peaky, heavy-tailed, non-Gaussian marginal distributions for the wavelet representation of an image. Moreover, large or small values also have a tendency to persist across scales.

Standard signal processing methods that model coefficients as either independent or jointly Gaussian do not capture this behavior. As described in [1], hidden Markov tree models provide a concise, tractable, and efficient means of modeling the statistical dependences and non-Gaussian statistics of wavelet transforms of images.

In this model, the wavelet coefficients in each subband are assumed to be distributed according to a mixture of two Gaussians with very different variances; the hidden variable for each wavelet coefficient determines from which Gaussian distribution it is considered to be sampled. The “narrow” state corresponds to smooth regions, the “wide” state to edges and ridges in the image. The extent of persistence through scale, within the quad tree illustrated in Figure 1, of the state of the hidden variable governing a specific spatial location and orientation, is modeled by a transition probability matrix, the entries of which are efficiently estimated via the Expectation Maximization algorithm.

4. CLASSIFICATION

After using wavelet transforms and hidden Markov trees to develop a suitable mathematical representation for a painting, we need to classify the paintings into fixed categories based on the feature we are trying to distinguish or select for (in this case, the style of the underpainting). We examined a variety of different classifiers; we found boosting classifiers gave us the highest cross-validation accuracy.

Boosting classifiers are ensemble methods that use a weighted majority vote among a set of base classifiers (in our case decision stumps). The weights are determined iteratively

by challenging the learner by the *hardest* examples, i.e. those examples that are misclassified in the previous rounds.

Initially, we employed the AdaBoost algorithm introduced by Freund and Schapire [9]. For binary classification, we found AdaBoost to be the most effective learning algorithm; however, our attempts to classify data into four different classes were more successful when we employed LogitBoost. The LogitBoost algorithm, formulated by Friedman, Hastie, and Tibshirani and fully described in [2], is known to be more robust against noise and outliers.

5. GOOSSEN VAN DER WEYDEN DATA SET

We were provided with fifteen detail views taken from several paintings by Flemish painter Goossen van der Weyden, presently in the Museum of Fine Arts of Antwerp (Belgium). For each of these, we were also provided with an infrared reflectogram, which shows the lines of underpaintings or sketches, together with features from the visible top layer. The underpaintings differed markedly; based on their style as well as the nature of the materials used, we distinguished four categories (see Figure 2). The aim of the project was to classify the paintings, using wavelet analysis and machine learning boosting algorithms, into categories corresponding to these four different types of underpaintings. That is, we wanted to examine to what extent the different types of underpaintings translated into classifiable differences in the final paintings that were executed on top of them by Goossen van der Weyden and his students.

5.1. Classification of the Underpainting

The underpaintings in this dataset were made either using dry materials, like black chalk, or liquid ones, like black ink. We distinguish the following four distinct classes of underpaintings:

1. Woodcut-type parallel lines indicating shadows: oblique, parallel lines using a dry medium, like fine charcoal.
2. Parallel lines indicating volume: underdrawings done with a liquid agent, like a fine brush; finer lines than for type 1, curved to indicate volume.
3. Cartoon-like sketch of features: underdrawings done with a liquid agent; sketchier and much less precise than those in the previous two categories, with no, or very few, parallel lines.
4. Graphical indication of feature placement: underdrawing done with a liquid medium and painted in a very free hand; these are not even sketches; a stenographic notation rather than a precise indication of what the artist intends to paint.




Fig. 2. Left: A painting detail and the corresponding infrared reflectometric (IRR) view, showing both underpainting lines and top layer features. The underpainting is best visible where the artist changed his mind between underpainting and top layer; e.g. position of nose and mouth in zone circled in red, or details of coiffe in zone circled in yellow. (Contrast in the IRR view has been enhanced to make lines stand out.) Right: The four classes of underpainting. From top to bottom: visible top layer, IRR view, detail of IRR view with underpainting highlighted in red, underpainting only (as it must have appeared before the top layer was painted).

These four underpainting styles are illustrated in Figure 2. Because the underpaintings in the first two groups are characterized by many fine (almost) parallel lines next to each other, whereas the lines are drawn more loosely in the last two groups, we also considered two large classes, with class FL (Fine Lines) consisting of the painting details in groups 1 and 2, and class LD (Loose Drawings) of those in groups 3 and 4. These two classes were used for binary classification.

5.2. Binary Classification Results

For the binary, or two-group, classification problem the AdaBoost algorithm yielded the most promising results; its performance improved considerably with the number of iterations. The results are summarized in Figure 3; the accuracy reaches a plateau of around 93% after about 1000 iterations of the algorithm.

5.3. Four-group Classification Results

AdaBoost was less successful in a direct classification of the data into its four different groups (as opposed to two successive binary classifications). However, with LogitBoost we reached a classification accuracy of 90% in about 400 iterations. (In contrast, AdaBoost achieved only 56.71% accuracy.)




Fig. 3. The table at left shows the percentage of data accurately classified into two different classes by the AdaBoost algorithm, run with varying numbers of iterations; the plot on the right shows accuracy values reaching a plateau around 1000 iterations.

The convex combination of base hypotheses used by AdaBoost and LogitBoost to construct their classifier contains hundreds of features, mostly biased towards the highest scales (those with the finest detail) in both the binary and the four-group classification results. This emphasis on the finest scales for the task of distinguishing between the different underpainting classes could suggest that the differences in execution (and possible skill-level) for these painting details lies in the finest details of the work, rather than in broader stylistic trends, which would be apparent at coarser scales.




Fig. 4. The table at left shows the percentage of data accurately classified into four different classes by the LogitBoost algorithm, run with varying numbers of iterations; the plot on the right shows accuracy values reaching a plateau around 400 iterations.

5.4. Blind Data Set Results

Apart from cross-validation, we also used a blind dataset in order to check the classifier constructed by LogitBoost in the supervised learning task of Section 5.3. We were provided with an additional seven unlabeled painting details by van der Weyden belonging to the same four classes of underdrawings. For each of these seven, the LogitBoost classifier determined

		Probability (in %)						
Painting detail	Class	A	B	C	D	E	F	G
1		0	8	4	2	8	1	8
2		3	31	6	13	18	14	42
3		96	28	82	60	53	37	49
4		1	33	8	25	21	48	1
Correct Class		3	2	3	4	1	4	2
Success?		V	V	V	V	X	V	V

Fig. 5. Classification results for seven GvdW images in the blind testset. Columns A through G correspond to the seven “new” painting details; in each column the most confident choice of class, as given by the LogitBoost classifier, is indicated with a double box; the second most confident choice (if its likelihood reaches into the double digits) by a single box. A large green V at the bottom of the column indicates the most likely choice is indeed correct; a smaller green v that the second most likely choice is the correct one. Example E is completely misclassified; its “correct” label is in fact LogitBoost’s *least* likely class for this example. (See, however, Figure 6).



Fig. 6. Left: top layer of painting detail E of the blind test in section 5.4; Middle: IRR view; Right: highlighted underpainting lines. These underpainting lines are reminiscent of classes 3 or 4 rather than 1.

the probabilities p_1, \dots, p_4 that it belonged to the four classes. Only after the result was obtained did we compare with the classification determined independently by one of us (M.M.) who had not participated in the machine learning part of this experiment. Of these seven, the algorithm scored the correct class as most likely in three cases, and second most likely in another three cases. (See Figure 5) In the remaining case the algorithm’s “*least* likely” answer corresponded to the independently attributed label.

A closer examination showed, however, that the “ground truth” label did *not* correspond to the correct description of the IRR view in the anomalous case, as shown in Figure 6. The IRR for this instance corresponds more to class 3 or 4 (and not to class 1, despite the “correct” label); these are, in fact, the classes labeled by the LogitBoost classifier as most likely for this example.

6. CONCLUSIONS

Our results in the classification of van der Weydens painting details in the supervised learning experiment described here, consistent with the style of their underdrawings, show great promise. Further work is planned to investigate and validate this in more depth. It is striking that in Figure 5, Class 3 dominates among the most likely picks by LogitBoost; further work is needed to make sure this is not due to a confounding factor. The authors plan to extend this study, using many more examples of paintings by GvdW and others, and replacing the Hidden Markov Tree model by a method that retains more spatial localization.

7. REFERENCES

- [1] H. Choi, J. Romberg, R. Baraniuk, and N. G. Kingsbury, *Hidden Markov Tree Modeling of Complex Wavelet Transforms*, Proc. International Conference on Acoustics, Speech and Signal Processing (ICASSP), 2000.
- [2] J. Friedman, T. Hastie, and R. Tibshirani, *Additive Logistic Regression: A Statistical View of Boosting*, Annals of Statistics, vol. 28, 1998.
- [3] A. Graps, *An Introduction to Wavelets*, Computing in Science and Engineering, vol. 2, no. 2, pp. 50-61, Summer, 1995.
- [4] S. Jafarpour, G. Polatkan, E. Brevdo, S. Hughes, A. Brasoveanu, and I. Daubechies, *Stylistic Analysis of Paintings using Wavelets and Machine Learning*, European Signal Processing Conference, 2009.
- [5] G. Hulin de Loo, *An Authentic Picture by Goossen van der Weyden and the Legend of S. Dymphna from Tongerlo*, The Burlington Magazine for Connoisseurs, Vol. 22, No. 115, Oct. 1912: 26-32.
- [6] B. Lytle and C. Yang, *Detecting Forged Handwriting with Wavelets and Statistics*, Rose Hulman Undergraduate Mathematics Journal, Terre Haute, IN, 2006.
- [7] S. Lyu, D. Rockmore, and H. Farid, *Wavelet Analysis for Authentication*, <http://www.scientificcommons.org/42465948>, 2008.
- [8] I. W. Selesnick, R. G. Baraniuk and N. G. Kingsbury, *The Dual-Tree Complex Wavelet Transform*, IEEE Signal Processing Magazine, vol. 22, no. 6, pp. 123-151, 2005.
- [9] Y. Freund and R. Schapire, *A decision-theoretic generalization of on-line learning and an application to boosting*, Journal of Computer and System Sciences, vol. 55, no.1, pp.19-139, 1997.