

Multiple visual features, regularization and machine learning for the authentication of Jackson Pollock’s drip painting

Mohammad Irfan,^a David G. Stork^b and Jim Coddington^c

^aDepartment of Computer Science, Stony Brook University, Stony Brook NY 11790 USA

^bRicoh Innovations, 2882 Sand Hill Road Suite 115, Menlo Park CA 94025 USA
and Department of Statistics, Stanford University, Stanford CA 94305 USA

^cMuseum of Modern Art, 11 West 53rd Street, New York, NY 10019 USA

ABSTRACT

Jackson Pollock’s action paintings—in which he dripped, poured, and splashed liquid paint onto canvases on the floor—are some of the most important works in American Abstract Expressionism. There are many works of doubtful attribution and outright fakes, including a celebrated work purchased for \$5 in the early 1980s, to a cache of works found in 2002 by Alex Matter, son of one of Pollock’s friends. Material studies of paint, support, priming, signatures and provenance (the documentary record of the sequence of ownership of a work) are not always definitive, and any additional, complementary evidence could be useful in authentication studies. To this end, Taylor and his colleagues (Taylor, Micolich and Jonas, 99) pioneered the application of fractal analysis to the dripped works of Pollock. In brief, the Taylor group used a box-counting algorithm to estimate fractal and scale-space signatures of Pollock’s works; they reported that such signatures differed sufficiently from “fake” Pollocks that their method could be used as part of authentication protocol. The fractal approach was criticized on a number of grounds, however, including that fractals could be “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 (Jones-Smith and Mathur, 2006). Irfan and Stork, however, answered those objections both theoretically and experimentally, and demonstrated that the classical paradigm of pattern classification, (Duda, Hart and Stork, 2001) based on multiple visual features and statistical estimation of classifier parameters, leads to improved recognition accuracy (Irfan and Stork, 2009). While a single feature led to near-chance classification (53% accuracy), five features led to 82% accuracy. Although their empirical results were encouraging, they were nevertheless preliminary. Our current work builds upon that work by training on more image data, and of higher resolution, of both genuine Pollocks and “fakes,” and employs sophisticated feature extraction, feature selection and classifier pruning algorithms commonly used in pattern recognition research. The final accuracy depends upon the amount of data, of course, and should exceed that reported by Irfan and Stork.

Keywords: Jan van der Heyden

1. INTRODUCTION

xxxx

xxx¹

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

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

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 squared Euclidean distance as a

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

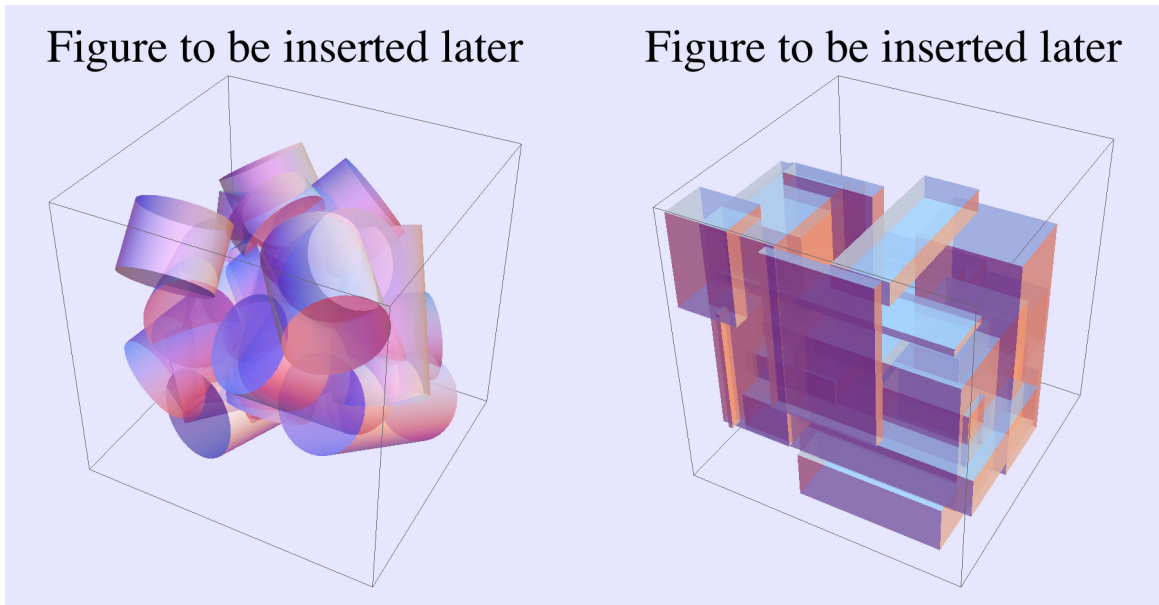


Figure 1. JunkFig

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.

2. RESULTS

ABSTRACT

A recent paper questioned the value of fractal- or scale-based visual features for the authentication of drip paintings, in particular those by the American Abstract Expressionist Jackson Pollock. The criticisms and negative recommendations arose from a failure to follow well-established principles and methodologies from statistical pattern recognition. This current paper demonstrates that several of the key negative arguments are invalid or at best incomplete, and points to empirical data that shows the value of such scale-space features in multi-feature statistical classifiers, ones that may be useful as part of a broader art historical authentication regime.

INTRODUCTION: FRACTAL IMAGE ANALYSIS OF DRIP PAINTINGS

Art authentication and related tasks—such as dating artworks, identifying artists who collaborated on different passages in a given artwork, etc.—are subtle, vexed and poorly understood. Scholars rely on artists’ signatures, provenance (the documentary record of ownership), chemical studies of media and pigments (oil, acrylic), carbon-14 dating, dendrochronology (dating wood panel supports by the pattern of growth rings), material studies of the support (paper, canvas), preparation (e.g., sizing), fingerprints, and the connoisseurship of experts intimately familiar with the artist, his oeuvre, his contemporaries, and so forth.[?] Such analyses are often not definitive, however, and thus any additional informative objective tests of an artwork could be valuable in authentication studies.

The Abstract Expressionist Jackson Pollock (1912–1956) is one of America’s most important artists, best known for his “action” paintings executed by dripping, pouring and splashing liquid paint onto horizontal canvases on the floor. There are a number of works of doubtful authorship and outright fakes, and hence any technique that might aid in such authentication would be valuable to the art community. Nearly a decade ago, Richard

Taylor and his colleagues introduced a fractal or scale-space analysis to this end.^{?,?} Their core algorithm estimated the fractal properties of candidate Pollock paintings by means of a box-counting procedure in which, after preprocessing, the image is divided into boxes of different sizes and the proportion of boxes containing any paint versus box size is plotted. The slope of this occupancy data on a log-log plot yields the painting’s fractal dimension. Taylor and his colleagues reported that genuine Pollock paintings generally yielded a characteristic two-legged shape on such a log-log plot whereas fake and other non-genuine Pollocks generally did not.

Recently, three physicists—Katherine Jones-Smith, Harsh Mathur and Lawrence S. Krauss (hereafter J-SMK)—levied strong criticisms of the Taylor et al. fractal approach, and by extension related scale-space methods, such as those based on multifractals.[?] In brief, their three main rebuttals, based on both mathematical analysis and empirical data, were:^{?,?,?,?}

Fractals provide no discriminative information J-SMK showed that some artists can create drip paintings whose fractal characteristics matched those of genuine Pollock paintings. They showed that fractal features were unreliable, and that the fractal signatures of genuine Pollocks could also be matched by those from highly artificial computer drawings of stars and other geometric figures—figures that do not resemble genuine Pollock or even fake Pollock paintings. Finally, they showed that a few fake paintings passed the Taylor et al. fractal criteria and a few genuine Pollock paintings failed the criteria.

Occlusion disrupts visual fractal properties J-SMK showed that even if one layer of paint (of one color) possessed fractal properties, such visual properties are almost always disrupted or destroyed by layers of subsequently applied paint (of different colors), which occlude the lower layer being analyzed.

Fractal analysis requires a broader range of scales J-SMK pointed out that true fractal properties of an image can be estimated only if the range of spatial scales used in the algorithms is sufficiently large, typically a factor of 100×, and that the Taylor et al. data did not span this requisite range.

For these and subsidiary reasons, J-SMK concluded:[?]

“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.*” [emphasis added]

This current paper relies on principles from statistical pattern recognition, the extensive body of research on feature selection and texture classification, and refers to published empirical data to rebut or show the limitations of these individual claims of J-SMK; it thus leads to a rejection of their overall negative conclusion and recommendation.

This paper is organized as follows. Section 3 reviews the basic methodology of statistical pattern classification—specifically feature selection—and proves that even uninformative or “useless” features can improve classification when used in conjunction with other features. These features need not even be *visual* features. The section also points to preliminary empirical research demonstrating precisely this fact for the case of scale-space features in Pollock authentication. It reviews the fact that the procedure of pattern classification always reduces information or projection of features onto a feature subspace. Accordingly, one can always construct highly atypical, artificial patterns that are nevertheless classified as a member of a given category; this fact, however, does not argue against the value of a classifier when applied to patterns from the appropriate distributions.

Section 4 shows that even though J-SMK are correct that the partial occlusion of one fractal pattern by another will disrupt or destroy the visible fractal properties of the layer partially occluded, this criticism is not relevant to the design of classifiers for Pollock authentication. First, the image pre-processing algorithm for *inpainting* can estimate and hence recover some (or in rare cases *all*) of the image information hidden by the occluding layers in paintings. More likely, such inpainting can recover visual properties that increase the usefulness of the layer information in classification, whether or not this information represents true (unoccluded) fractal properties. Furthermore, there is no need for the visible portion to be a true fractal to be of use to a

classifier anyway; all that is necessary is that the visible information be useful in predicting the category of the image. It is methodologically invalid to rule out the possible value of a feature (such as a fractal feature) if preprocessing can—even in principle—permit such a feature to improve classification accuracy. In fact a recent study showed that such partial (unoccluded) image information can indeed be recovered through preprocessing.

Section 5 considers the J-SMK claim that the range of spatial scales that has been used in box-counting algorithms is too small to recover true fractal characteristics. The brief rebuttal to this claim follows that in Sect. 4: there is no need for the feature to be a true fractal for it to be useful in classification. The section also refers to recent empirical evidence that shows the value of such “corrupted” information in the classification of Pollock paintings.

Section 6 summarizes the analyses and the rejection of the negative recommendations of J-SMK and argues instead that the traditional methodology from statistical pattern classification based on multiple features, possibly including fractal or scale-space feature extractors, properly trained on the appropriate distributions may indeed serve as part of a broader effort of art historical authentication of drip paintings.

3. REBUTTAL TO THE CLAIM THAT BECAUSE FRACTAL FEATURES ARE “UNINFORMATIVE” THEY SHOULD NOT BE USED IN AUTHENTICATION

This section considers the three primary arguments by J-SMK that fractal features are uninformative and hence should not be used in classifiers for Pollock authentication. Recall that J-SMK showed that both genuine and fake Pollocks could share the same fractal properties, that a few Pollocks failed the fractal test, and that highly artificial images passed this test; thus they argued that the fractal feature was “useless” or “uninformative” and should not be used “even as a supplement to other analysis.”

3.1 Genuine and fake Pollock works have similar fractal features

Consider the J-SMK “uninformative” fractal claim first from the perspective of feature selection in pattern classification theory. Figure 2 shows hypothetical distributions of two categories in a two-dimensional visual feature space $x_1 \times x_2$, where patterns in category ω_1 might represent genuine Pollocks and ω_2 might represent fakes. The densities in a lower-dimensional space, here one-dimensional, are given by

$$p(x_1|\omega_i) = \int p(x_1, x_2|\omega_i) dx_2, \tag{1}$$

a process called “marginalization.” (Equation 1 shows the full class-conditional probability density functions marginalized over the x_2 feature.)

Suppose x_1 in Fig. 2 represents some fractal feature. When marginalized over the other, x_2 feature, both categories have the *same* distributions—thus the x_1 feature, taken alone, provides no discriminative information and is then “useless,” much as J-SMK report. Likewise, in the case shown the x_2 feature is “useless.” Nevertheless, a classifier based on *both* features can be extremely accurate, as illustrated by the black oval decision boundary.

More generally, if the class-conditional densities are obey $p(x_1|\omega_1) = p(x_1|\omega_2)$, then feature x_1 *taken alone* indeed provides no discriminative power—it is “useless.” This does not mean, however, that $p(x_j|\omega_1) = p(x_j|\omega_2)$ for $j \neq 1$, that is, for other features. More importantly, it does not imply that $p(\mathbf{x}|\omega_1) = p(\mathbf{x}|\omega_2)$ in the full feature space and that the x_1 feature cannot improve classifier accuracy. Any difference in distribution means the Bayes error rate (the lowest error rate achievable by any classifier) is lower than that of chance performance. This same principle generalizes to multi-dimensional subspaces Ξ^i and Ξ^j , where the full input space is $R^n = \Xi^i \times \Xi^j$ and $i + j = n$. Perhaps paradoxically, a classifier based on an arbitrary number of features *each* of which is “useless” can have extremely high accuracy, as for instance the multi-dimensional version of the exclusive-OR problem shown in Fig. 2.

Note too that the x_2 feature in Fig. 2 need not represent a *visual* feature. It might represent some material property of the painting in question, for example whether the canvas is sized (coated—plotted as low x_2 values) or not. Thus the data in that figure would be summarized: Genuine Pollock paintings that are sized have low

Figure to be inserted later

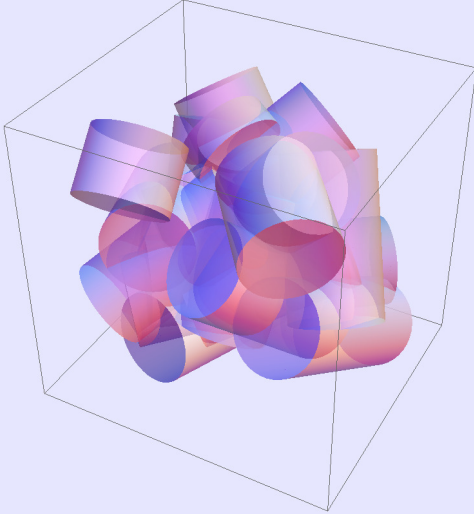


Figure to be inserted later




Figure 2. In this idealized two-category classification problem, the white points represent patterns in one category, ω_1 (e.g., genuine Pollocks), and the black points another category, ω_2 (non-Pollocks), in a two-dimensional feature space $x_1 \times x_2$. (This is the statistical generalization of the simplest non-linearly separable problem, the exclusive-OR problem.) The projections of each of the two distributions onto the x_1 axis of the two class-conditional distributions are the same, and thus any classification based on this feature alone would yield a 50% accuracy, i.e., chance performance. The x_1 feature, taken alone, is thus “useless.” Likewise, the projections onto the x_2 are the same, and hence x_2 , taken alone, is similarly “useless.” However, a nonlinear classifier based on *both* features (such as the one described by the black oval decision boundary) gives nearly 100% accuracy. Thus, even if x_1 represented some measure of the fractal properties of Pollock’s drip paintings and *taken alone* provides no information for classification, this feature might be useful when used in a classifier based on multiple features.

fractal feature value and those that are not sized have high fractal feature value while (for some reason) fake Pollock paintings have the converse. As before, either feature alone is “useless,” but here the *sole visual feature* describing fractal properties is sufficient to enable the two-feature classifier to have high accuracy.

Figure 3 illustrates a more common and realistic case, where the underlying probability distributions are Gaussians with the same covariance matrix but different means, here separated in just the x_2 direction. As in Fig. 2, the distributions marginalized over the x_2 feature are the same—that is, $p(x_1|\omega_1) = p(x_1|\omega_2)$ —and hence the x_1 feature taken alone is “useless.” Unlike Fig. 2, the projections onto the x_2 axis are not the same, though they overlap significantly. Both these examples show, however, that the “useless” feature improves classification dramatically when used in a multi-feature classifier.

As the introduction to the leading book on feature selection in pattern recognition stresses: “features that are not individually relevant may become relevant in the context of others.” [?, p. 9]

The class-conditional densities in Fig. 3 are bi-variate Gaussians with the same (arbitrary) covariance matrix Σ but different means, i.e., $p(\mathbf{x}|\omega_i) \sim N(\mu_i, \Sigma)$. The Bayes optimal decision boundary in this case is linear and given by [1, p. 40]

$$\mathbf{w}^t(\mathbf{x} - \mathbf{x}_0) = 0, \tag{2}$$

where $P(\omega_i)$ are the prior probabilities,

$$\mathbf{w} = \Sigma^{-1}(\mu_1 - \mu_2), \tag{3}$$

and

Figure to be inserted later




Figure to be inserted later




Figure 3. A two-category classification problem where the class-conditional distributions are two-dimensional Gaussians or normal distributions with the same covariance but different means: $p(\mathbf{x}|\omega_i) \sim N(\mu_i, \Sigma)$. The one-dimensional densities $p(x_1|\omega_i)$, given by Eq. 1, are shown below, and analogously for $p(x_2|\omega_i)$ at the left. Here, the Bayes error rate for classification based solely on the x_1 feature is 50%—chance performance. The Bayes error rate for classification based solely on the x_2 feature is roughly 15%. The Bayes optimal decision boundary in the full $x_1 \times x_2$ space is marked by the diagonal black line, and its corresponding Bayes error rate is roughly 1.8%. In short, the inclusion of the “useless” x_1 feature lowers the Bayes error from 15% to 1.8%.

$$\mathbf{x}_0 = \left[\frac{1}{2} - \frac{\ln[P(\omega_1)/P(\omega_2)]}{(\mu_1 - \mu_2)^t \Sigma^{-1} (\mu_1 - \mu_2)} \right] (\mu_1 - \mu_2). \quad (4)$$

This Bayes decision boundary is shown as the black line.

More generally, there is an infinite number of distributions that have the same projection onto a given feature axis (or indeed set of axes) but for which the Bayes classifier must use the “useless” feature to achieve lowest error. If the full n -dimensional densities are known, then the Bayes error can never become lower as additional features are used in the classifier. [1, Sect. 3.7 and Fig. 3.3]

Even highly noisy or unreliable data is frequently useful in pattern classification and most classifiers can even learn in the presence of “nasty” or “malicious” noise.^{?,?} One simply cannot claim a classifier is useless because some feature it computes is noisy or unreliable.

The visual authentication of Pollock paintings is an instance of the well-explored field of *visual texture classification* and one would expect that proponents and opponents of any methods applied to Pollock authentication would build upon the results of this discipline, specifically its universal use of *multiple* features for classification. Note especially that an early overview of the state-of-the-art in visual texture recognition considered eighteen leading methods and found that only one method used as few as five features, but others use up to 47 features, with the average being 18 features. [?, 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.

A recent empirical study showed the value of such “useless” scale-space visual features for the authentication of drip paintings. Irfan and Stork trained standard classifiers—specifically the Perceptron and nearest-neighbor classifiers—to distinguish genuine from fake Pollocks using a fractal feature plus four other features (Levy dimension, genus, and two features based on oriented visual energy).[?] They found that in one classifier while a scale-space feature alone provided only slightly better than chance performance (52.4% accuracy) and that

the other features provided 76.2% accuracy, all five features together provided 81.0% accuracy—not unlike the case illustrated in Fig. 3. Moreover, in the nearest-neighbor classifier, the basic fractal feature provided some classification accuracy, 66.7%, which was improved to 76.2% by the inclusion of other visual features. In short, scale-space features may indeed be of some use in multi-feature classifiers for authenticating Pollock paintings. In short, the empirical evidence, collected using appropriate protocols from statistical pattern classification, demonstrates both that the fractal (or scale-space) feature taken *alone* provides some classification benefit in the simple linear classifier, and that in both classifiers tested, the fractal feature improves the accuracy in the presence of other features. More importantly, it shows that classifiers designed using the methods from statistical pattern classification can perform better than chance, and thus may indeed have value in authentication studies.

For these reasons, the J-SMK empirical analyses of Pollocks and recently discovered works in the collection of Alex Matter of doubtful authorship⁷ based on a *single* feature are not sufficient to support their explicit claim that the fractal feature should not be part of *any* Pollock authentication regimen, “even as a supplement.”

3.2 Highly artificial images satisfy the fractal criterion for genuine Pollocks

J-SMK showed highly artificial sketches—Jones-Smith’s *Gross pebbles* and *Mixed stars* (2006), created in *Photoshop*—that match the fractal characteristics of genuine Pollocks and they argue, therefore, against the value of such features in Pollock authentication studies.

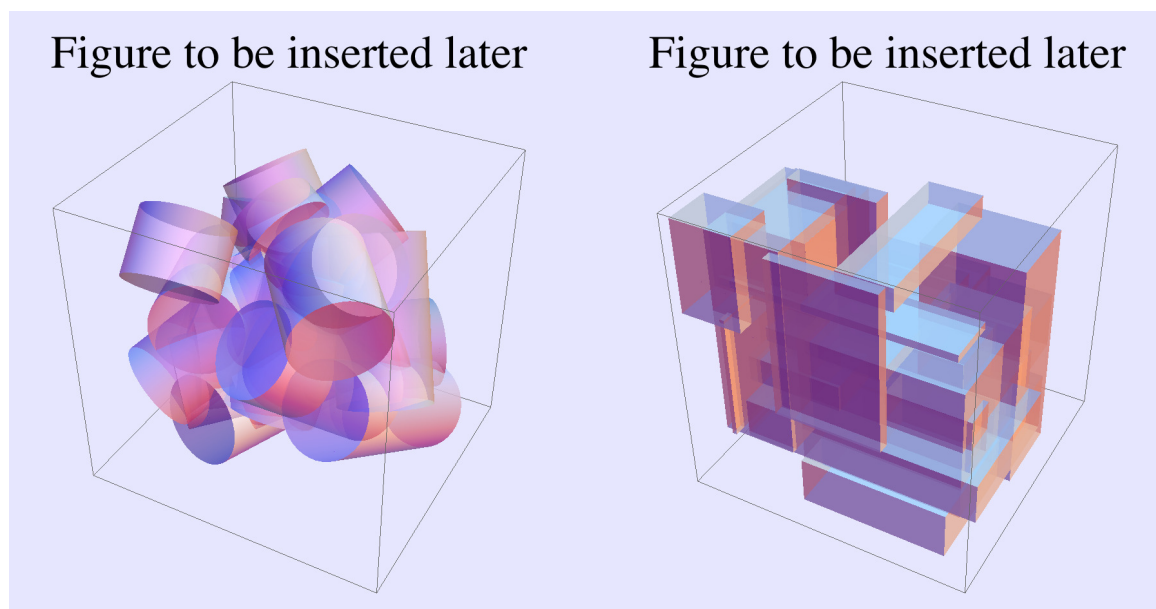


Figure 4. Although patterns generally lie in a high-dimensional space (three are illustrated), classification occurs in a lower-dimensional subspace or manifold, often nonlinear, such as the two-dimensional manifold here. The black line is the decision boundary defined on this manifold. A pattern far from the distributions used to train the classifier (large black dot) will nevertheless be placed in the same category as patterns that appear rather differ from it. The existence of such patterns does not invalidate the accuracy or usefulness of the classifier for the target distributions for which it was trained, here the small dots close to the manifold.

Consider first this claim from the perspective of the theory of pattern classification. Figure 4 illustrates the basic phenomenon. With very rare exceptions, pattern classification involves the first step of feature extraction, or representation of a pattern in a high-dimensional space in a space of lower dimensions—typically much lower. Classification decisions and decision boundaries are then defined in this feature space. In this common framework, the feature space typically passes through the data. Common Principal Components Analysis (PCA) finds the linear subspace that minimizes the squared distances of the patterns to the subspace, for example. This feature space and the decision boundary defined within it are created for the *conditions and distributions for the classifier is intended*, here, dripped paintings made physically at a particular range of scales, colors, and so forth. Proper




Figure 5. The classification results (in gray) of a highly trained and widely used English language optical character recognition (OCR) system applied to characters *not* from the target distributions of patterns. The Japanese katakana, Japanese hiragana, mathematical symbols and Greek lower-case letters are all classified as English characters that bear scant resemblance. Such eccentric or deviant test patterns, which the classifier was never meant to classify, always exist in any complex pattern classification domain, and such classification results do not imply that the classifier is useless or inappropriate for its target problem. In short, the fact that J-SMK can construct highly artificial computer sketches—regardless the “intent” of their creator—that are classified as Pollock paintings in no way invalidates such a classifier when applied to the proper, target classification distributions, in particular dripped paintings of the requisite size, physical composition, physical mass, viscosity, colors, and so on.

pattern recognition methodology is to choose samples *IID* (*independent identically distributed*), that is, the test patterns should be drawn from the same distribution as the training data in the high-dimensional subspace. That is, the training and test patterns must be from the *same* distributions.¹ Of course, one can always construct patterns that deviate from these distributions that are nevertheless classified by the system, but that does not imply that the classifier is useless for the distributions for which it was designed.

The Taylor et al. method was never meant to address digital images, made with a computer mouse, with lines having no splatter, and so on. To state that a test pattern selected non-IID does not appear similar to patterns chosen IID from the category to which the test pattern is assigned somehow undermines the classifier criterion is misleading at best. Any rejection of a classifier relying on such logic is invalid.

Consider now a practical illustration of this rebuttal to the J-SMK claim. Figure 5 shows the classification results of a highly trained, widely used, accurate English language optical character recognition (OCR) system when presented with patterns quite far from the relevant set, that is, *not* IID. The fact that each non-English character or mathematical symbol is classified as an apparently dissimilar English character in no way obviates the usefulness of the classifier for the target distributions for which the classifier is intended and on which it was trained. It matters not the motivations or intent or belief of the agent providing the novel test pattern; if the novel test pattern was not selected IID from the appropriate full distributions, or violate the conditions used in creating the classifier (*viz.*, that the works are physical paintings, for example), the classification results are not representative of the system as it was designed and intended.

3.3 Some genuine Pollocks fail the fractal criterion and some fake Pollocks pass the criterion

J-SMK showed that certain Pollock paintings, such as *The wooden horse: Number 10A* (1948) and *Free form* (1946), failed the Taylor et al. fractal criterion and that some fake drip paintings passed the criterion. Of course,

virtually no pattern recognition system yields perfect classification; the classifier need only be better than chance to provide some authentication value, and this requires more analysis than small-sample non-IID demonstrations. (Elsewhere in the realm of art authentication, the fact that some genuine signatures are deemed by experts to be fakes and some fake signatures are deemed to be genuine does not mean that signatures should never be used in art authentication, of course. All that is necessary is that such expert judgements be more accurate than chance, especially when used in a regimen of several techniques. The proper way to address this question is through careful empirical work on a large corpus of appropriate paintings. Only if such a trained classifier—here using multiple features—has an error rate indistinguishable from chance would it mean the classifier is useless and that a particular feature should not be used in authentication studies. Small-sample “tests”— more properly called “demonstrations”— using improper sampling methods, must be supplanted by statistical tests using IID samples and techniques such as corss-validation and leave-one-out evaluation. As mentioned, preliminary tests using multiple features, cross-validation, and other standard techniques from statistical pattern classification have shown the value of such features in classifiers. More extensive tests along these lines are in progress.[?]

4. REBUTTAL TO THE CLAIM THAT BECAUSE OCCLUSION DISRUPTS FRACTAL PROPERTIES, BOX-COUNTING SHOULD NOT BE USED

J-SMK point out that the fractal properties of the visible portion of one image are disrupted or destroyed if that image is partially occluded by another image. In the case of Pollock authentication, the fractal properties of the visible portion of one layer of dripped paint (green, say) are disrupted by a subsequently applied layer of another color (purple, say). J-SMK conclude thus that it is inappropriate that Taylor et al. applied their box-counting algorithm to the visible portions (color separations), and that the subsequent scale-space analysis is thus inappropriate.

Most classifiers employ *preprocessing* of sensed signals before feature extraction and ultimate classification [1, p. 2] and the automatic visual authentication of Pollock paintings is no exception. Indeed, J-SMK explore the Taylor et al. preprocessing step of color separation before concluding that fractal features are not useful. But one must consider other appropriate preprocessing methods before ruling out the possible value of any particular feature.

As such, there are two responses to this criticism by J-SMK: The first is that image preprocessing algorithms can estimate and hence recover some of the occluded information and the second is that the classifier may not need an accurate estimate of the full scale-space properties anyway. Shahram, Stork and Donoho developed the *De-pict* algorithm, which estimates the images of partially occluded layers of brush strokes, and demonstrated its power on *Self portrait in a grey felt hat* by Vincent van Gogh.[?] In brief, the core steps of the *De-pict* algorithm are (cf. Fig. 6):

- Step 1: Estimate the current number of layers** Use a clustering algorithm (such as *k*-means) to estimate the number of colors, and hence number of layers.
- Step 2: Determine which layer is currently on top** Compute the variance of several spatial statistics of each of the color layers; the color layer with the smallest such variance is likely on top, since its strokes have not been occluded and disrupted.
- Step 3: Output the current top layer** Save the image of the current top layer for scale-space analysis.
- Step 4: Digitally remove the pixels of the top layer** Remove all pixels of the color associated with the top layer identified in **Step 2**.
- Step 5: Fill in the regions exposed in Step 4** Digitally *inpaint* to fill in the previously occluded portions of the hidden layers.[?]
- Step 6: Iterate** If there remain pixels from the layers in the original image, let $k \leftarrow k - 1$ and **Go to Step 1**. (Iteration should occur *k* times, once for each layer.)




Figure 6. The goal of the `De-pict` algorithm is to recover partially hidden layers of brush strokes, including poured and dripped lines of paint. The left figure shows a schematic of a small passage in a final dripped painting where a black paint drip line occludes a lower line of gray paint. The `De-pict` algorithm clusters the image pixels by color to estimate the number of colors and hence layers (here two). Then it estimates the spatial statistics of each color layer to estimate which layer is on top—the more uniform these statistics, the more likely the layer is on top. Then the top layer is digitally removed, leaving a region of undefined color (middle figure). Then the strokes of remaining colors are *in-painted* into this region (right figure). Typical inpainting algorithms minimize a functional of the curvature of similar visible edges entering the region.

Thus, **Step 3** yields images of the estimated hidden portions of occluded layers; such images may indeed prove useful sources of features in classifiers for Pollock paintings.

The second response to the J-SMK criticism is that there is no need for the extracted information to be fractal anyway. Indeed, Irfan and Stork showed empirically that the scale-space features of partly occluded layers of drip paintings are indeed useful for classification of Pollock paintings even if they might not conform to the rigid definition of a fractal.[?]

5. REBUTTAL TO THE CLAIM THAT A NARROW SCALE RANGE PRECLUDES THE VALUE OF BOX-COUNTING FEATURES

As mentioned, Taylor and his colleagues estimated fractal dimensions through a box-counting algorithm. J-SMK pointed out that the range of scales employed by Taylor et al. is generally too narrow to reliably estimate true fractal properties: “analysis of a box-counting curve over less than two orders of magnitude, in the absence of any theoretical prediction of scaling behavior, is an insufficient criterion by which to establish fractality. *Fractal dimensions determined over such a limited range are meaningless.*” [?, emphasis added] While this criticism is technically valid, it is also irrelevant to the problem of pattern recognition, at least in theory. What is essential in pattern recognition is not that a feature conform to some crisp mathematical definition, but rather that it can be specified and computed in an objective, repeatable way and that its use improves classification accuracy. For instance, the features learned at the hidden layer of a traditional three-layer neural network are often extremely complicated and hard to interpret. [1, Sect. 6.5.1] All that is needed is that the information be useful in a classifier, a matter that can be answered only empirically. In fact, Irfan and Stork recently showed that such box-counting features indeed improved classification of Pollock versus non-Pollock paintings.[?]

6. CONCLUSION

The above analyses rebuts each of the primary arguments of J-SMK rejecting the use of fractal features as part of an authentication regimen for drip paintings such as by Jackson Pollock. Nevertheless their recommendation that fractal analysis, or more properly scale-space analysis, is not (yet) appropriate for the high-stakes challenges of art authentication is of course prudent regardless. Much work remains to be done before computer image methods can provide assistance to art scholars on Pollock authentication. The empirical results of Irfan and Stork, for example, were based on a small set of data, admittedly poor “fakes” of somewhat low resolution, and so on.² Reliable studies require more image data of higher resolution, more and better features, and the application of sophisticated techniques from machine learning such as boosting, bagging and cross-validation.¹ A promising approach is to train pattern recognition systems on both visual and *non-visual* features, such as material properties of the paint and media.

Even if fractal and image-based computer methods show value, they must be used carefully, with a deep understanding of their strengths and limitations, to extend and refine—not replace—expert judgements and as part of a broader authentication regimen based on material studies, connoisseurship, and so on. Nevertheless, there seems to be an adequate theoretical foundation and preliminary empirical evidence that such ongoing research may lead to systems that provide real value to the art community.^{?,?} xxx

ACKNOWLEDGMENTS

The second author thanks the Getty Research Institute, Santa Monica, for access to its superb reference library.

REFERENCES

1. R. O. Duda, P. E. Hart, and D. G. Stork, *Pattern classification*, John Wiley and Sons, New York, NY, Second ed., 2001.
2. R. P. Taylor, A. P. Micolich, and D. Jonas, “Fractal analysis of Pollock’s drip paintings,” *Nature* **399**, p. 422, 1999.
3. B. Mandelbrot, *The fractal geometry of nature*, W. H. Freeman, New York, NY, 1982.
4. 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.
5. K. Jones-Smith and H. Mathur, “Fractal analysis: Revisiting Pollock’s drip paintings,” *Nature* **444**(doi: 10.1038/nature05398), pp. E9–E10, 2006.
6. 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.
7. 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.
8. E. G. Landau and C. Cernuschi, eds., *Pollock/Matters*, McMullen Museum of Art, Boston, MA, 2007.
9. 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.
10. R. C. Gonzalez and R. E. Woods, *Digital image processing*, Prentice Hall, Upper Saddle River, NJ, Third ed., 2007.
11. D. G. Stork and J. Coddington, eds., *Computer image analysis in the study of art*, vol. 6810, SPIE/IS&T, Bellingham, WA, 2008.
12. D. G. Stork, J. Coddington, and A. Bentkowska-Kafel, eds., *Computer vision and image analysis in the study of art*, vol. xxx, SPIE/IS&T, Bellingham, WA, 2010.