Multiple visual features for the computer authentication of Jackson Pollock's drip paintings: Beyond box-counting and fractals

Mohammad 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
and Department of Statistics, Stanford University, Stanford CA 94305 USA

ABSTRACT

Drip paintings by the American Abstract Expressionist Jackson Pollock have been analyzed through computer image methods, generally in support of authentication studies. The earliest and most thoroughly explored methods are based on an estimate of a "fractal dimension" by means of box-counting algorithms, in which the painting's image is divided into ever finer grids of boxes and the proportion of boxes containing some paint is counted. The plot of this proportion (on a log-log scale) reveals scaling or fractal properties of the work. These methods have been extended in a number of ways, including multifractal analysis, where an information measure replaces simple box paint occupancy. Recent studies suggest that it is unlikely that any single measure, including those based on such box counting, will yield highly accurate authentication; for example, a broad class of highly artificial angular sketches created in software reveal the same "fractal" properties as genuine Pollock paintings. Others have argued that this result precludes the value of such fractal-based features for such authentication. We show theoretically that even if a visual feature (taken alone) is "uninformative," such a feature can enhance discrimination when it is combined in a classifier with other features—even if these other features are themselves also individually uninformative. We describe simple classifiers for distinguishing genuine Pollocks from fakes based on multiple features such as fractal dimension, topological genus, "energy" in oriented spatial filters, and so forth. We trained linear-discriminant and nearest-neighbor classifiers using these features and found that our classifiers gave slightly improved recognition accuracy on human generated drip paintings. Most importantly, we found that although fractal features, taken alone might have low discriminative power, such features improved accuracy in multi-feature classifiers. We conclude that it is premature to reject the use of visual features based on box-counting statistics for the authentication of Pollock's dripped works, particularly if such measures are used in conjunction with multiple features, machine learning and art material studies and connoisseurship.

Keywords: fractal image analysis, Jackson Pollock, painting analysis, pattern recognition, box-counting algorithm, art authentication

1. INTRODUCTION

The American Abstract Expressionist Jackson Pollock (1912–1956) is best known for his so-called "action" paintings, executed by dripping, pouring and splashing liquid paint onto horizontal canvases on the floor (Fig. 1).[1] There is a large number of works of doubtful authorship and outright fakes executed in this drip technique. For instance, a large, unsigned work executed in this drip technique was purchased in a thrift store in San Bernadino CA in the early 1990s by Teri Horton and has been attributed to Pollock by some art scholars, based on thumbprint and other evidence;[2] this attribution has been contested or rejected by other art scholars, based in large part on traditional connoisseurship. Likewise, a cache of 32 small dripped works was found in the possession of Alex Matter in 2003, believed by some art scholars to be by Pollock's hand, a view contested by other art scholars.[3]

Authentication based on signatures, provenance (the documentary record of sales of a work), and traditional material studies of media, support, preparation (e.g., sizing) are not always definitive—for Pollock's oeuvre and

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

Figure 1. Jackson Pollock, *Untitled (Mural)* (1950) 183 × 244 cm (oil on canvas). This work is typical of Pollock's drip oeuvre, where the artist dripped and poured liquid paint onto the large canvas on the floor in sweeping arcs. (Pollock-Krasner Foundation and ArtStor.)

indeed those of many other artists—and thus any additional informative objective test could be quite valuable. We focus on image based tests, that is, ones based on high-resolution color or multi-spectral photographs of a work. Such image studies have shown promise when applied to the works of other painters. For instance, Lyu, Rockmore and Farid applied wavelet analysis to high-resolution images of the six portraits in Perugino's *The holy family* and claimed to identify three different "hands," that is, the work of three different artists.[4] They speculate that such image based methods might show promise in further attribution and authentication studies as well. Likewise, Chris Johnson and members of an international consortium have described a number of methods, including wavelet based ones, for the authentication of van Gogh's paintings.[5]

Pollock's distinctive abstract drip paintings do not avail themselves to such a wavelet analysis or indeed most traditional methods from computer vision and forensic image analysis. Richard Taylor, inspired by the apparent range of scales of structures in Pollock's drip paintings, pioneered the use of fractal analysis for their authentication.[6,7] A fractal is a mathematical structure that exhibits self-similarity—portions of the object at one scale have nearly the same structure as does the object at a different scale. Taylor estimated the "fractal dimension" of Pollock's works by means of a box-counting algorithm[6] (see Sect. 2, below). Such use of fractals—and even the claim that a true fractal dimension can be estimated from a Pollock painting—have been criticized by a number of scholars, however.[8] We stress, too, that decades of theoretical and empirical research in visual pattern recognition show that it is unlikely that the use of any *single* such feature will be highly reliable for any complex classification task, such as authenticating paintings by Pollock.[9,10]

We are unaware of research adequately exploring techniques of automatic classification based on *multiple* visual features as applied to the problem of authenticating Pollock's paintings. This is the approach we present here. We begin in Sect. 2 with a brief description of the box-counting algorithm, fractals, and related mathematical measures, as well as some of the criticisms of these technique for the authentication of Pollock's drip paintings. We demonstrate theoretically that even a single feature that provides no discriminative information may provide discriminative information when used in conjunction with other features. Then we describe our visual features and classification methods in Sects. 3 and 4, respectively. Section 5 contains our empirical classification results. We conclude in Sect. 6 with lessons learned and extentions for further use in Pollock studies.

2. FRACTALS FOR AUTHENTICATING POLLOCK'S DRIP PAINTINGS

A *fractal* is a geometric object that shows self-similarity, and can be characterized by a possibly a real-valued "fractal dimension" D, which could be non-integer.[11] Many complex natural objects, such as coastlines, clouds, certain crystals, and rivers show fractal properties or fractal behavior, as do most natural images and some visual art.[12] Richard Taylor and his colleagues introduced the use of fractals for authenticating the drip paintings of Pollock[6] and have explored its use in a number of studies.[13,14] In brief, these researchers advocated the use of a box-counting algorithm to estimate such "fractal" properties, as we shall discuss in Sect. 2.2.[3] (We mention in passing that other stochastic characterizations of Pollock's works have been explored such as $1/f$ noise[15] and *multifractals*[16,17] though with only modest theoretical or empirical benefits. Accordingly, we shall concentrate on fractals even though our results and recommendations apply to other such measures as well.)

Figure 2. Remco Teunen, *Untitled* (oil on canvas). For our analyses we used paintings such as this as an admittedly poor representative of a "fake" Pollock; no serious art scholar or even casual viewer would ever mistake this painting for a genuine Pollock. Nevertheless, such a painting has complex features and has a number of visual features that can be used in the study of the authentication of Pollock paintings.

Proponents claimed to have shown value in the scale-space analysis for authentication, but...

> ...[t]aken in isolation, these [fractal] results are not intended to be a technique for attributing a poured painting to Jackson Pollock. However, the results may be useful when coupled other important information such as provenance, connoisseurship and materials analysis.[18,19]

Jones-Smith, Mathur and Krauss levied a number of criticisms of the fractal and box-counting approach for the authentication of Pollock's drip paintings:[8,20]

Non-uniqueness of fractal characteristics These critics created a number of highly artificial images that matched the fractal or scale-space properties of genuine Pollocks—images that appear very different from genuine Pollocks.

Inadequacy of box-counting algorithm They also pointed out that the range of spatial scales the Taylor team used in its box-counting method (roughly two orders of magnitude) was narrower than that traditionally used to estimate a fractal dimension, in short, that the properties of Pollocks estimated are not true fractals anyway.

Contamination of fractal properties by occlusion They noted further that even if each layer of different colored paint in a Pollock painting individually exhibited fractal characteristics, the exposed (i.e., visible) portion of a layer need not exhibit fractal characteristics. They thus claimed that estimating a fractal dimension from the visible image of each layer was inappropriate for Pollock authentication.

For these central reasons these scholars conclude:

> Our data make it clear that the fractal criteria of Taylor *et al.* should play *no* role whatsoever in authenticity debates. Given the complete lack of correlation between artist and fractal characteristics that we have found, in particular, the failure of fractal analysis to detect deliberate forgery, it is clear that box-counting data are not useful even as a supplement to other analysis." [20, emphasis in the original]

We show below that the Jones-Smith *et al.* conclusions and negative recommendations—that fractal analysis "...should play *no* role whatsoever in authenticity debates"—are unjustified or at least premature on both theoretical and empirical grounds.

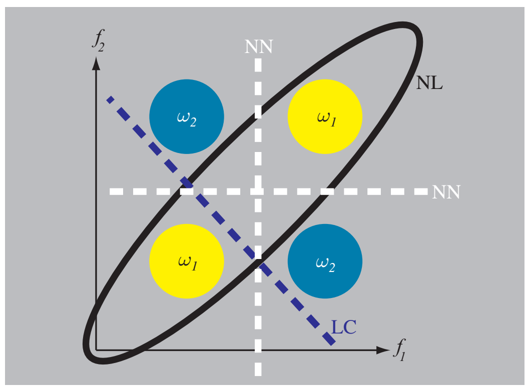

Figure 3. The yellow disks represent the locations of patterns in category ω_1 in the two-dimensional feature space $f_1 \times f_2$, and the dark-blue disks represent patterns in category ω_2. We assume these categories have equal prior probabilities, $P(\omega_1) = P(\omega_2) = 0.5$. (This is an example of the exclusive-OR problem.) The projections of these probability densities onto the f_1 axis are the same for the two categories, and thus any classifier based solely on feature f_1 would have a 50% error rate, the worst possible for a dichotomizer where the categories have equal priors; that is, the Bayes error rate is 50%. Therefore, taken alone, measurements of f_1 of any test pattern are uninformative and useless for classification. The projection of the densities onto the other feature axis, f_2, are similarly uninformative, and the classification rate based on feature f_2 alone is also 50%. Notice, though, that one particular classifier based on *both* features, shown by the (nonlinear) elliptical decision boundary marked **NL**, has 0% error rate—perfect classification. A nearest-neighbor classifier, marked **NN**, here would split the input space into quadrants, classifying all patterns in the upper-right and lower-left quadrants as ω_1 and elsewhere ω_2, and thus also have a 0% error rate. Even a simple linear classifier, shown by the straight line marked **LC**, yields a 25% error rate. In short, even if f_1 represented some measure of the "fractal" properties of Pollock's drip paintings, and even if *taken alone* provides no information for classification, this feature might be useful when used in a classifier based on multiple features.

2.1 Response to claimed drawbacks of non-uniqueness

We first address the Jones-Smith *et al.* "non-uniqueness" claim from a theoretical standpoint. Figure 3 shows hypothetical distributions of two categories in a two-dimensional feature space $f_1 \times f_2$, where patterns in category ω_1 could represent genuine Pollocks and ω_2 could represent fakes, possibly even the computer generated drawings of Jones-Smith *et al.* In this hypothetical case suppose f_1 represented some fractal property. Note especially that here both categories have the *same* distributions when projected onto the f_1 axis, and thus this feature, taken alone, provides no discriminative information whatsoever. This f_1 feature is "useless," much as Jones-Smith *et al.* claim of fractal features in Pollock authentication studies.

Nevertheless, when used in conjunction with another feature, here f_2, f_1 can indeed be useful. The general phenomenon by the classification boundaries illustrated in Fig. 3—the classification benefit of multiple feature, even individually useless features—occurs frequently throughout real-world pattern classification, especially in problems with complicated nonlinear decision boundaries and multiple features. The benefit of multiple features is never as great as in the extreme case of the exclusive-OR problem just discussed, but the benefit can be large and statistically significant. Nor does each feature need be uninformative for such a benefit. In the case that the probability densities of the two categories are known, under extremely lax conditions, the more features are used the lower the resulting classification error. Stated more precisely, if the full class-conditional probability distributions are known, the inclusion of more features cannot increase the Bayes error rate. [10, cf. Chapters 1, 2 & 9 and Fig. 3.3]

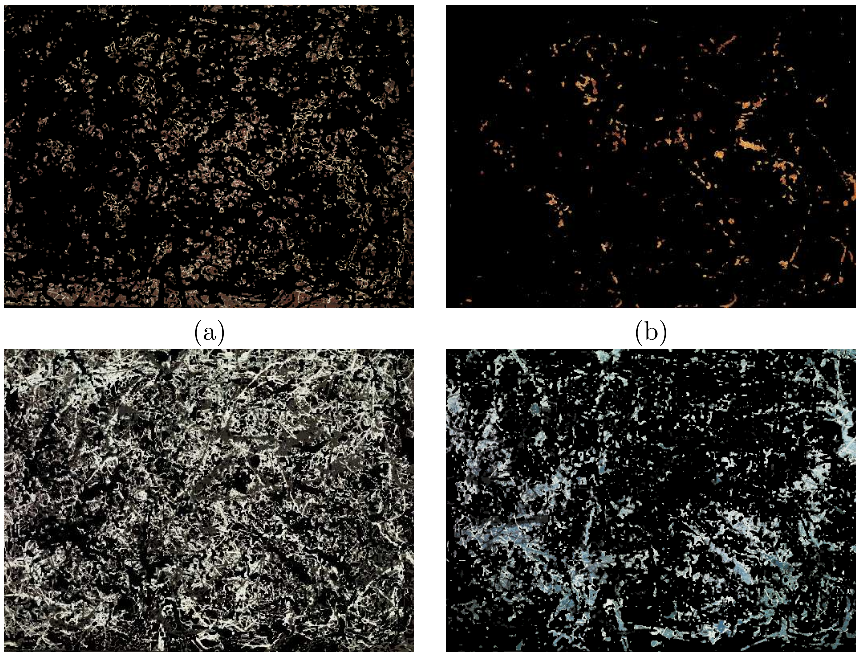

(a) (b) (c) (d)

Figure 4. Color layer segmentation of Jackson Pollock, *Mural* of Fig. 1 based on color clustering. (a) Brown layer. (b) Red layer. (c) Gray layer. (d) Blue layer. Note that these layers do not correspond to the specific layers of paint as applied by the artist, nor is it necessary that they do so for the extracted visual features to be useful in classification.

Note that even if the "fractal" dimension f_1 taken alone were uninformative, it could be useful when combined with *non-visual* features, for instance material studies or provenance, such as advocated by Micolich *et al.* Consider again Fig. 3 where we assume ω_1 represents genuine Pollocks and ω_2 fakes or forgeries, that the f_1 feature is some quantitative measure of fractal dimension. Now suppose feature f_2 represents some material property of the work, for instance the paint medium: *acrylic* works are represented at the bottom of the graph while *oil* works at the top. (Other such material features could be whether the canvas was coated/sized or instead unsized, the support was duck canvas or paper, the canvas was mounted onto a stretcher or not, and so forth.) In this hypothetical case, the distributions in Fig. 3 could be summarized thus: when Pollock painted with oil, his paintings had a high fractal dimension; when he painted with acrylics, his paintings had a low fractal dimension. Forgers (for some reason) had the converse. A nonlinear classifier using the single visual fractal feature as well as the material information could in principal have perfect accuracy distinguishing Pollocks from fakes, at least in this theoretical example. As such, the negative recommendation by Jones-Smith, Mathur and Krauss, above[20] must be rejected, at least in theory.

We shall see in Sect. 5 that our preliminary analyses show that the Jones-Smith *et al.* negative recommendation based on their "non-uniqueness" claim must be rejected on empirical grounds as well.

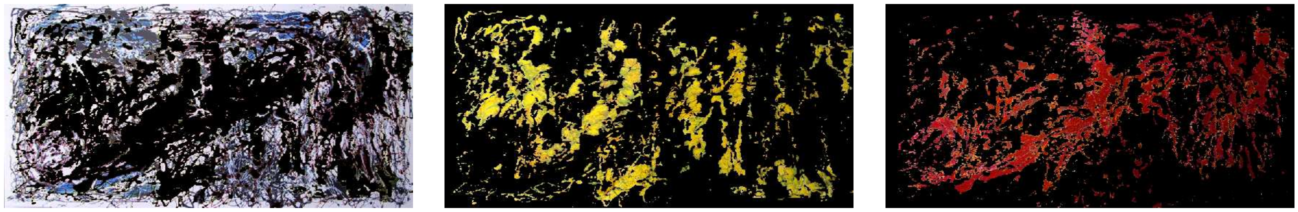

(a) (b) (c)

Figure 5. Color layer segmentation of Remco Teunen, *Untitled*, of Fig. 2 based on color clustering. (a) Blue-gray layer. (b) Yellow layer. (c) Red layer. Just as in Fig. 4, it is not necessary that this color segmentation to recover the actual number and full layers created by Teunen for the extracted visual information to be useful for classification.

2.2 Response to claimed inadequacies of the box-counting algorithm

The traditional notion of topological dimension refers to the number of numerical values needed to specify a point on an object. A line or line segment has topological dimension $D = 1$, because each point on the line can be specified uniquely by a single number. A plane has topological dimension $D = 2$, because each point requires a *pair* of numbers, e.g., one for the horizontal position, one for the vertical position. Likewise for a spatial volume such as a cube has topological dimension is $D = 3$, and so on for higher, abstract spaces. A fractal dimension is a generalization of topological dimension to the case of non-integer values.[11]

Taylor and his colleagues computed a "fractal dimension" by means of a box-counting algorithm.[7] This algorithm is applied to a monochromatic drip painting or segmented color layer as follows. One preprocesses the painting to segment color layers. Then for each layer, one splits the image into a set of large boxes and counts the proportion of these boxes that contain the image of any paint. Next, the size of the boxes is decreased, the paint occupancy estimated again, and so on. In nearly every case, the smaller the boxes, the higher the proportion of unoccupied boxes. The slope of a log-log plot of this occupancy versus box size is related to the painting's fractal dimension.[6] (There are a number of technical subtleties—such as the need for averaging over box spatial offsets—that need not concern us here.)

Jones-Smith *et al.* have argued that the typical range of scales used by Taylor and his colleagues (roughly two orders of magnitude) is too small to be used to estimate a true fractal dimension. Instead, a scale-space metric based on box-counting, including the reasonable statement about the need for large range in scale. While this criticism may be true, it is irrelevant, at least in theory, for the problem of pattern recognition. What is essential in pattern recognition is not the linguistic term applied to a feature, but rather that it can be specified and computed in an objective, repeatable way.

We eschew the debate over the term "fractal" in this context except to state we seek to transcend the limitations of this term. Visual features in computer vision and pattern recognition are often extremely complicated, resulting from successive nonlinear image processing at different scale; such features need not conform to some crisp mathematical form in order to be useful. In classical multi-layer neural networks, for example, the hidden layers compute (intermediate) features that are nearly always complicated and hard to interpret, let alone express in a crisp mathematical form. What matters is not whether a feature has some particular idealized mathematical description conforming to its use in a different domain, but instead that it can be computed, and of course that it improves classification accuracy in a given classifier architecture.[10, 21]

2.3 Response to claimed drawbacks of contamination of fractal properties by occlusion

Pollock executed his paintings by pouring paint, layer atop layer, alternating colors, and even revisiting colors. Even if each single layer, taken alone, were to exhibit fractal characteristics, the occlusion of one layer by subsequent layers would disrupt the estimated mathematical properties of the lower layer. That is, the visible portion of the lower layer would differ unobstructed. Jones-Smith and Mathur showed that the fractal properties of one layer could be corrupted or obscured by a subsequent layer and argued that this made fractal measures inappropriate for the case of Pollock studies.[8]

Note that, just as in Sect. 2.2, from the perspective of statistical pattern recognition there is no requirement that overlapped recover the actual drip patterns hidden by higher higher layers of drips nor their fractal characteristics. All that is needed is that the information be useful in a classifier. This can be addressed only empirically.

A simple example: Suppose that for the paintings of artist A the lower layer consisted of lines oriented primarily at $\pm 45°$ to the vertical and that for the paintings of artist B the lower layer were $\pm 30°$ to the vertical. Suppose further that for both artists these lower layers were partially obscured by upper layers of horizontal and vertical lines. In each case, the upper layer would disrupt the statistics but the orientation information extracted from the lower layer might nevertheless be useful in a classifier.

Artist	Title (year)
Jackson Pollock	*Alchemy* (1947)
Jackson Pollock	*Blue Poles* (1953)
Jackson Pollock	*Cathedral* (1947)
Jackson Pollock	*Convergence* (1952)
Jackson Pollock	*Eyes in the Heart* (1946)
Jackson Pollock	*Lucifer* (1947)
Jackson Pollock	*Mural* (1950)
Taylor Winn	*Tissue* (2005)
Taylor Winn	*Virus* (2005)
Taylor Winn	*Surge* (2006)
Jean-Paul Riopelle	*Espagne* (1951)
Jean-Paul Riopelle	*Pavane* (1954)

Table 1. Paintings used in our classification analyses.

3. OTHER VISUAL FEATURES

Given an image of a painting and the number of colors used in the painting as inputs, we need to extract different color layers of the painting for the purpose of feature extraction. This problem is, to some extent, independent of the spatial analysis. We segmented the image based on color using the MATLAB image processing toolbox. First, we convert the input image to the a^*b^* color space and then apply the k-means clustering algorithm based on Euclidean distance. Last, we segment the original image by the colors that we have obtained from clustering. This approach does not work well when the number of colors is large, however.[22] This color layer separation step is illustrated in Figs. 4 and 5.

We now turn to an empirical demonstration that use of other features can lead to better classification results. The features we tested were two based on prior fractal work, and three others. These were:

f_1**: fractal dimension** The fractal dimension, as estimated by the box-counting method proposed by Taylor and his colleagues, as outlined in Sect. 2.2.

f_2**: multifractal feature** The multifractal feature is a measure of how well the color layers show multifractality or information dimension. The norm of the log-log line fit of occupancy gives a good measure of how well a color layer shows the property of information dimension. We take the average of the line fits corresponding to the different color layers, excluding the worst line fit, as a feature. Murieka *et al.* showed that for Pollock paintings the information dimension of a single color layer is less than the Levy dimension of the same layer.[16] The difference between the two dimensions is a measure of the multifractality of the painting.

f_3**: Genus** The Euler number e is the difference between the number of separate objects and the number of holes. We have used the normalized Euler number as a feature.

f_4**: Oriented energy A** This is the sum of the differences between the oriented energies at $x°$ and $-x°$, for all $x \in \{10, 20, ..., 80\}$ where our convention is $0°$ represents the vertical. Oriented energies are found by

Artist	Title (year)	Classifier Result
Jackson Pollock	*White Light* (1954)	Pollock
Jackson Pollock	*Number 05* (1950)	Pollock
Jackson Pollock	*Number 28* (1950)	Pollock
Jackson Pollock	*Reflection of the Big Dipper* (1947)	Pollock
Taylor Winn	*Sunstorm* (2006)	Non-Pollock
Taylor Winn	*Circus* (2006)	Non-Pollock
Taylor Winn	*Tres Pez* (2005)	Non-Pollock
Taylor Winn	*Untitled II* (2005)	Non-Pollock
Jean-Paul Riopelle	*La Fort Ardente* (1955)	Pollock

Table 2. Classification of the test patterns by the linear classifier.

convolving the input image with oriented Gabor filters[23,24] at every 10° interval and then counting the number of pixels in the resulting image.

f_5: **Oriented energy B** The last feature is the difference between the sum of the number of local maximas and minimas on either side of the vertical direction in the plot of the oriented energies as a function of the angles.

These five features comprise the feature vector for each image. These features are admittedly somewhat arbitrary, based on informal considerations; we did not perform an extensive search for optimal features (see Sect. 6). The values of the different features might span different ranges, inappropriately making some features "more important" than others, skewing recognition rates. Thus we need to make sure that all features have roughly equal importance. We have achieved this by traditional standardization of the feature values: for each feature, we have transformed the feature values so that the feature values are distributed over the data set with mean 0.0 and variance 1.0. [10, Sect. 6.8.3]

4. CLASSIFIER DESIGN

We explored several well known classifier and machine learning techniques.[10] In broad overview, in a linear classifier such as the Perceptron each feature value is multiplied by a separate, learned, weight and summed. If the resulting value is greater than zero, the pattern is classified as ω_1, and otherwise ω_2. In the nearest-neighbor classifier, a pattern of unknown category (the test pattern) is assigned the category label of its nearest training pattern. We first describe the feature extraction process. Then we discuss how the features are used in the Perceptron and the nearest-neighbor algorithms.

After extracting the color layers of the image, we have applied the fractal analysis technique of Taylor *et al.* on each color layer. Determination of the levy and drip dimensions largely depends on the image size. If the image is too small, which is often the case in our study, then we cannot determine the drip dimension correctly. Furthermore, this analysis requires the color layers to be separated correctly. Previous studies have hand-picked the colors to achieve this but we built a rough automatic segmentation tool. Thus we have to deal with the inaccuracy in the color layer separation. For example, we have separated three primary colors, roughly brown, gray and red, of Pollock's *Blue Poles* using the simple k-means clustering technique. Then we have applied the fractal analysis technique on each of the three color layers. The norms of the line fits for this painting are 0.10, 0.07 and 0.18, respectively. Notice the relatively high magnitude of the third norm. This is due to the fact that the third color layer (red) is not as dense as the other two color layers and it does not show significant fractal behavior. Inaccurate color-space segmentation is the main reason for this phenomenon. We have observed that this type of undesired behavior happens quite frequently. To deal with this, we take the average of the norms of the line fits, excluding the highest norm, as the feature of fractal behavior.

To design the linear classifier we start with a set of training vectors—a set of feature vectors whose class (Pollock or non-Pollock) we already know. The decision boundary is a hyperplane in feature space defined by

$$\sum_{i=1}^{5} w_i f_i + b = 0 \quad \text{or}$$
$$\mathbf{w}^t \mathbf{f} + b = 0,$$

where the f_i are the values of the five features and weights w_i and bias b are initialized arbitrarily and we denote vectors by boldface. We trained our linear classifier using the Perceptron algorithm. For each training vector $\mathbf{v}$, if $\mathbf{v}$ is a genuine Pollock and $\mathbf{w}^t\mathbf{v} + b < 0$ then we update $\mathbf{w}$ and b by adding ρv and ρN^2 to these respectively, where ρ is the learning rate, taken as 0.5 and N is the maximum Euclidean norm among all the training vectors. Conversely, if $\mathbf{v}$ is a non-Pollock and $\mathbf{w}^t\mathbf{v} + b > 0$ then we update $\mathbf{w}$ and b by subtracting $\rho\mathbf{v}$ and ρN^2 from these respectively. We iterate over the set of training vectors, with learning decay, until convergence. [10, Chapter 5]

After the weights are learned, test patterns are classified by the linear rule:

$$\text{If } \mathbf{w}^t\mathbf{t} + b > 0, \text{ then } \quad \mathbf{t} \in \text{Pollock}$$
$$\text{otherwise} \quad \mathbf{t} \notin \text{Pollock.}$$

For the nearest-neighbor classifier we have preprocessed the feature vectors by the standardization method that we have discussed earlier. The nearest-neighbor classifier is given a set of training vectors and it classifies a test feature vector to the class of the nearest training vector. We have used the squred Euclidean distance as a

Test pattern	Nearest neighbors	Distance
White Light (Pollock)	*Lucifer* (Pollock)	2.49
	Alchemy (Pollock)	3.52
	Surge (Winn)	4.85
Number 05 (Pollock)	*Alchemy* (Pollock)	4.57
	Lucifer (Pollock)	7.12
	Surge (Winn)	11.36
Number 28 (Pollock)	*Eyes in the Heart* (Pollock)	3.34
	Lucifer (Pollock)	3.55
	Surge (Winn)	4.43
Reflection of the Big Dipper (Pollock)	*Blue Poles* (Pollock)	1.62
	Convergence (Pollock)	2.10
	Surge (Winn)	2.65
Sunstorm (Winn)	*Surge* (Winn)	3.25
	Convergence (Pollock)	6.72
	Pavane (Riopelle)	7.18
Circus (Winn)	*Surge* (Winn)	6.13
	Virus (Pollock)	7.34
	Convergence (Pollock)	8.19
Tres Pez (Winn)	*Surge* (Winn)	2.10
	Blue Poles (Pollock)	3.40
	Convergence (Pollock)	4.08
Untitled II (Winn)	*Surge* (Winn)	2.81
	Eyes in the Heart (Pollock)	7.22
	Convergence (Pollock)	9.44
La Fort Ardente (Riopelle)	*Convergence* (Pollock)	0.29
	Mural (Pollock)	0.96
	Pavane (Riopelle)	1.06

Table 3. Classification of the test patterns by the nearest-neighbor classifier.

measure of proximity among the feature vectors. We have also considered other versions of the standard nearest neighbor classifier, namely k-nearest neighbor classifier, where we compute the $k \geq 1$ training vectors that are closest to the test vector and we classify the test vector to the most representative class among these k training vectors.

5. RESULTS

We have trained the Perceptron-based linear classifier with the paintings listed in Table 1. For each of these paintings, we have separated three primary color layers. After the training phase, we have tested the linear classifier with nine paintings, listed in Table 2. For the nearest-neighbor classifier, we have used the same training patterns as listed in Table 1 and Table 3 shows typical results. As evident from Tables 2 and 3, both these classifiers give very good accuracy in our small experiment.

Table 5 shows our overall results related to the central question at hand: the value of multiple features. First, the linear classifier out performed the nearest-neighbor classifier when many features were used. The classification accuracy using a larger feature set was usually superior to (and never inferior to) the accuracy based on a subset of those features.

Classifier	$\{f_1\}$	$\{f_2\}$	$\{f_3, f_4, f_5\}$	$\{f_2, f_3, f_4, f_5\}$	$\{f_1, f_3, f_4, f_5\}$	$\{f_1, f_2, f_3, f_4, f_5\}$
Linear	76.2%	52.4%	47.6%	52.4%	76.2%	81.0%
Nearest-neighbor	66.7%	71.4%	47.6%	66.7%	66.7%	76.2%

Table 4. Classifier accuracy for different feature sets. In each case, the classifier performace was estimated through leave-one-out training. [10, Sect. 9.6.3] Several entries show the classification benefit of multiple features. For instance, in both the linear and the nearest-neighbor classifiers, the use of $\{f_1, f_2, f_3, f_4, f_5\}$ is superior to $\{f_1\}$, to $\{f_2\}$ and to $\{f_3, f_4, f_5\}$ feature sets used alone. Nevertheless, for any classifier and feature f_i, the recognition is generally higher if that feature is part of a larger set of features: in short, the use of multiple features is generally (but not universally) superior to the use of fewer features.

6. CONCLUSIONS

Our contributions to the problem of image analysis for authentication of Jackson Pollock's drip paintings are both theoretical and empirical. We have shown theoretically that a single "uninformative" feature can nevertheless be useful in classification when used in conjunction with other features—even similarly "uninformative" features. We showed that multiple features in fact outperform a single feature and that the "fractal" feature provides some classification improvement, albeit small.

The theoretical result that the use of more features never reduces the Bayes error rate, described in Sect. 2.1, applies to the case that the full distributions are known (or estimated) and can be implemented in the classifier. In theory this might require an arbitrarily large data set—which cannot exist for an artist with a finite oeuvre—though the more (error free) training data that is used, the more closely these distributions can be estimated.

Clearly, our empirical work is preliminary and the non-Pollock paintings we tested were not of sufficiently high quality that our current results are applicable to the authentication debates in the art community. Accordingly, our specific classification results must be extended and applied to larger corpora of high-resolution images with more features, and exploiting the full power of current classification methods: sophisticated feature selection and pruning, cross validation, boosting and bagging, and so on.[10]

We believe our current results justify exploring such next steps in Pollock authentication studies.

ACKNOWLEDGMENTS

We would like to thank Richard Taylor and Katherine Jones-Smith for helpful discussions and Ms. Jones-Smith for access to the "fake" Pollock image data. We also thank the Pollock-Krasner Foundation for permission to reproduce the Pollock painting and Remco Teunen for permission to reproduce his painting.

REFERENCES

1. K. Varnedoe and P. Karmel, *Jackson Pollock*, Museum of Modern Art, New York, NY, 2007.
2. P. P. Biro, "Forensics and microscopy in authenticating works of art," *Infocus* (1), pp. 4–17, 2006.
3. E. G. Landau and C. Cernuschi, eds., *Pollock/Matters*, McMullen Museum of Art, Boston, MA, 2007.
4. S. Lyu, D. Rockmore, and H. Farid, "A digital technique for art authentication," *Proceedings of the National Academy of Sciences* **101**(49), pp. 17006–17010, 2004.
5. C. R. Johnson, E. Hendriks, I. J. Berezhnoy, E. Brevdo, S. M. Hughes, I. Daubechies, J. Li, E. Postma, and J. Z. Wang, "Image processing for artist identification," *IEEE Signal Processing magazine* **25**(4), pp. 37–48, 2008.
6. R. P. Taylor, A. P. Micolich, and D. Jonas, "Fractal analysis of Pollock's drip paintings," *Nature* **399**, p. 422, 1999.
7. R. P. Taylor, "Order in Pollock's chaos," *Scientific American* **287**(66), pp. 116–121, 2002.
8. K. Jones-Smith and H. Mathur, "Fractal analysis: Revisiting Pollock's drip paintings," *Nature* **444**(`doi: 10.1038/nature05398`), pp. E9–E10, 2006.
9. T. Y. Young and K. Sun Fu, *Handbook of pattern recognition and image processing*, Academic Press, New York, NY, 1986.
10. R. O. Duda, P. E. Hart, and D. G. Stork, *Pattern classification*, John Wiley and Sons, New York, NY, Second ed., 2001.
11. B. Mandelbrot, *The fractal geometry of nature*, W. H. Freeman, New York, NY, 1982.
12. C. Redies, J. Hasenstein, and J. Denzler, "Fractal-like image statistics in visual art: Similarity to natural scenes," *Spatial Vision* **21**(1–2), pp. 137–148, 2007.
13. 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.
14. 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," *Pattern Recognition Letters* **28**(6), pp. 695–702, 2007.
15. 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.
16. J. R. Mureika, C. C. Dyer, and G. C. Cupchik, "On multifractal structure in non-representational art," *Physical Review* **72**(4), p. 046101, 2005.
17. 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.
18. A. Abbott, "Fractals and art: In the hands of a master," *Nature* **439**, pp. 648–650, 2006.
19. 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.
20. 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.
21. I. Guyon, S. Gunn, M. Nikravesh, and L. A. Zadeh, eds., *Feature extraction: Foundations and applications*, Springer, New York, NY, 2006.
22. 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.
23. T. S. Lee, "Image representation using 2D Gabor wavelets," *Transactions on Pattern Analysis and Machine Intelligence (PAMI)* **18**(10), pp. 959–971, 1996.
24. D. G. Stork and H. R. Wilson, "Do Gabor functions provide appropriate descriptions of visual cortical receptive fields?," *Journal of the Optical Society of America A* **7**(8), pp. 1362–1373, 1990.