Boosting multi-feature visual texture classifiers for the authentication of Jackson Pollock's drip paintings

Mahmoud Al-Ayyoub,[a] Mohammad T. Irfan[a] and David G. Stork[b]

[a]Department of Computer Science, Stony Brook University, Stony Brook NY 11794 USA
[b]Ricoh Innovations, 2882 Sand Hill Road Suite 115, Menlo Park CA 94025 USA

ABSTRACT

Early attempts at authentication Jackson Pollock's drip paintings based on computer image analysis were restricted to a single "fractal" or "multi-fractal" visual feature, and achieved classification nearly indistinguishable from chance. Irfan and Stork pointed out that such Pollock authentication is an instance of *visual texture recognition*, a large discipline that universally relies on *multiple* visual features, and showed that modest, but statistically significant improvement in recognition accuracy can be achieved through the use of multiple features. Our work here extends such multi-feature classification by training on more image data and images of higher resolution of both genuine Pollocks and fakes. We exploit methods for feature extraction, feature selection and classifier techniques commonly used in pattern recognition research including Support Vector Machines (SVM), decision trees (DT), and AdaBoost. We extract features from the fractality, multifractality, pink noise patterns, topological genus, and curvature properties of the images of candidate paintings, and address learning issues that have arisen due to the small number of examples. In our experiments, we found that the unmodified classifiers like Support Vector Machines or Decision Tree alone give low accuracies (60%), but that statistical boosting through AdaBoost leads to accuracies of nearly 75%. Thus, although our set of observations is very small, we conclude that boosting methods can improve the accuracy of multi-feature classification of Pollock's drip paintings.

Keywords: Jackson Pollock, drip painting analysis, machine learning, pattern recognition, texture recognition, image processing, fractal analysis, multi-fractal analysis, AdaBoost

1. INTRODUCTION

Jackson Pollock's action paintings of driped and poured liquid paint are important works in Abstract Expressionism. There are many works of doubtful attribution and outright fakes, including one celebrated work—*Teri's find*, purchased for $5 in the early 1980s—as well as a cache of 32 works found in 2002 by Alex Matter, son of one of Pollock's friends. Traditional authentication methods based on studies of paint, support, priming, signatures and provenance (the documentary record of the sequence of ownership of a work) are not always definitive—for Pollock on indeed many other artists—and thus any additional, complementary evidence could be useful in authentication studies. To this end, Taylor and his colleagues introduced a box-counting algorithm to estimate fractal and scale-space signatures of Pollock's works.[1–3] This team reported that such mathematical "signatures" generally differed between genuine and fake Pollocks and thus that their method could be used as part of authentication protocol. The fractal approach has been criticized on a number of grounds including claims that fractals could be uninformative or "useless," that the range of spatial scales used by Taylor et al. was too narrow to infer true fractal properties, and that their image processing precluded the estimation of true fractal properties.[4, 5] Irfan and Stork, however, answered those objections both theoretically and experimentally, and demonstrated that the classical paradigm of texture classification based on *multiple* visual features and statistical estimation of classifier parameters, leads to improved recognition accuracy.[6]

Although the Irfan/Stork empirical results were encouraging, they were nevertheless preliminary. Our current project builds upon that work by training on more image data (of both genuine Pollocks and fakes) and images of higher resolution. We employ feature extraction, feature selection and classifier selection techniques commonly used in pattern recognition research and several supervised classification frameworks, such as support vector machines, decision trees, and boosting through AdaBoost, to the Pollock classification problem.

Send correspondence to David G. Stork, artanalyst@gmail.com.

In Sect. 2 we review the motivation for multiple visual features in texture recognition and describe some of the visual features we employ in our study. In Sect. 3 we discuss the classification algorithms and methods applied to Pollock authentication. We conclude in Sect. 4.

2. VISUAL FEATURES FOR POLLOCK AUTHENTICATION

Texture classification is a very well developed branch of pattern classification, and perhaps the most important general result from the field is that it is essential that *multiple* visual features be used for classification. An overview of research in visual texture recognition noted that of eighteen leading methods one method used as few as five features but others used up to 47 features, with the average being 18 features. [7, p. 279] Clearly a texture recognition program (for Pollock authentication) based on but a *single* visual feature is incommensurate with norms and vast body of research in the relevant subdiscipline of pattern classification.[6,8]

It is quite difficult to predict which features will and will not be useful in any particular texture classification problem, and Pollock authentication is no different. In designing a classifier for the authentication task, we would like to incorporate features that would give us good intra-class similarities (that is, the genuine Pollock paintings, in general, have similar feature values. We further desire that the features would give good inter-class distance (that is, in general, the feature values would be different for Pollock and non-Pollock paintings). We tried various features, including a variant on the fractal feature explored by Taylor and others: estimates of fractal behavior, multifractality, pink noise spectrum, topological genus, and curvature properties of the images. Of course other features could be used as well.

The first, preprocessing, step is to separate the color layers of the painting being considered.[1,4,6] For practical reasons we do not consider all the different colors used in the painting, but rather consider only a few colors that have been mostly used in the painting. For example, Fig. 1 shows Pollock's *Number 1* (1948), and Fig. 2 shows the black layer of this painting obtained by color thresholding.

Fractal behavior

As mentioned above, Richard Taylor and his colleagues argue that Pollock's paintings may show some mathematical or "fractal" order.[9] The use of fractal properties in authenticating Pollock's paintings has its proponents as well as its opponents.[4] In pattern recognition and machine learning, the strict mathematical definitions of some terms come second to the usefulness of the relevant features. Furthermore, the inclusion of some feature—even a "useless" feature—contaminated or not, may improved classification, as pointed out by Irfan and Stork.[6,8] Thus, it is reasonable to extract some measure of fractal behavior as a feature of a drip painting. However, instead of using the fractal dimensions, we use the extent to which a painting shows fractal properties, as described below.

The Taylor et al. fractal analysis works on each isolated color layer roughly as follows. We overlay each color layer with a grid of squares, and count the number of non-empty squares, that is the number of squares containing some amount of paint in them. We vary the grid size from the smallest possible to the largest possible, and plot the count of non-empty squares against the side length of the squares in the log-log scale. If the color layer has an exact fractal property then the plot will be a straight line; the slope of the straight line is called its fractal dimension. This method is commonly known as the *box counting method*. Taylor and his colleagues have shown that the color layers in some of Pollock's drip paintings reveal two slopes, one at a lower resolution (when the side length of the squares is roughly 1.5 cm or smaller) and the other at the higher resolutions. Taylor has named these two fractal properties *drip dimension* and *levy dimension*, respectively.

Figure 3 shows the box counting plot for this black layer of *Number 1*. The slope of the best-fit line fit gives a measure of fractal dimension for the black layer of the image. We do not use the slope directly as a feature of the painting but instead the norm of the residuals of the line fit, normalized by the number of data points in the plot. For example, the norm of the residuals of the line fit (before normalizing) shown in Fig. 3 is 0.0465. We take the average of this quantity over all color layers being examined and use the resulting scalar as the fractality feature in our classifiers.

Rockmore and his colleagues have reported that he color layers in some of Pollock's drip paintings show multifractal properties.[10,11] We follow their lead and extract a multifractality feature from an isolated color layer for our classification task. As described above, we overlay each color layer with boxes in box counting method. However, instead of counting how many boxes have any paint in them, we measure Shannon information quantity

Figure 1. Jackson Pollock, *Number 1* (1948), 173 × 264 cm, oil on canvas, Museum of Modern Art.

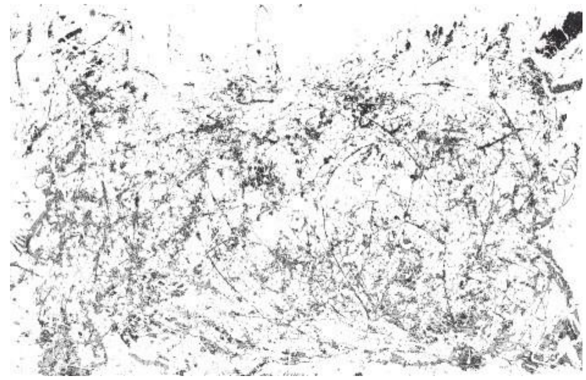

Figure 2. The black color layer of *Number 1* obtained by simple, semi-automatic thresholding.

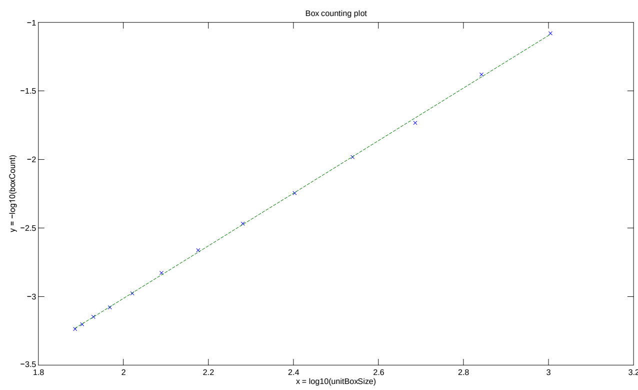

Figure 3. The box counting plot for the image shown in Fig. 2. The slope of the best-fit line is the image's fractal dimension or levy dimension of that black layer. In our experiments, we use more data points than shown here and use the norm of the residuals of the line fit, scaled by the number of data points, as our fractality feature.

of each box and then sum that quantity over all the boxes to yield the Shannon entropy measure for that box size. We then plot this Shannon entropy measure against the box sizes in log-log scale. The resulting plot will give us a straight line if the corresponding color layer shows multifractal properties, the slope of that straight line is called the *information dimension* of the color layer. We do not adopt the information dimension as a feature but instead compute the extent to which a color layer shows the information dimension property by using of the norm of the line fit (normalized by the number of data points in the plot). We take the average of this quantity for the color layers of the painting under consideration, and this gives us the single scalar multifractality feature of the painting.

The plot for the black color layer of *Number 1* is shown in Fig. 4. The best line fit for the plot is the information dimension of the corresponding image. However, we use the norm of the residuals of the line fit, which is 0.0765 in this case, to derive the multifractality feature of the painting.

Topological genus

We have analyzed the topology of several Pollock paintings. First, we have converted the input image to grayscale. Then we have separated the relatively dark layer of the image and applied Wiener filtering to smooth the image. We have computed the Euler number e as the difference between the number of objects and the number of holes in the resulting image. This number e is then used as the topological genus feature of the image. Typically, the Pollock paintings show larger values (both negative and positive) for this feature compared to the non-Pollock ones, although there are many exceptions.

Oriented energy

Analysis of oriented energy has revealed some interesting properties of several Pollock paintings. Taking the

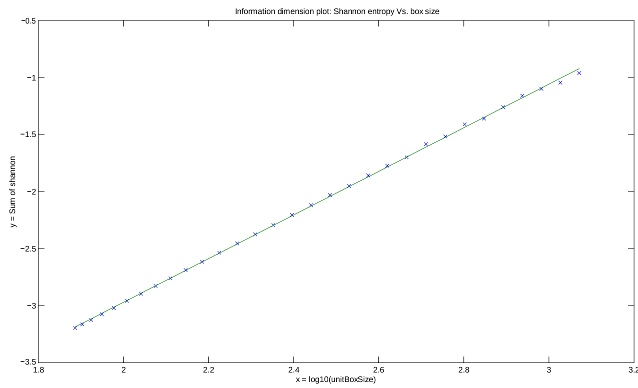

Figure 4. The Shannon entropy measures Vs. box sizes plot in log-log scale for the image shown in Fig. 2. The slope of the line is termed the information dimension of the color layer. The norm of the residuals of this line fit is 0.0765, which is used in computing our multifractality feature.

convolution of the Gabor filter of an input image at a certain orientation with the input image, we get a measure of the "use" of that orientation in the input image. We have analyzed several Pollock and non-Pollock paintings at every 10 degree angle, beginning at the vertical direction as the 0 degree angle. For each orientation we have measured the fraction of pixels that appears in that orientation. The plot of this measure against the orientations are nearly symmetric with respect to the vertical direction. Figure 5 shows the plot for Pollock's *Blue poles*.

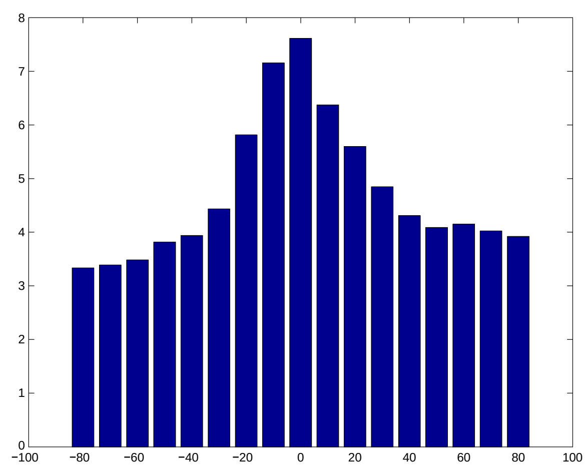

Figure 5. The "energy" in a Gabor-like filter versus orientation for a color layer in Pollock's *Blue Poles*.

However, many other Pollock paintings do not have similar characteristics. Thus, although oriented energy plots can be informative in analyzing the angular properties of abstract paintings, we have not observed any good intra-class similarities among genuine Pollock paintings with respect to the oriented energy plots, and therefore, we have not used this feature in our classification task.

Connected components

For each of the color layers of the input image, we have analyzed the histogram of the size of connected components of that layer. We have found that for both Pollock and non-Pollock paintings the histogram is weighted toward the smaller size of the connected components, even after discarding very small components. This may be due to the occlusion among different color layers, which would contribute to the breaking of the same stroke into multiple ones. Accurate recovery of a color layer from occlusion could possibly give us different results, a matter for future research.[12]

We have plotted the histogram of connected component sizes in a logarithmic scale. That is, first, we first compute $\lfloor log_2(r \cdot c) \rfloor$ bins, where $r \times c$ is the image size. Then for each component we have computed the number $\lfloor log_2(s) \rfloor$, where s is the size of that component, and we have thrown the component in the $\lfloor log_2(s) \rfloor$-th bin. In the end we have plotted the frequency of components in those bins. This feature, taken alone, was not discriminative between Pollock and non-Pollock images.

Pink noise spectrum

Pink noise, also known as $1/f$-noise, has been attributed to the pleasant perception of musical melodies. Recently it has been shown that Pollock's paintings also have underlying pink noise patterns at low scales (below 10-30cm).[13] The pleasant visual perception of Pollock's paintings may be attributed to this property. We have extracted two features of drip paintings based on how well the paintings show pink noise patterns. To compute these features, we first convert the image of a painting to an 8-bit grayscale version I. Then the image I is randomly sampled in horizontal direction, left to right and then right to left and so on, to obtain a large number of samples (at most 100,000 in our experiment). At each step of this horizontal sampling, we sample a number from the Gaussian distribution with mean 0 and variance 1, round that sampled number, and use this rounded number as the vertical displacement for the next horizontal coordinate(cf. Fig. 6). The sampling procedure gives us a sequence of sampled pixels x_i, $1 \leq i \leq N$, where N is the total number of samples. The value of each element in this sequence is in between 1 and 256, both inclusive. Let $\bar{x}$ be the mean of this sequence. We then compute an integrated sequence y_i, $1 \leq i \leq N$, as follows:

$$y_i = \sum_{j=1}^{i} (x_j - \bar{x}). \tag{1}$$

Let us select an interval length L for the sequence y_i. For each interval in the sequence y_i, we determine the *trend* of the subsequence in that interval by computing a line fit. This gives us a new N-element trend sequence, where the i-th element v_i comes from the line fit of the interval containing y_i.

We then compute a *detrended* sequence z_i, $1 \leq i \leq N$, as $z_i = y_i - v_i$. The *fluctuation* in the sampled sequence for interval length L is defined as:

$$F(L) = \frac{1}{N} \sum_{i=1}^{N} z_i^2. \tag{2}$$

We can vary the interval length L and plot $\log F(L)$ against $\log L$. We say that I shows pink noise property if this plot is a straight line with a slope of 1. In our classification task, we have worked with both the norm of the line fit and the slope to compute the pink noise features of the paintings.

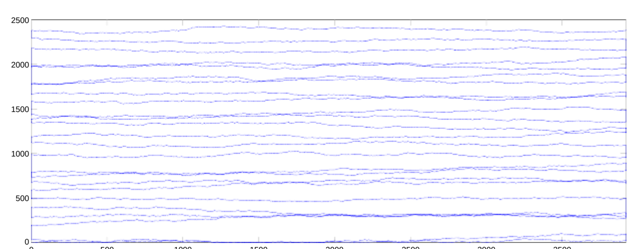

Figure 6. The horizontal sampling of *Number 1*.

Figure 7 illustrates the pink noise property of *Number 1*. The dimensions of the image of the painting are 3908×2538. Using the sampling procedure we have obtained 100,000 samples in total. We have varied the length of intervals L from 10 to 136 (the upper limit corresponds to approximately 10cm of the original painting). In the plot of Fig. 7, the values of L has been increased in a geometric progression. The line fit has a norm of 0.0712

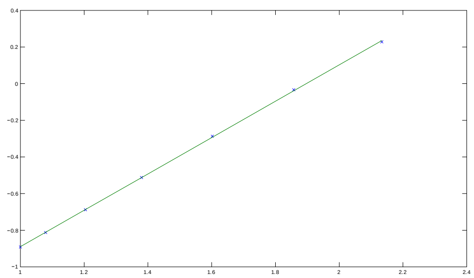

Figure 7. Pink noise pattern in Pollock's Number 1: the norm of the line fit is 0.0712, and the line fit has a slope of 1.06.

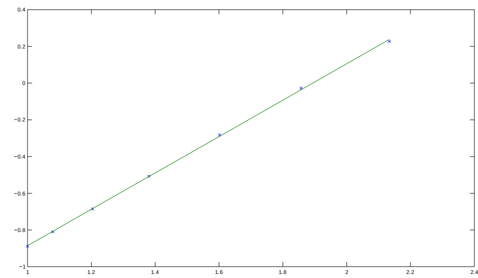

Figure 8. Isotropy property of *Number 1*: after rotating the original painting 90^0 clockwise, similar pink noise pattern as the original image has been found.

and a slope of 1.06. These two values have been used in extracting the pink noise features of the painting. The details are given in the next subsection.

Isotropy

We use the term isotropy with reference to the pink noise pattern just described. A painting is isotropic if it shows pink noise properties both in its original orientation and the orientation obtained by rotating the painting 90^0 clockwise. It has been reported that Pollock's drip paintings are isotropic.[13] We extract the isotropy feature from the painting by rotating the painting and then applying the same procedure as described above, only the number of horizontal strips and the total number of samples being different. We then take the average of the two slopes (for two orientations) to use it as a feature. We also take the average of the norms of the two line fits as another feature.

Curvature

Analysis of curvature can significantly contribute to the authentication studies of Pollock's paintings. Besides the artistic intuition of the painter, many physical factors might be held responsible for the curvatures of certain drips. Examples of such physical factors are the height from which the drip has been originated, the motion of the hand of the painter, and even the length of the hand of the painter. Of course, the directional preference, if any, of the painter would be a prime factor in creating drips of certain curvature.

To analyze the curvature feature of a painting, we use the Curevelet tool,[14] which is a non-adaptive technique for multi-scale representation of curves. The first step is to separate the color layers and choose the most prominent layer for analysis. We then run the Curvelet tool on the chosen layer to obtain the Curvelet coefficients. These coefficients usually lie in the range $[-1, 1]$, which we divide into 200 intervals, each of length 0.01. We then compute the distribution of these coefficients over the intervals, which appears to be a normal distribution with zero-mean. Thus, we report the standard deviation of this distribution.

3. EXPERIMENTAL CLASSIFICATION RESULTS

First, we focus on collecting example images of authentic Pollock paintings and drip paintings by other artists, for example, Taylor Winn, Marcel Barbaeu, and Swarez. It should be noted here that the image resolution (or magnification) of the paintings can have an effect on the classification. Ideally, we would like to have the same amount of magnification for all the paintings. Moreover, the magnification should be high enough for the fractal, multifractal, and pink noise characterizations to apply to a good degree of accuracy. However, our research so far indicates that sample images having these desirable resolution properties are very much hard to obtain but xxx.

3.1 Feature Selection

We have a collection of proposed features for the authentication tasks. Our preliminary investigation showed that not all of these features improve the classification significantly.[6] Thus, we are interested in selecting a small set of important features and avoids overfitting.[15]

Initial inspection of preliminary results suggests that the following six features are the most useful.

1. Fractalility
2. Multifractalility
3. Pink noise pattern (slope of the line fit)
4. Pink noise pattern (norm of the residuals of the line fit)
5. Topological genus
6. Curvature profile

Thus, we will limit our discussion below to these features.

3.2 Classifier Design

The preliminary work of Irfan and Stork worked with very simple classifiers like the nearest neighbor and the perceptron to settle the debate on using certain features.[6] We extend that work by actually building a classifier for the authentication task as a whole.

Support Vector Machine (SVM)
We start with Support Vector Machines to give some intuition on the hardness of this problem. Given the complete data set of 42 paintings, each represented by the six features discussed above, an SVM with a linear kernel chose 38 (out of 42) examples to be support vectors (among those were 19 out of Pollock's 21 paintings). We used a quadratic kernel for the SVM chose 24 paintings (including 13 Pollock paintings) to be support vectors. These results suggest that each painting show enough unique characteristics to be interesting on its own right, and it is very hard to find a small subset of representative examples based on which predictions can be made.

We performed cross validation to estimate the classification error of SVM with a quadratic kernel. The leave-one-out method gives an error estimate of 26.19%, while 10-fold and 20-fold cross validations give error estimates of 25.81% and 26.05% (averaged over 100 runs), respectively.

Decision Tree (DT)
Taken the small size of the data set into account, decision trees would be a natural choice of a classifier. Now, to control the complexity of the DT, we use a *splitting parameter*, which bounds the number of non-leaf nodes in the DT. We use the well-known Classification and Regression Tree (CART) algorithm for building a decision tree, whose size is bounded by the splitting parameter, to fit the training data.

Once again, we perform cross validation to estimate the error of this decision tree approach, while varying the splitting parameter. Figure 9 shows the error estimates that we get when using three different types of cross validation, viz., 10-fold, 20-fold, and leave-one out, for different values of the splitting parameter. We note that each data point in the plot is the average of 100 runs of the algorithm. We have two important observations on the decision tree-based approach. First, it is clear from Fig. 9 that the decision tree alone does not even reach 70% accuracy. In fact, the leave-one-out cross validation gives the lowest error estimate of 35.71% when the splitting parameter is 5. Secondly, we can see the effect of the complexity of the decision tree on the classification error. We have high error estimates at the two extrimities of the complexity of the decision tree. We have the lowest error estimates when the number of non-leaf nodes is in between 3 and 5.

AdaBoost with decision trees as the weak learner. Since the results with decision trees were not satisfactory, we have resorted to the famous technique of AdaBoost with the hope of a better classifier. We use a version of AdaBoost, called ModestBoost,[16] which has been claimed to have more generalization capability and better stability than the two well-known versions of it, viz., RealBoost and GentleBoost. We have used the same CART method (as the weak learner) as described above, with the additional capability of dealing with weighted observations. We can control the complexity of this weak learner by varying the splitting parameter, which directly controls the maximum size of the decision trees. We also vary the maximum number of iterations the AdaBoost algorithm invokes the weak learner.

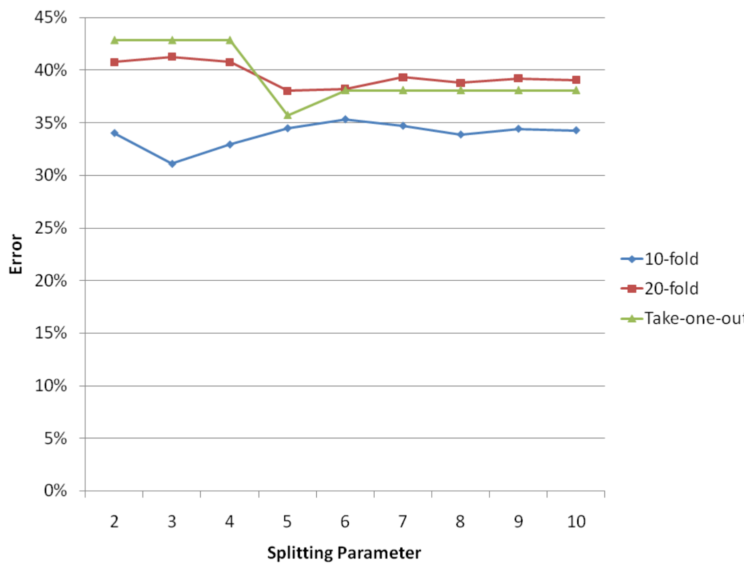

Figure 9. Classification error in a decision tree as a function of tree complexity (splitting parameter).

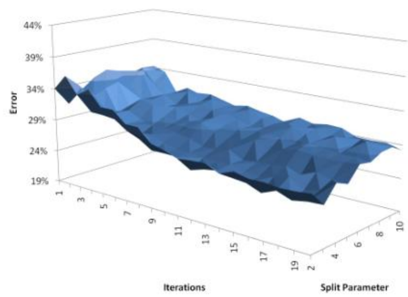

Figure 10. AdaBoost with 10-fold cross-validation

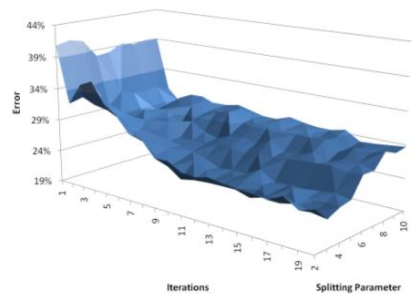

Figure 11. AdaBoost with 20-fold cross-validation

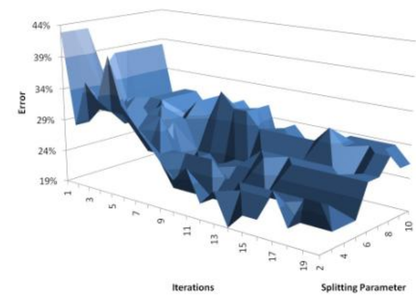

Figure 12. AdaBoost with leave-one-out cross-validation

Figures 10, 11, and 12 show the error estimates when AdaBoost is repeatedly calling DT with different values for the splitting parameter. As a general trend in the plots, the error tends to decrease with the increase in the maximum number of iterations (number of times DT, the weak learner, is invoked). Our experiments show that we can achieve higher than 80% accuracy (measured using the leave-one-out cross-validation) for certain values of the DT splitting parameter and the AdaBoost maximum number of iterations parameter. Furthermore, the accuracy is always close to 80% when the splitting parameter is not too small and not too large and the maximum number of iterations is close to 20.

Figure 13 shows the effect of the complexity of the weak learner on the overall error estimate of AdaBoost. For all three cross-validation methods, we find that if the splitting parameter is too low or too high then the error estimate is large. Obviously, high values of this splitting parameter lead to overfitting.

4. CONCLUSIONS

We have applied several supervised classification techniques for image-based authentication of the paintings of Jackson Pollock. One of the main challenges in classification—including texture based recognition of Pollock's drip paintings—is to find "good" features. We have explored six features in this project, none of which taken alone distinguish genuine from fake Pollock's painting. Nevertheless, when used in conjunction, these features yield classification better than chance, at least on our restricted dataset. In our experiments, we have found that the simple classifiers like SVM or DT alone do not give good only modest performance. In particular, DT alone gives error estimates of roughly 40%. However, when we have used DT as a weak learner in AdaBoost, we have

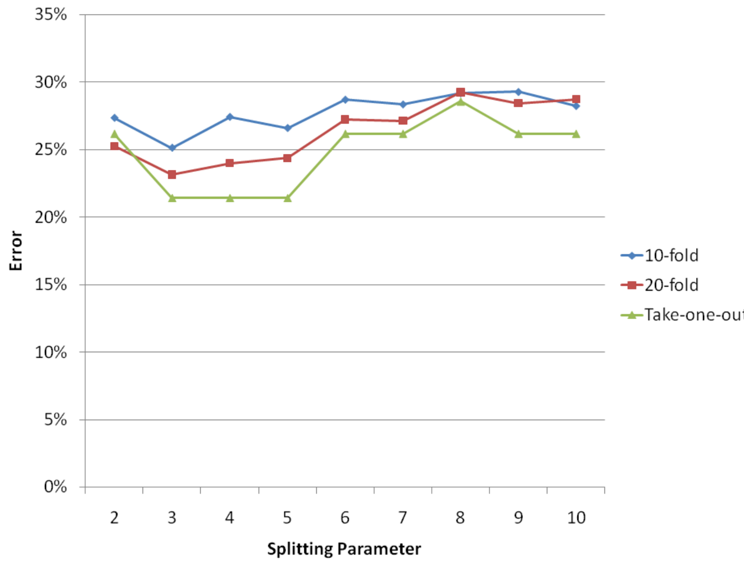

Figure 13. AdaBoost can improve the performance of a weak learner, especially at lower-complexity models (smaller splitting parameter.

obtained error estimates of roughly 25%. Thus, although our set of observations is very small, we conclude that boosting methods can significanlty improve Pollock classification accuracy. This result comports with theoretical studies and a wide range of empirical studies showing the benefit of boosting techniques.[17]

There are several future directions that could lead to improved image-based authentication methods:

- A larger corpus of high-resolution images of both Pollock and non-Pollock paintings is essential before image-based Pollock authentication can be valuable to the art community.
- A generative probabilistic graphical model is another direction that we may consider. This might be particularly helpful if we cannot find "good" features for the classification task. Furthermore, the physical process of drip paintings can help build the generative model. Nevertheless, learning issues would be challenging in this setting.
- A related problem that has direct application in the authentication task is the *color layer separation problem*. The main challenge of this problem is to recover each color layer used in the drip paintings from occlusion. Graph cut methods, as frequently applied in similar computer vision problems, might be a good direction to pursue.[18]

There is a long history of connoisseurship, including art authentication, and new computer-based methods have shown modest, if promising, benefits. There seem to be evidence that such computer-based techniques, when closely guided and constrained by art historical context, may find use as a supplement to traditional techniques and thus provide value in humanistic scholarship of art.[19, 20]

ACKNOWLEDGMENTS

We thank Jim Coddington of the Museum of Modern Art, Richard Taylor of the University of Oregon, and Dan Rockmore of Dartmouth College for helpful discussions.

REFERENCES

1. R. P. Taylor, A. P. Micolich, and D. Jonas, "Fractal analysis of Pollock's drip paintings," *Nature* **399**, p. 422, 1999.
2. R. P. Taylor, R. Guzman, T. P. Martin, G. D. R. Hall, A. P. Micolich, D. Jonas, B. C. Scannell, M. S. Fairbanks, and C. A. Marlow, "Authenticating Pollock paintings using fractal geometry," *Nature* **399**, p. 422, 2006.
3. A. P. Micolich, B. C. Scannell, M. S. Fairbanks, T. P. Martin, and R. P. Taylor, "Comment on 'Drip Paintings and Fractal Analysis' by K. Jones-Smith, H. Mathur and L. M. Krauss," *Nature* , 2007.
4. K. Jones-Smith and H. Mathur, "Fractal analysis: Revisiting Pollock's drip paintings," *Nature* **444**(doi: 10.1038/nature05398), pp. E9–E10, 2006.
5. K. Jones-Smith, H. Mathur, and L. M. Krauss, "Reply to Comment on 'Drip Paintings and Fractal Analysis' by Micolich *et al.*," *Nature* (doi: arXiv:0712.165v1), 2008.
6. M. Irfan and D. G. Stork, "Multiple visual features for the computer authentication of Jackson Pollock's drip paintings: Beyond box-counting and fractals," in *SPIE Electronic Imaging: Image processing: Machine vision applications II*, K. S. Niel and D. Fofi, eds., **7251**, pp. 72510Q1–11, SPIE/IS&T, Bellingham, WA, 2009.
7. T. Wagner, "Texture analysis," in *Handbook of computer vision and applications: Signal processing and pattern recognition*, B. Jähne, H. Haußecker, and P. Geißler, eds., **2**, pp. 275–308, Academic Press, New York, NY, 1999.
8. D. G. Stork, "Learning-based authentication of Jackson Pollock's drip paintings," *SPIE Newsroom* **May 27, 2009**(doi: 10.1117/2.1200905.1643), 2009.
9. R. P. Taylor, "Order in Pollock's chaos," *Scientific American* **287**(66), pp. 116–121, 2002.
10. J. Coddington, J. Elton, D. Rockmore, and Y. Wang, "Multifractal analysis and authentication of Jackson Pollock's paintings," in *Computer image analysis in the study of art*, D. G. Stork and J. Coddington, eds., **6810**, pp. 68100F–1–12, IS&T/SPIE, (Bellingham, WA), 2008.
11. J. R. Mureika, C. C. Dyer, and G. C. Cupchik, "On multifractal structure in non-representational art," *Physical Review* **72**(4), p. 046101, 2005.
12. M. Shahram, D. G. Stork, and D. Donoho, "Recovering layers of brushstrokes through statistical analysis of color and shape: An application to van Gogh's *Self portrait with grey felt hat*," in *Computer image analysis in the study of art*, D. G. Stork and J. Coddington, eds., **6810**, pp. 68100D–1–8, SPIE/IS&T, Bellingham, WA, 2008.
13. J. Alvarez-Ramirez, C. Ibarra-Valdez, E. Rodríguez, and L. Dagdug, "1/f-noise structures in Pollock's drip paintings," *Physica A* **387**(1), pp. 281–295, 2008.
14. E. J. Candès and F. Guo, "New multiscale transforms, minimum total variation synthesis: Applications to edge-preserving image reconstruction," *Signal Processing* **82**(11), pp. 1519–1543, 2002.
15. R. O. Duda, P. E. Hart, and D. G. Stork, *Pattern classification*, John Wiley and Sons, New York, NY, Second ed., 2001.
16. A. Vezhnevets and V. Vezhnevets, "'Modest AdaBoost'–Teaching AdaBoost to generalize better," *Graph-icon* **12**(5), pp. 987–997, 2005.
17. R. E. Schapire, "The Boosting approach to machine learning: An overview," in *Nonlinear estimation and classification*, D. D. Denison, M. H. Hansen, C. C. Holmes, B. K. Mallick, and B. Yu, eds., **171**, pp. 149–172, Springer, (New York, NY), 2003.
18. Y. Kuang, D. G. Stork, and F. Kahl, "Improved curvature-based inpainting applied to fine art: Recovering van Gogh's partially hidden brush strokes," in *Computer vision and image analysis of art II*, D. G. Stork, J. Coddington, and A. Bentkowska-Kafel, eds., (San Francisco, CA), 2011.
19. D. G. Stork and J. Coddington, eds., *Computer image analysis in the study of art*, vol. 6810, SPIE/IS&T, Bellingham, WA, 2008.
20. D. G. Stork, J. Coddington, and A. Bentkowska-Kafel, eds., *Computer vision and image analysis in the study of art*, vol. 7531, SPIE/IS&T, Bellingham, WA, 2010.