Dorée Duncan Seligmann
Avaya Labs

Computer Vision, Image Analysis, and Master Art: Part 3

David G. Stork
Ricoh Innovations

Marco F. Duarte
Rice University

As we saw in the first two parts of this series, we can apply computer vision and image analysis algorithms to art, notably realist paintings, to shed light on a number of problems in art history. We saw first how new uncalibrated methods for estimating perspective transformations lets us transform and view rendered objects from different positions and thus compare perspective-aligned passages within a single painting.[1] In the second part of the series,[2] we saw how algorithms for inferring the direction of illumination based on shading along an occluding contour can also be applied to paintings, revealing much about the working methods of artists, including the Baroque masters Georges de la Tour and Michelangelo Merisi da Caravaggio.

In this final part, we consider the problem of quantifying shape and form—or more specifically, quantifying the differences between the shapes of different contours. We shall see that we can use such quantitative methods to address a claim about the working methods of the early Renaissance master Jan van Eyck.

Representing shape

Everyone has an informal understanding of the meaning of shape, but even if we restrict our consideration to 2D plane figures, it's difficult to describe, represent, and quantify shape. We might use colloquial terms such as elongated, squashed, and so forth, but these are imprecise.

Art historians might describe a brushstroke or pencil line with terms such as fluid, jagged, or cramped. Moreover, the shape of a form or contour might depend on the scale: at a large scale, a contour might appear smooth and rounded, while at a smaller, finer scale, angular and jagged. In applying computers to the understanding of art, we'll need powerful representations and methods that not only describe the wealth of shapes we find in art, but also remain useful to art historians.

Let's restrict our attention to 2D binary shapes, as in Henri Matisse's elegant *Nu bleu IV* in Figure 1. In visual pattern recognition, shape representations fall into two general categories: region-based and contour-based. Some of the simple geometric region-based descriptions of a region *R* that have proven useful in image analysis include an exhaustive list of the pixels that comprise *R*, area (or, equivalently, the number of pixels in *R*), height and width, and a histogram of the region's projections onto principal or coordinate axes.[3] For further discussion of these descriptions, see the "Region-Based Descriptions" and "Contour-Based Descriptions" sidebars.

Shape descriptions and the analysis of art

The more closely a shape representation is matched to problems in art, the more useful it will be to art historians. Some of the terms noted in this article and the sidebars are easily understood and already used by art historians (direction, height, width, and elongatedness), but others seem quite removed from the concepts of art and thus of little value to such scholars (such as signature, higher moments, and chain codes). For instance, elongation might be a useful feature for expressing just one of the many differences between the elongated, stretched portraits by the Italian painter and sculptor Amedeo Modigliani

Editor's Note

This three-part series describes how a variety of methods adapted from computer vision, image analysis, and pattern recognition can be applied to visual arts and help answer questions in art history. In this final installment, David Stork discusses how shapes can be described. He outlines the challenges in quantifying shape and form analysis, and describes how techniques to compare shapes may be used to compare paintings to, for example, determine how a copy was made.

—*Dorée Duncan Seligmann*

1070-986X/07/$25.00 © 2007 IEEE

Published by the IEEE Computer Society

(1884–1920) and Domenikos Theotocopolous—better known as “El Greco” (d. 1614)—and the squashed, bulbous portraits of the Colombian painter Fernando Botero (b. 1932).

Complex or highly mathematical shape descriptions are occasionally useful in studying art. Consider the abstract expressionist Jackson Pollock’s *Lavender Mist* (see http://www.ibiblio.org/wm/paint/auth/pollock/lavender-mist/pollock.lavender-mist.jpg), a classic example of this artist’s later works, which he executed by pouring and dripping paint onto a canvas on the floor. It’s not obvious how we might describe in any detailed way this work’s complex, chaotic form, and any inherent structure. Nevertheless, recent research has done just that, through fractal analysis.

The fractal dimension describes how certain properties of a shape—such as the area and length of a boundary—change as a function of scale. Consider the total length of the coastline of England, as judged on maps of difference scales. On a map with a large scale, the coastline looks somewhat smooth. Suppose you count the number of 50-mile-long “steps” along the coastline needed to surround the island on the map. Denote that total distance as L_{50}.

Figure 1. Henri Matisse,* Nu bleu IV*, Nice, 1952, gouache on cut and pasted paper, 103 × 74 cm. (Photo credit: François Fernandez. Direction des Musées de France, gift of Jean Matisse, 1978, on deposit at the Musée Matisse, Nice.)

Next, imagine a larger, more detailed map of England. Repeat the aforementioned procedure,

Region-Based Descriptions

The following terms are used in image analysis for geometric region-based descriptions:

- *Eccentricity*: the dimensionless ratio of the length of the maximum chord, *K*, to the maximum chord perpendicular to *K*.
- *Elongatedness*: the ratio of the length to width of the minimal-area bounding rectangle.
- *Rectangularity*: the maximum ratio of the area’s region to that of a bounding rectangle over all possible rectangle orientations.
- *Direction*: the orientation of the minimum bounding rectangle, expressed as an angle with respect to the vertical.
- *Moments*: the sums over the region *R* of polynomials of the positions of pixels (assuming that regions have uniform areal density).
- *Euler-Poincaré characteristic or genus,* θ: the genus of a region *R* describes its topology or its connectedness and its number of holes.
- *Convex hull*: the shape of the minimal convex region containing *R*.
- *Signature*: for each point *p* on the region’s boundary, the normal chord is the line segment entirely within *R* that’s perpendicular to *R*’s boundary at *p*. The signature of *R* is the scalar length of this normal chord as a function of arc length around the contour.
- *Fractal dimension*: the fractal dimension of a complex region describes how its area or its length varies as a function of the scale.[1]

Some of these measures are invariant to transformations such as rotations and scale (for example, genus and eccentricity), while others are not (such as area and projections). *Mereology* is the study of the relationships between parts and wholes, and in pattern classification we often use hierarchical representations (trees or graphs) to describe how the full region might be broken down into subregions, sub-subregions, and so on.

Reference

1. B.H. Kaye, *A Random Walk through Fractal Dimensions*, 2nd ed., VCH Publishers, 1994.

Contour-Based Descriptions

Some of the contour-based descriptions of shape's boundary, *B*, include the following:

- *List of pixels*: this list might be expressed in rectilinear coordinates, polar coordinates, or other coordinates.
- *Boundary length*: the simple overall length of a boundary.
- *Curvature* (or radius of curvature): the full description of contour *B* consists of the radius of curvature as a function of position along the contour.
- *Bending energy*: the energy required to bend a stiff rod into the contour's shape (this measure is closely related to the curvature).
- *Chain code*: a chain code represents the successive directions of one-pixel steps along a boundary. Thus if U, D, R, and L represent a step up, down, right, and left, respectively, then the code UUUURRRDDDDLLL represents a simple 4- × 3-pixel vertically oriented rectangle.
- *Basis expansion*: we can represent a boundary as the linear combination of spatial basis functions—for instance B-splines, Fourier descriptors, and so on.
- *Connected segments*: a list of the segments of a polygonal shape.

but now count the number of 10-mile-long steps. The total length of the coastline you measure now, L_{10}, will surely be longer than L_{50} because you must step along the more complex, convoluted coastline visible in the second map.

Likewise, for ever-more detailed maps and shorter steps, the total length of the coastline you measure will be greater and greater. In fact, there is no simple, single answer to the question, What is the length of the coastline of England?

For many complex shapes, we can describe the increase of the measured boundary length as the step size's power-law function. For a simple square, this length doesn't increase with decreasing scale, but for the coast of England, it does. The fractal dimension of a boundary relates to the exponent in such a power law. A square's fractal dimension is low while the coast of England's is fairly high.

Another method for determining a region's fractal dimension is based on area. Here, rather than estimating a boundary's length by stepping along the boundary with ever-smaller step sizes, we can estimate the region's area.

Specifically, place a square grid of a given resolution over the region and count the number of component squares that contain any part of the region. Then, multiply that number by the area of each square. For a coarse square grid (for instance, consisting of a single square), the region's estimated area is large. For ever-finer and finer square grids, however, it becomes more likely that a small component square doesn't contain any part of the region. Thus, the region's total estimated area will be smaller. As with length, there's often a power-law dependency of the total area as a function of scale (component square size) for complex regions.

Richard Taylor and his colleagues have used this latter approach to estimate the fractal dimension of Pollock's drip paintings. (More specifically, they determined how the estimated fractal dimension depended on the scale.) They found that genuine Pollock paintings had a fairly distinctive signature or dependency in the fractal dimension as a function of scale—one that differed somewhat from those measured in counterfeit Pollock paintings.[4]

Jan van Eyck's *Portrait of Cardinal Niccolò Albergati*

If a digital image of the oil work is reduced and overlapped with a digital image of the silverpoint, we find excellent correspondence—that is, excellent fidelity. The central question is thus: How was that copy made? If the scale or magnification had been 1.0, then a number of methods might have been used—for instance, tracing onto thin paper then retracing onto the oak panel support. Another method is pouncing, where the artist pierces the original with small holes along the contours, places the original over the copy support, and forces charcoal dust through the holes to thereby mark the copy.

But neither of those methods can explain magnification (or minification) of a copy, such as we find in the van Eyck oil, which is roughly 40 percent larger in scale than the silverpoint. Alas, we have no documentary records of his method related to these works, though we have records of a number of techniques from that time and indeed earlier. One such technique employs the *Reductionszirkel* or reducing compass, a simple hinged mechanical device. The artist adjusts the separation between two of the legs to be the same as some chosen points on the original (such as the two eyes), then uses two other, mechanically linked legs, to mark the separation on the scaled copy. Another technique is copying or enlarging simply "by eye," that is, without any aids.

Recently, the artist David Hockney hypothesized that as early as the beginning of the Renaissance some painters executed their works by secretly tracing over images that were optically projected by a concave mirror or lens onto their supports (such as canvas, paper, and oak panel).[5] The Albergati portraits have been adduced as evidence for the theory, specifically through the claim that van Eyck copied the silverpoint using an epidiascope, or opaque projector.[6] There has been unanimous rejection by the independent scholarly community of this general tracing claim, at least for the early Renaissance,[7] and for the Albergati portraits in particular, in part based on the dramatic discovery of tiny pinprick holes in the silverpoint that indicated mechanical (not optical) methods were used.[8,9]

Another, admittedly partial, test of the tracing claim centers on verifying that a talented realist artist using mechanical devices from van Eyck's era can indeed achieve a fidelity that we find between the two works in Figure 2. Thus, we need principled methods for testing whether the distances between the silverpoint and a portrait made by mechanical methods is roughly the same as that between the silverpoint and van Eyck's oil.

How shall we measure such a difference between shapes?

Comparing shapes via the Chamfer distance

The computer vision and pattern recognition community developed the Chamfer distance, which is a principled method for quantifying the difference between the shapes of two contours.[10] Let's denote these contours C_1 and C_2. In brief, we take each point (pixel) on C_2 and find the nearest point on C_1, sum these distances, and divide them by the length of C_2. Thus, the Chamfer distance between C_1 and C_2 is roughly the average distance between a point on C_1 and its nearest point on C_2. (Of course, there may be several points on C_1 that have the same nearest point on C_2.)

We can illustrate the Chamfer distance calculation using the distance transform. The distance transform of a curve C assigns to every point in an area the distance to the nearest point on C. Figure 3 shows the distance transform of a skeletonized version of a portion of the Albergati silverpoint (which we denote as C_1): low values (blue) are close to the points on C_1, and red ones farther away. We can compute the Chamfer distance between another curve (C_2) and C_1 by convolving the pixel locations of C_2 with the distance transform of C_1 and then normalizing.

Figure 2. Jan van Eyck's **Portrait of Cardinal Niccolò Albergati**, *oil on wood, 34.1 × 27.3 cm (c. 1432). (Image credit: Kunsthistorisches Museum, Wien oder KHM, Wien.)*

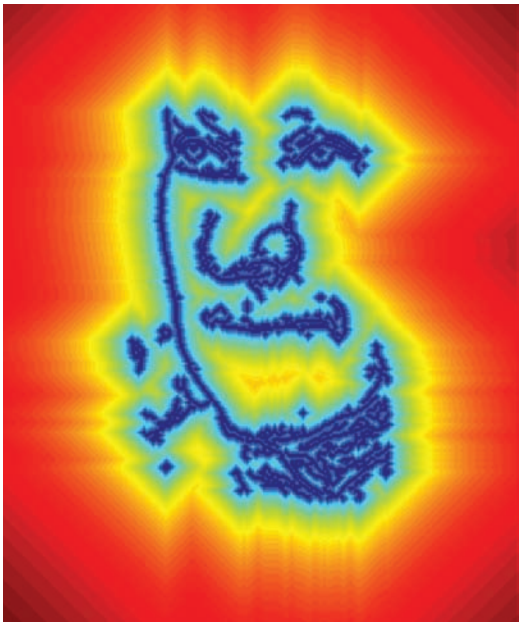

Figure 3. The distance transform applied to a line-thinned version of a portion of Jan van Eyck's **Albergati** *portrait silverpoint, C_1. The color of each pixel represents the perpendicular distance to a point on C_1. A measure of the Chamfer distance of one contour to another is the pixel average over C_2 of the shortest distance to a point on C_1.*

As part of our investigation of David Hockney's tracing theory,[5] we computed the Chamfer distance from the Albergati silverpoint and the following works:

- the van Eyck oil copy, and
- a copy/enlargement done by a professional artist using mechanical devices from van Eyck's day.

The full procedure required us to use thinning algorithms, to perform affine transformations (rotations, displacements, and uniform scaling), and other matters described in further detail elsewhere.[11] We found that a modern professional artist, using only mechanical copying/enlarging devices known from the time of van Eyck, could indeed achieve a fidelity (expressed as a Chamfer distance) roughly the same as van Eyck.

Our experimental result, historical information, and physical evidence (pinprick holes), led us to reject as unpersuasive the claim that van Eyck used an optical projector when copying/enlarging this work.

Future directions

There remains much work to be done on analyzing Jackson Pollock's drip paintings, for instance. A single shape descriptor such as the fractal dimension, taken alone, is unlikely to give the most reliable information for discriminating genuine Pollocks from forgeries. Perhaps additional shape descriptors such as curvature and connectedness (so-called connected components) will help discrimination. Moreover, any classification must rely on a sufficiently large statistical sample of genuine Pollocks and forgeries, and sophisticated pattern recognition algorithms. These are directions that might prove fruitful in the near future.

As the three articles in this series have shown, techniques from computer vision, image analysis, and pattern recognition hold promise for use in the study of visual arts. Art historians and computer vision experts will need to collaborate to understand the problems in art history that might be tackled or answered by computer methods, and to understand the power and the limitations of analytical methods. **MM**

Acknowledgments

We thank the Matisse Museum, Nice, France, for permission to reproduce *Nu bleu IV* and the Kunsthistorische Museum, Vienna, Austria, the home of the oil version of Jan van Eyck's *Portrait of Cardinal Niccolò Albergati*, analyzed here.

References

1. D.G. Stork, "Computer Vision, Image Analysis, and Master Art: Part 1," *IEEE MultiMedia*, vol. 13, no. 3, 2006, pp. 16-20.
2. D.G. Stork and M.K. Johnson, "Computer Vision, Image Analysis, and Master Art, Part 2," *IEEE MultiMedia*, vol. 13, no. 4, 2006, pp. 12-17.
3. M. Sonka, V. Hlavac, and R. Boyle, *Image Processing, Analysis and Machine Vision*, 2nd ed., PWS Publishing, 1999.
4. R.P. Taylor, A.P. Micolich, and D. Jonas, "Fractal Analysis of Pollock's Drip Paintings," *Nature*, vol. 399, 1999, p. 422.
5. D. Hockney, *Secret Knowledge: Rediscovering the Lost Techniques of the Old Masters*, Viking Studio, 2001.
6. D. Hockney and C.M. Falco, "Quantitative Analysis of Qualitative images," *Proc. SPIE Electronic Imaging*, SPIE Press, 2005.
7. *Early Science and Medicine*, special issue on optics, instruments and painting, 1420–1720: Reflections on the Hockney-Falco thesis, vol. 10, no. 2, 2005.
8. T. Ketelsen et al., "New Information on Jan van Eyck's Portrait Drawing in Dresden," *Burlington Magazine*, vol. CXLVII, no. 1224, 2005, pp. 169-175.
9. D.G. Stork, "Did Jan van Eyck Build the First 'Photocopier' in 1432?" *Proc. SPIE Electronic Imaging, Color Imaging IX: Processing, Hardcopy, and Applications*, R. Eschbach and G.G. Marcu, eds., SPIE Press, 2004, pp. 50-56.
10. M.A. Butt and P. Maragos, "Optimum Design of Chamfer Distance Transforms," *IEEE Trans. Image Processing*, vol. 10, no. 7, 1998, pp. 1477-1484.
11. D.G. Stork and M. Duarte, "Fidelity Analysis of Mechanically Aided Copying/Enlarging of Jan van Eyck's *Portrait of Niccolò Albergati*," *Proc. SPIE Electronic Imaging*, SPIE Press, 2007, in press.

Readers may contact David G. Stork at artanalyst@gmail.com.

Contact Artful Media editor Dorée Duncan Seligmann at doree@avaya.com.