

Digital Art Forensics

Siwei Lyu¹, Daniel Rockmore^{1,2}, and Hany Farid^{1,†}
Department of Computer Science¹ and Department of Mathematics²
Dartmouth College
Hanover NH 03755

We describe a computational technique for digitally authenticating works of art. This approach builds statistical models of an artist from a set of authenticated works. Additional works are then authenticated against this model. The statistical model consists of first- and higher-order wavelet statistics. We show preliminary results from our analysis of thirteen drawings by Pieter Bruegel the Elder. We also present preliminary results showing how these techniques may be applicable to determining how many hands contributed to a single painting.

[†] Correspondence should be addressed to H. Farid. 6211 Sudikoff Lab, Department of Computer Science, Dartmouth College, Hanover NH 03755. tel/fax: 603.646.2761/1672; email: farid@cs.dartmouth.edu.

1 Introduction

It probably wasn't long after the creation of paintings, sculptures, and other art forms that a lucrative business in art forgeries was found. And it probably wasn't long after this that techniques for detecting art forgeries emerged. Much of this work has been based on physical analyses (e.g., chemical dating, x-ray, etc.) [?]. With the advent of powerful digital technology it seems that computational tools can begin to provide new insights and tools into the art and science of art forgery detection (e.g., [5, 12, 11]).

We present a computational tool for analyzing prints, drawings and paintings for the purpose of characterizing their authenticity. More specifically we begin with high-resolution digital scans of a drawing or painting, perform a multi-scale, multi-orientation image decomposition (e.g., wavelets), construct a compact model of the statistics within this decomposition, and look for consistencies or inconsistencies across (or within) different drawings or paintings.

We first describe the underlying statistical model and then show preliminary results from our analysis of thirteen drawings by Pieter Bruegel the Elder and a painting by Perugino.

2 Wavelet Statistics

The decomposition of images using basis functions that are localized in spatial position, orientation, and scale (e.g., wavelets) has proven extremely useful in a range of applications (e.g., image compression, image coding, noise removal, and texture synthesis). One reason for this is that such decompositions exhibit statistical regularities that can be exploited (e.g., [9, 8, 2]). Described below is one such decomposition, and a set of statistics collected from this decomposition.

The decomposition is based on separable quadrature mirror filters (QMFs) [13, 14, 10]. As illustrated in Figure 1, this decomposition splits the

frequency space into multiple scales and orientations. This is accomplished by applying separable lowpass and highpass filters along the image axes generating a vertical, horizontal, diagonal and lowpass subband. For example, the horizontal subband is generated by convolving with the highpass filter in the horizontal direction and lowpass in the vertical direction, the diagonal band is generated by convolving with the highpass filter in both directions, etc. Subsequent scales are created by subsampling the lowpass by a factor of two and recursively filtering. The vertical, horizontal, and diagonal subbands at scale $i = 1, \dots, n$ are denoted as $V_i(x, y)$, $H_i(x, y)$, and $D_i(x, y)$, respectively. Shown in Figure 3 is a three-level decomposition of the image of Dartmouth Hall shown in Figure 2.

Given this image decomposition, the statistical model is composed of the mean, variance, skewness and kurtosis of the subband coefficients at each orientation and at scales $i = 1, \dots, n - 2$. These statistics characterize the basic coefficient distributions. In order to capture the higher-order correlations that exist within this image decomposition, these coefficient statistics are augmented with a set of statistics based on the errors in an optimal linear predictor of coefficient magnitude.

As described in [2], the subband coefficients are correlated to their spatial, orientation and scale neighbors. For purposes of illustration, consider first a vertical band, $V_i(x, y)$, at scale i . A linear predictor for the magnitude of these coefficients in a subset of all possible neighbors may be given by:

$$\begin{aligned} |V_i(x, y)| &= w_1|V_i(x - 1, y)| + w_2|V_i(x + 1, y)| \\ &+ w_3|V_i(x, y - 1)| + w_4|V_i(x, y + 1)| \\ &+ w_5|V_{i+1}(\frac{x}{2}, \frac{y}{2})| + w_6|D_i(x, y)| \\ &+ w_7|D_{i+1}(\frac{x}{2}, \frac{y}{2})|, \end{aligned} \quad (1)$$

where w_k denotes scalar weighting values, and $|\cdot|$ denotes magnitude. This particular choice of spatial, orientation, and scale neighbors was employed in our earlier work on detecting traces of digital tampering in images [4]. Here we employ an iterative brute-force search (on a per subband




Figure 1: An idealized multi-scale and orientation decomposition of frequency space. Shown, from top to bottom, are levels 0,1, and 2, and from left to right, are the lowpass, vertical, horizontal, and diagonal subbands.



Figure 2: An image of Dartmouth Hall.



Figure 3: Shown are the absolute values of the subband coefficients at three scales and three orientations for an image of Dartmouth Hall, Figure 2. The residual lowpass subband is shown in the upper-left corner.

and per image basis) for the set of neighbors that minimizes the prediction error within each sub-band.

Consider again the vertical band, $V_i(x, y)$, at scale i . We constrain the search of neighbors to a 3×3 spatial region at each orientation subband and at three scales, namely, the neighbors:

$$\begin{aligned}
 &V_i(x - c_x, y - c_y), H_i(x - c_x, y - c_y), \\
 &D_i(x - c_x, y - c_y), \\
 &V_{i+1}(\frac{x}{2} - c_x, \frac{y}{2} - c_y), H_{i+1}(\frac{x}{2} - c_x, \frac{y}{2} - c_y), \\
 &D_{i+1}(\frac{x}{2} - c_x, \frac{y}{2} - c_y), \\
 &V_{i+2}(\frac{x}{4} - c_x, \frac{y}{4} - c_y), H_{i+2}(\frac{x}{4} - c_x, \frac{y}{4} - c_y), \\
 &D_{i+2}(\frac{x}{4} - c_x, \frac{y}{4} - c_y),
 \end{aligned}$$

with $c_x = \{-1, 0, 1\}$ and $c_y = \{-1, 0, 1\}$, where $c_x, c_y \neq 0$. From these 80 possible neighbors, the iterative search begins by finding the single most predictive neighbor (e.g., $V_{i+1}(x/2 - 1, y/2)$)¹. This neighbor is held fixed and the next most predictive neighbor is found. This process is repeated five more times to find the optimally predictive neighborhood. On the k^{th} iteration, the predictor coefficients (w_1, \dots, w_k) are determined as follows. Let the vector \vec{V} contain the coefficient magnitudes of $V_i(x, y)$ strung out into a column vector, and the columns of the matrix Q contain the chosen neighboring coefficient magnitudes also strung out into column vectors. The linear predictor then takes the form:

$$\vec{V} = Q\vec{w}, \quad (2)$$

where the column vector $\vec{w} = (w_1 \dots w_k)^T$. The predictor coefficients are determined by minimizing the quadratic error function:

$$E(\vec{w}) = [\vec{V} - Q\vec{w}]^2. \quad (3)$$

This error function is minimized by differentiating with respect to \vec{w} :

$$dE(\vec{w})/d\vec{w} = 2Q^T[\vec{V} - Q\vec{w}], \quad (4)$$

¹Integer rounding is used when computing the spatial positions of a parent, e.g., $x/2$ or $x/4$.

setting the result equal to zero, and solving for \vec{w} to yield:

$$\vec{w} = (Q^T Q)^{-1} Q^T \vec{V}. \quad (5)$$

The log error in the linear predictor is then given by:

$$\vec{E}_v = \log_2(\vec{V}) - \log_2(|Q\vec{w}|). \quad (6)$$

Once the full set of neighbors is determined additional statistics are collected from the errors of the final predictor - namely the mean, variance, skewness, and kurtosis. This entire process is repeated for each oriented subband, and at each scale $i = 1, \dots, n - 2$, where at each subband a new set of neighbors is chosen and a new linear predictor estimated.

For a n -level pyramid decomposition, the coefficient statistics consist of $12(n - 2)$ values, and the error statistics consist of another $12(n - 2)$ values, for a total of $24(n - 2)$ statistics. These values represent the measured statistics of an artist and, as described below, are used to classify or cluster drawings or paintings.

3 Bruegel

Pieter Bruegel the Elder (1525/30-1569) was perhaps one of the greatest Dutch artists. Of particular beauty are Bruegel’s landscape drawings. We choose to begin our analysis with Bruegel’s work not only because of their exquisite charm and beauty, but also because Bruegel’s work has recently been the subject of renewed study and interest [7]. As a result many drawings formerly attributed to Bruegel are now considered to belong to others. As such, we believe that this is a wonderful opportunity to test and push the limits of our computational techniques.

We digitally scanned (at 2400 dpi) eight authenticated drawings by Bruegel and five forgeries from 35mm color slides, Figure 4 (slides were provided courtesy of the Metropolitan Museum of Art [7]). These color (RGB) images, originally of size 3894×2592 , were cropped to a central

Num.	Title	Artist
3	Pastoral Landscape	Bruegel
4	Mountain Landscape with Ridge and Valley	Bruegel
5	Path through a Village	Bruegel
6	Mule Caravan on Hillside	Bruegel
9	Mountain Landscape with Ridge and Travelers	Bruegel
11	Landscape with Saint Jermove	Bruegel
13	Italian Landscape	Bruegel
20	Rest on the Flight into Egypt	Bruegel
<hr/>		
7	Mule Caravan on Hillside	-
120	Mountain Landscape with a River, Village, and Castle	-
121	Alpine Landscape	-
125	Solicitud Rustica	-
127	Rocky Landscape with Castle and a River	-

Figure 4: Authentic (top) and forgeries (bottom). The first column corresponds to the catalog number in [7].

2048×2048 pixel region, converted to grayscale² (gray = $0.299R + 0.587G + 0.114B$), and autoscaled to fill the full intensity range $[0, 255]$. Shown in Figure 5 are examples of an authentic drawing and a forgery.

For each of 64 (8×8) non-overlapping 256×256 pixel region in each image, a five-level, three-orientation QMF pyramid is constructed, from which a 72-length feature vector of coefficient and error statistics is collected, Section 2.

In order to determine if there is a statistical difference between the eight authentic drawings and the five forgeries, we first computed the Hausdorff distance [6] between all 13 pairs of images. The resulting 13×13 distance matrix was then subjected to a multidimensional scaling (MDS)

²While converting from color to grayscale results in a significant loss of information, we did so in order to make it more likely that the measured statistical features and subsequent classification was more likely to be based on the artist’s strokes, and not on simple color differences.




Figure 5: Authentic #6 (top) and forgery #7 (bottom), see Table 4.




Figure 6: Results of analyzing 8 authentic Bruegel drawings (blue circles) and 5 forgeries (red squares). Note how the forgeries lie significantly outside of the bounding sphere of authentic drawings.

with a Euclidean distance metric [3]. Shown in Figure 6 is the result of visualizing the projection of the original 13 images onto the top-three MDS eigenvalue eigenvectors. The blue circles correspond to the authentic drawings, and the red squares to the forgeries. For purely visualization purposes, the wire-frame sphere is rendered at the center of mass of the eight authentic drawings and with a radius set to fully encompass all eight data points. Note that all five forgeries fall well outside of the sphere. The distances of the authentic drawings to the center of the sphere are 0.34, 0.35, 0.55, 0.90, 0.56, 0.17, 0.54, and 0.85. The distances of the forgeries are considerably larger at 1.58, 2.20, 1.90, 1.48, and 1.33 (the means of these two distance populations are statistically significant: $p < 1^{-5}$ (one-way anova)). Even in this reduced dimensional space, there is a clear difference between the authentic drawings and the forgeries.

4 Perugino

Pietro di Cristoforo Vannucci (Perugino) (1446-1523) is well known as a portraitist and a fresco painter, but perhaps he is best known for his altarpieces. By the 1490s Perugino maintained a workshop in Florence as well as in Perugia and was quite prolific. Shown in Figure 7 is the painting *Madonna With Child* by Perugino. As with many of the great Renaissance paintings, however, it is likely that Perugino only painted a portion this work - apprentices did the rest. To this end, we wondered if we could uncover statistical differences amongst the faces of the individual characters.

The painting (at the Hood Museum, Dartmouth College) was photographed using a large-format camera (8×10 inch negative) and drum-scanned to yield a color $16,852 \times 18,204$ pixel image. As in the previous section this image was converted to grayscale. The facial region of each of the six characters was manually localized. Each face was then partitioned into non-overlapping 256×256 regions and auto-scaled into the full intensity range



Figure 7: Madonna With Child by Perugino. How many hands contributed to this painting?




Figure 8: Results of analyzing the Perugino painting. The numbered data points correspond to the six faces (from left to right) in Figure 7. Note how the three left-most faces (1-3) cluster, while the remaining faces are distinct. This clustering pattern suggests the presence of four distinct hands.

[0, 255]. This partitioning yielded (from left to right) 189, 171, 189, 54, 81, and 144 regions. The same set of statistics as described in the previous section was collected from each of these regions. Also as in the previous section, we computed the Hausdorff distance between all six faces. The resulting 6×6 distance matrix was then subjected to MDS. Shown in Figure 8 is the result of visualizing the projection of the original six faces onto the top-three MDS eigenvalue eigenvectors.

The numbered data points correspond to the six faces (from left to right) in Figure 7. Note how the three left-most faces cluster, while the remaining faces are distinct. The average distance between faces 1 – 3 is 0.61, while the average distance between the other faces is 1.79. This clustering pattern suggests the presence of four distinct hands, and is consistent with the views of some art historians [1].

5 Discussion

We have presented a computational tool for digitally authenticating or classifying works of art. This technique looks for consistencies or inconsistencies in the first- and higher-order wavelet statistics collected from drawings or paintings (or portions thereof). We showed preliminary results from our analysis of thirteen drawings by Pieter Bruegel the Elder and a painting by Perugino. There is no doubt that much work remains to refine and further test these results, but we are very hopeful that these techniques will eventually play an important role in the ever-growing field of art forensics.

Acknowledgments

D. Rockmore has been supported by grant AFOSR F49620-00-1-0280. H. Farid has been supported by an Alfred P. Sloan Fellowship, an NSF CAREER Grant (IIS-99-83806), a Department of Justice Grant (2000-DT-CS-K001), and a departmental NSF Infrastructure Grant (EIA-98-02068).

References

- [1] Personal correspondence with Barton Thurber, Hood Museum, Dartmouth College.
- [2] R.W. Buccigrossi and E.P. Simoncelli. Image compression via joint statistical characterization in the wavelet domain. *IEEE Transactions on Image Processing*, 8(12):1688–1701, 1999.
- [3] T. Cox and M. Cox. *Multidimensional Scaling*. Chapman & Hall, London, 1994.
- [4] H. Farid and S. Lyu. Higher-order wavelet statistics and their application to digital forensics. In *IEEE Workshop on Statistical Analysis in Computer Vision (in conjunction with CVPR)*, Madison, WI, 2003.
- [5] I Fortes, R. Morales-Bueno, J.L. Perez de la Cruz, F. Triguero, and M.A. Comino. Obtaining mondrian-style paintings through probabilistic regular grammars. In *International Conference on Information Systems, Analysis and Synthesis*, 1998.
- [6] D.P. Huttenlocher, G.A. Klanderman, and W.J. Rucklidge. Comparing images using the hausdorff distance. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 15(9):850–863, 1993.
- [7] N.M. Orenstein, editor. *Pieter Bruegel the Elder*. Yale University Press, New Haven and London, 2001.
- [8] R. Rinaldo and G. Calvagno. Image coding by block prediction of multiresolution subimages. *IEEE Transactions on Image Processing*, 4(7):909–920, 1995.
- [9] J. Shapiro. Embedded image coding using zerotrees of wavelet coefficients. *IEEE Transactions on Signal Processing*, 41(12):3445–3462, 1993.
- [10] E.P. Simoncelli and E.H. Adelson. *Subband image coding*, chapter Subband transforms, pages 143–192. Kluwer Academic Publishers, Norwell, MA, 1990.
- [11] R. Taylor. Pollock, Mondrian and the nature: Recent scientific investigations. *Chaos and Complexity Letters*, in press.
- [12] R. Taylor, A.P. Micolich, and D. Jones. Fractal analysis of Pollock’s drip paintings. *Nature*, 399:422, 1999.
- [13] P.P. Vaidyanathan. Quadrature mirror filter banks, M-band extensions and perfect reconstruction techniques. *IEEE ASSP Magazine*, pages 4–20, 1987.
- [14] M. Vetterli. A theory of multirate filter banks. *IEEE Transactions on ASSP*, 35(3):356–372, 1987.