

Invariant Histograms

Daniel Brinkman¹
Department of Applied Mathematics
and Theoretical Physics
University of Cambridge
Cambridge CB3 0WA
England
D.Brinkman@damtp.cam.ac.uk

Peter J. Olver¹
School of Mathematics
University of Minnesota
Minneapolis, MN 55455
U.S.A.
olver@math.umn.edu
<http://www.math.umn.edu/~olver>

Abstract

We introduce and study a Euclidean-invariant distance histogram function for curves. For a sufficiently regular plane curve, we prove that the cumulative distance histograms based on discretizing the curve by either uniformly spaced or randomly chosen sample points converge to our histogram function. Robustness of the curve histogram function under noise and pixelization of the curve is also established. We argue that the histogram function serves as a simple, noise-resistant shape classifier for regular curves under the Euclidean group of rigid motions. Extensions of the underlying ideas to higher-dimensional submanifolds, as well as to area histogram functions invariant under the group of planar area-preserving affine transformations, are discussed.

1 Introduction.

Given a finite set of points contained in \mathbb{R}^n , equipped with the usual Euclidean metric, consider the histogram formed by the mutual distances between all distinct pairs of points. An interesting question, first studied in depth by Boutin and Kemper, [5, 6], is to what extent the distance histogram uniquely determines the point set. Clearly, if the point set is subjected to a rigid motion — a combination of translations, rotations, and reflections — the interpoint distances will not change, and so two rigidly equivalent finite point sets have identical distance histograms. However, there do exist sets that have identical histograms but are not rigidly equivalent. (The reader new to the subject may enjoy trying to find an example before proceeding further.) Nevertheless, Boutin and Kemper proved that, in a wide range of situations, the set of such counterexamples is “small” — more precisely, it forms an algebraic subvariety of lower dimension in the space of all point configurations. Thus, one can say that, usually, the distance histogram uniquely determines a finite point set up to rigid equivalence. This motivates the use of the distance histogram as a simple, robust, noise-resistant signature that can be used to distinguish most rigidly inequivalent finite point sets, particularly those that arise as landmark points on an object in a digital image.

The goal of this paper is to develop a comparable distance histogram function for continua — specifically curves, surfaces, and higher-dimensional submanifolds of Euclidean spaces. Most of the paper, including all proofs, will concentrate on the simplest scenario: a “fully

¹Supported in part by NSF Grant DMS 08-07317.

regular” bounded plane curve. Regularity, as defined below, does allow corners, and so, in particular, includes polygons. We will approach this problem using the following strategy. We first sample the curve using a finite number of points, and then compute the distance histogram of the sampled point set. Stated loosely, our main result is that, as the curve becomes more and more densely sampled, the appropriately scaled cumulative distance histograms converge to an explicit function that we name the *global curve histogram function*. Alternatively, computing the histogram of distances from a fixed point on the curve to the sample points leads, in the limit, to a *local curve histogram function*, from which the global version can be obtained by averaging over the curve. Convergence of both local and global histograms is rigorously established, first, for uniformly sampled points that are separated by a common arc length distance; second, for randomly sampled points; and, finally, for “pixelized curves”, where we discretize using pixel coordinates in a digital representation of the curve.

The global curve histogram function can be computed directly through an explicit arc length integral. By construction, it is invariant under rigid motions. Hence, a basic question arises: does the histogram function uniquely determine the curve up to rigid motion? While there is ample evidence that, under suitably mild hypotheses, such a result is true, we have been unable to establish a complete proof, and so must state it as an open conjecture. A proof would imply that the global curve histogram function, as approximated by its sampled point histograms, can be unambiguously employed as an elementary, readily computed classifier for distinguishing shapes in digital images, and thus serve as a much simpler alternative to the joint invariant signatures proposed in [17]. Extensions of these ideas to subsets of higher-dimensional Euclidean spaces, or even general metric spaces, are immediate. Moreover, convergence in sufficiently regular situations can be established along the same lines as the planar curve case treated here.

Following Boutin and Kemper [5], we also consider area histograms formed by triangles whose corners lie in a finite point set. In two dimensions, area histograms are invariant under the group of equi-affine (meaning area-preserving affine) transformations. We exhibit a limiting area histogram function for plane curves that is also equi-affine invariant, and propose a similar conjecture. Generalizations to other transformation groups, e.g., similarity, projective, conformal, etc., of interest in image processing and elsewhere, [11, 18], are worth developing. The corresponding discrete histograms will be based on suitable joint invariants — for example, area and volume cross ratios in the projective case — which can be systematically classified by the equivariant method of moving frames, [17]. Analysis of the corresponding limiting histograms will be developed elsewhere.

Our study of invariant histogram functions has been motivated in large part by the potential applications to object recognition, shape classification, and geometric modeling. Discrete histograms appear in a broad range of powerful image processing algorithms: shape representation and classification, [1, 25], image enhancement, [23, 25], the scale-invariant feature transform (SIFT), [12, 20], object-based query methods, [24], and as integral invariants, [13, 21]. They provide lower bounds for and hence establish stability of Gromov–Hausdorff and Gromov–Wasserstein distances, underlying an emerging new approach to shape theory, [14, 15]. Local distance histograms underly the method of shape contexts, [3]. The method of shape distributions, [19] for distinguishing three-dimensional objects relies on a variety of invariant histograms, including local and global distance histograms, based on the fact

that objects with different Euclidean-invariant histograms cannot be rigidly equivalent; the converse, however, was not addressed. Indeed, there are strong indications that the distance histogram alone is insufficient to distinguish surfaces, although we do not know explicit examples of rigidly inequivalent surfaces that have identical distance histograms.

2 Distance Histograms.

Let us first review the results of Boutin and Kemper, [5, 6], on distance histograms defined by finite point sets. For this purpose, our initial setting is a general metric space V , equipped with a distance function $d(z, w) \geq 0$, for $z, w \in V$, satisfying the usual axioms.

Definition 1. The *distance histogram* of a finite set of points $P = \{z_1, \dots, z_n\} \subset V$ is the function $\eta = \eta_P: \mathbb{R}^+ \rightarrow \mathbb{N}$ defined by

$$\eta(r) = \# \{ (i, j) \mid 1 \leq i < j \leq n, d(z_i, z_j) = r \}. \quad (2.1)$$

In this paper, we will restrict our attention to the simplest situation, when $V = \mathbb{R}^m$ is endowed with the usual Euclidean metric, so $d(z, w) = \|z - w\|$. We say that two subsets $P, Q \subset V$ are *rigidly equivalent*, written $P \simeq Q$, if we can obtain Q by applying an isometry to P . In Euclidean geometry, isometries are *rigid motions*: the translations, rotations, and reflections generating the Euclidean group, [27]. Clearly, any two rigidly equivalent finite subsets have identical distance histograms. Boutin and Kemper’s main result is that the converse is, in general, false, but is true for a broad range of generic point configurations.

Theorem 2. Let $\mathcal{P}^{(n)} = \mathcal{P}^{(n)}(\mathbb{R}^m)$ denote the space of finite (unordered) subsets $P \subset \mathbb{R}^m$ of cardinality $\#P = n$. If $n \leq 3$ or $n \geq m + 2$, then there is a Zariski dense open subset $\mathcal{R}^{(n)} \subset \mathcal{P}^{(n)}$ with the following property: if $P \in \mathcal{R}^{(n)}$, then $Q \in \mathcal{P}^{(n)}$ has the same distance histograms, $\eta_P = \eta_Q$, if and only if the two point configurations are rigidly equivalent: $P \simeq Q$.

In other words, for the indicated ranges of n , unless the points are constrained by a certain algebraic equation, and so are “non-generic”, the distance histogram uniquely determines the point configuration up to a rigid motion. Interestingly, the simplest counterexample is not provided by the corners of a regular polygon. For example, the corners of a unit square have 4 side distances of 1 and 2 diagonal distances of $\sqrt{2}$, and so its distance histogram has values $\eta(1) = 4$, $\eta(\sqrt{2}) = 2$, while $\eta(r) = 0$ for $r \neq 1, \sqrt{2}$. Moreover, this is the only possible way to arrange four points with the given distance histogram. A simple nongeneric configuration is provided by the corners of the kite and trapezoid quadrilaterals shown in Figure 1. Although clearly not rigidly equivalent, both point configurations have the same distance histogram, with nonzero values $\eta(\sqrt{2}) = 2$, $\eta(2) = 1$, $\eta(\sqrt{10}) = 2$, $\eta(4) = 1$. A striking one-dimensional counterexample, discovered in [4], is provided by the two sets of integers $P = \{0, 1, 4, 10, 12, 17\}$ and $Q = \{0, 1, 8, 11, 13, 17\}$, which, as the reader can check, have identical distance histograms, but are clearly not rigidly equivalent.

To proceed, it will be more convenient to introduce the (renormalized) *cumulative distance histogram*

$$\Lambda_P(r) = \frac{1}{n} + \frac{2}{n^2} \sum_{s \leq r} \eta_P(s) = \frac{1}{n^2} \# \{ (i, j) \mid d(z_i, z_j) \leq r \}, \quad (2.2)$$

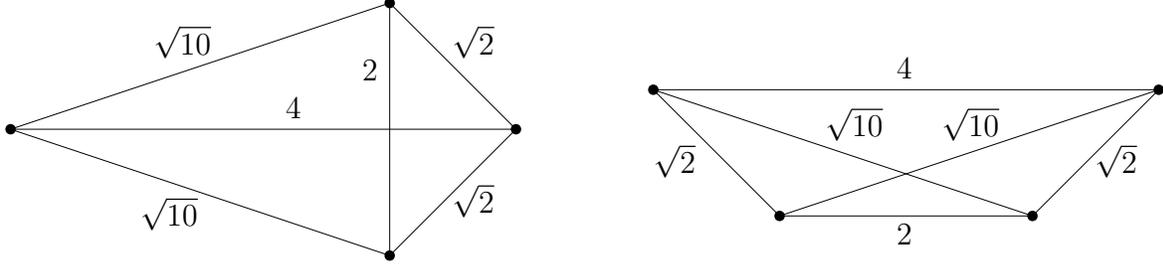


Figure 1: Kite and Trapezoid.

where $n = \#P$. We note that we can recover the usual distance histogram (2.1) via

$$\eta(r) = \frac{1}{2} n^2 [\Lambda_P(r) - \Lambda_P(r - \delta)] \quad \text{for sufficiently small } \delta \ll 1. \quad (2.3)$$

We further introduce a *local distance histogram* that counts the fraction of points in P that are within a specified distance r of a given point $z \in \mathbb{R}^m$:

$$\lambda_P(r, z) = \frac{1}{n} \# \{ j \mid d(z, z_j) \leq r \} = \frac{1}{n} \#(P \cap B_r(z)), \quad (2.4)$$

where

$$B_r(z) = \{ v \in V \mid d(v, z) \leq r \}, \quad S_r(z) = \partial B_r(z) = \{ v \in V \mid d(v, z) = r \} \quad (2.5)$$

denote, respectively, the ball (in the plane, the disk) of radius r centered at the point z and its bounding sphere (circle). Observe that we recover the cumulative histogram (2.2) by averaging its localization:

$$\Lambda_P(r) = \frac{1}{n} \sum_{z \in P} \lambda_P(r, z) = \frac{1}{n^2} \sum_{z \in P} \#(P \cap B_r(z)). \quad (2.6)$$

In this paper, we are primarily interested in the case when the points lie on a curve. Until the final section, we restrict our attention to plane curves: $C \subset V = \mathbb{R}^2$. A finite subset $P \subset C$ will be called a set of *sample points* on the curve. We will assume throughout that the curve C is bounded, rectifiable, and, usually (although not necessarily), closed. Further mild regularity conditions will be introduced below. We use $z(s)$ to denote the arc length parametrization of C , measured from some base point $z(0) \in C$. Let

$$l(C) = \int_C ds < \infty \quad (2.7)$$

denote the curve's length, which we always assume to be finite.

Our aim is to study the limiting behavior of the cumulative histograms constructed from more and more densely chosen sample points. It turns out that, under reasonable assumptions, the discrete histograms converge, and the limiting function can be explicitly characterized as follows.

Definition 3. Given a curve $C \subset V$, the *local distance histogram function* based at a point $z \in V$ is

$$h_C(r, z) = \frac{l(C \cap B_r(z))}{l(C)}, \quad (2.8)$$

i.e., the fraction of the total length of the curve that is contributed by those parts contained within the disk of radius r centered at z . The *global distance histogram function* of C is obtained by averaging the local version over the curve:

$$H_C(r) = \frac{1}{l(C)} \int_C h_C(r, z(s)) ds. \quad (2.9)$$

Observe that both the local and global distance histogram functions have been normalized to take values in the interval $[0, 1]$. The global function (2.9) is invariant under rigid motions, and hence two curves that are rigidly equivalent have identical global histogram functions. An interesting question, which we consider in some detail towards the end of the paper, is whether the global histogram function uniquely characterizes the curve up to rigid equivalence.

Modulo the definition of “fully regular”, to be presented in Section 4, our main result can be stated as follows. (See Section 5 for details on how “randomly chosen points” are selected.)

Theorem 4. *Let C be a regular plane curve. Then, for both uniformly spaced and randomly chosen sample points $P \subset C$, the cumulative local and global histograms converge to their continuous counterparts:*

$$\lambda_P(r, z) \longrightarrow h_C(r, z), \quad \Lambda_P(r) \longrightarrow H_C(r), \quad (2.10)$$

as the number of sample points goes to infinity.

3 Local Histogram Functions.

Our proof of Theorem 4 begins by establishing convergence of the local histograms. In this section, we work under the assumption that the sample points are uniformly spaced with respect to arc length along the curve.

Let us recall some basic terminology concerning plane curves, mostly taken from Guggenheimer’s book, [10]. We will assume throughout that $C \subset \mathbb{R}^2$ has a piecewise C^2 arc length parametrization $z(s)$, where s belongs to a bounded closed interval $[0, L]$, with $L = l(C) < \infty$ being its overall length. The curve is always assumed to be *simple*, meaning that there are no self-intersections, and either *closed*, i.e., a Jordan curve, or what we will call a *curve segment* that has distinct endpoints $z(0) \neq z(L)$. By convention, we will also designate a single point to be a segment of length 0. We use $t(s) = z'(s)$ to denote the unit tangent, and¹ $\kappa(s) = z'(s) \wedge z''(s)$ the signed curvature at the point $z(s)$. Under our assumptions, both $t(s)$ and $\kappa(s)$ have left- and right-hand limiting values at their finitely many discontinuities.

¹The symbol \wedge denotes the two-dimensional cross product, which is the scalar $v \wedge w = v_1 w_2 - v_2 w_1$ for $v = (v_1, v_2)$, $w = (w_1, w_2)$.

A point $z(s) \in C$ where either the tangent or curvature is not continuous will be referred to as a *corner*.

A closed curve is called *convex* if it bounds a convex region in the plane. A curve segment is *convex* if the region bounded by it and the straight line segment connecting its endpoints is a convex region. A curve segment is called a *spiral arc* if the curvature function $\kappa(s)$ is continuous, strictly monotone², and of one sign, i.e., either $\kappa(s) \geq 0$ or $\kappa(s) \leq 0$. Keep in mind that, by strict monotonicity, $\kappa(s)$ is only allowed to vanish at one of the endpoints of the spiral arc.

Definition 5. A plane curve is called *regular* if it is piecewise C^2 and the union of a finite number of convex spiral arcs, circular arcs, and straight lines.

Thus, any regular curve has only finitely many *corners*, finitely many *inflection points*, where the curvature has an isolated zero, and finitely many *vertices*, meaning points where the curvature has a local maximum or minimum, but is not locally constant. In particular, polygons are regular, as are piecewise circular curves, also known as *biarcs*, [16]. (But keep in mind that our terminological convention is that polygons and biarcs have corners, not vertices!) Examples of irregular curves include the graph of the infinitely oscillating function $y = x^5 \sin 1/x$ near $x = 0$, and the nonconvex spiral arc $r = e^{-\theta}$ for $0 \leq \theta < \infty$, expressed in polar coordinates.

Theorem 6. *If C is a regular plane curve, then there is a positive integer m_C such that the curve's intersection with any disk having center $z \in C$ and radius $r > 0$, namely $C \cap B_r(z)$, consists of at most m_C connected segments. The minimal value of m_C will be called the circular index of C .*

Proof: This is an immediate consequence of a theorem of Vogt, [26] — see also [10, Exercise 3-3.11] — that states that a convex spiral arc and a circle intersect in at most 3 points. Thus, $m_C \leq 3j + 2k$, where j is the number of convex spiral arcs, while k is the number of circular arcs and straight line segments needed to form C . *Q.E.D.*

Example 7. Let C be a rectangle. A disk $B_r(z)$ centered at a point $z \in C$ will intersect the rectangle in either one or two connected segments; see Figure 2. Thus, the circular index of a rectangle is $m_C = 2$.

For each positive integer n , let $P_n \subset C$ denote a collection of uniform sample points separated by a common arc length spacing $\Delta l = L/n$. Observe that when C is a closed curve, n equals the number of sample points, while when C is a nonclosed curve segment, there are $n + 1$ sample points including the two endpoints.

Proposition 8. *Let C be a regular curve. Then, for any $z \in C$ and $r > 0$, the corresponding cumulative local histograms based on uniformly spaced sample points $P_n \subset C$ converge:*

$$\lambda_n(r, z) = \lambda_{P_n}(r, z) \longrightarrow h_C(r, z) \quad \text{as } n \rightarrow \infty. \quad (3.1)$$

²Guggenheimer, [10], only requires monotonicity, allowing spiral arcs to contain circular subarcs, which we exclude. Our subsequent definition of regularity includes curves containing finitely many circular arcs and straight line segments.

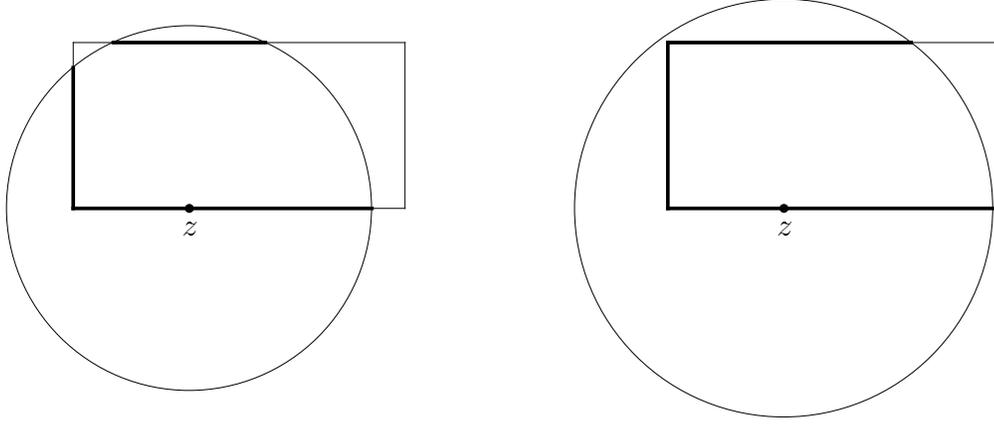


Figure 2: Intersections of a Rectangle and a Disk.

Proof. We will prove convergence by establishing the bound

$$|h_C(r, z) - \lambda_n(r, z)| \leq \frac{m_C \Delta l}{L}, \quad (3.2)$$

where m_C is the circular index of C .

By assumption, since $z \in C$, the intersection $C \cap B_r(z) = S_1 \cup \dots \cup S_k$ consists of k connected segments whose endpoints lie on the bounding circle $S_r(z)$, where $1 \leq k \leq m_C$. Since the sample points are uniformly spaced by $\Delta l = L/n$, the number of sample points n_i contained in an individual segment S_i can be bounded by

$$(n_i - 1) \Delta l \leq l(S_i) < (n_i + 1) \Delta l.$$

Summing over all segments, and noting that

$$\sum_{i=1}^k n_i = \#(P_n \cap B_r(z)) = n \lambda_n(r, z), \quad \sum_{i=1}^k l(S_i) = l(C \cap B_r(z)) = L h_C(r, z),$$

we deduce that

$$L \lambda_n(r, z) - k \Delta l \leq L h_C(r, z) < L \lambda_n(r, z) + k \Delta l,$$

from which (3.2) follows. *Q.E.D.*

Example 9. Let C be a circle of radius 1. A set of n evenly spaced sample points $P_n \subset C$ form the corners of a regular n -gon. Using the identification $\mathbb{R}^2 \simeq \mathbb{C}$, the cumulative histogram of P_n is given by

$$\lambda_n(r, z) = \frac{1}{n} \# \{ j \mid 1 \leq j \leq n, |e^{2\pi i j/n} - z| < r \}.$$

On the other hand, the local histogram function (2.8) for a circle is easily found to have the explicit form

$$h_C(r, z) = \frac{1}{\pi} \cos^{-1} \left(1 - \frac{1}{2} r^2 \right), \quad (3.3)$$

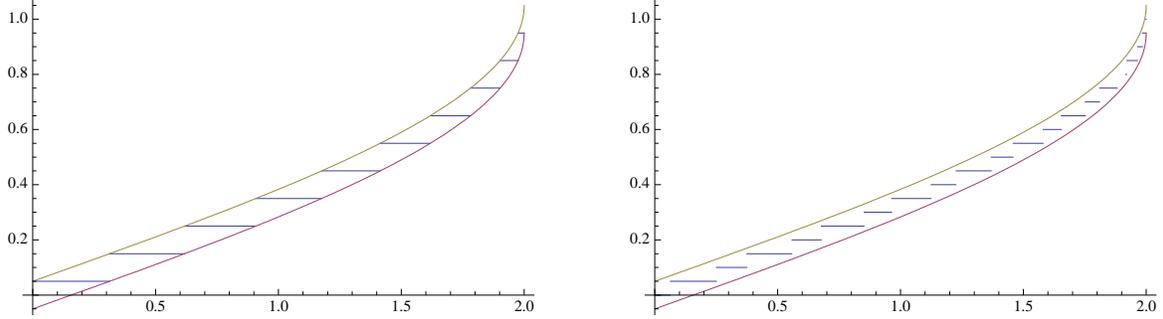


Figure 3: Local Histogram Functions for a Circle.

which, by symmetry, is independent of the point $z \in C$.

In Figure 3, we plot the discrete cumulative histogram $\lambda_n(r, z)$ for $n = 20$, along with the bounds $h_C(r, z) \pm \Delta l / (2\pi)$ coming from (3.2), which reflect the fact that a circle has circular index $m_C = 1$. In the first plot, the center z coincides with a data point, while the second takes z to be a distance .01 away, as measured along the circle. Observe that the discrete histogram stays within the indicated bounds at all radii, in accordance with our result.

4 Global Histogram Functions.

We now turn our attention to the convergence of the global histograms. Again, we work under the preceding regularity assumptions, and continue to focus our attention on the case of uniformly spaced sample points $P_n \subset C$.

First, we observe that the local histogram function $h(s) = h_C(r, z(s))$ is piecewise continuous as a function of s . Indeed, $h(s)$ is continuous unless the circle of radius r centered at $z(s)$ contains one or more circular arcs that belong to C , in which case $h(s)$ has a jump discontinuity whose magnitude is the sum of the lengths of such arcs. By our assumption of regularity, C contains only finitely many circular arcs, and so $h(s)$ can have only finitely many jump discontinuities. On the other hand, regularity implies that the global histogram function is everywhere continuous.

Therefore, the global histogram integral (2.9) can be approximated by a Riemann sum³ based on the evenly spaced data points:

$$H_C(r) = \frac{1}{L} \int_C h_C(r, z(s)) ds \approx \frac{1}{L} \sum_{z \in P_n} h_C(r, z) \Delta l. \quad (4.1)$$

Since C has finite length, $\Delta l = L/n \rightarrow 0$ as $n \rightarrow \infty$, and so the Riemann sums converge. On the other hand, (3.1) implies that the local histogram function can be approximated by the (rescaled) cumulative point histogram $\lambda_n(r, z)$, and hence we should be able to approximate

³When C is not closed, the right hand side is technically not a Riemann sum, since it includes contributions from both endpoints. However, this does not affect the convergence. An alternative approach is to assign the endpoints a weighting of $1/2$ when forming the histogram functions.

the Riemann sum in turn by

$$\frac{1}{L} \sum_{z \in P_n} \lambda_n(r, z) \Delta l = \frac{1}{n} \sum_{z \in P_n} \lambda_n(r, z) = \Lambda_n(r), \quad (4.2)$$

using the first equality of (2.6). This will imply the global convergence result (2.9).

To bound the approximation error, we proceed as follows. Recall, [2], that the *total variation* of a function $f(x)$ on the interval $[a, b]$ is defined as

$$\mathcal{V}_a^b[f] = \sup \sum_{i=1}^{n-1} |f(x_{i+1}) - f(x_i)|, \quad (4.3)$$

where the supremum is taken over all subdivisions $a \leq x_1 < x_2 < \dots < x_n \leq b$. If $f(x)$ is monotone, then $\mathcal{V}_a^b[f] = |f(b) - f(a)|$. If $f(x)$ is continuous and piecewise continuously differentiable, then $\mathcal{V}_a^b[f] = \int_a^b |f'(x)| dx$ is the integral of the absolute value of its derivative. A function is said to have *bounded variation* if $\mathcal{V}_a^b[f] < \infty$. In particular, all piecewise C^1 functions are of bounded variation, as are all functions that are monotone on a finite number of subintervals. Indeed, any function of bounded variation is the difference of two monotone functions. The *cumulative property* states that if $a \leq b \leq c$, then $\mathcal{V}_a^c[f] = \mathcal{V}_a^b[f] + \mathcal{V}_b^c[f]$.

Curiously, we were unable to find an exact statement and proof in the literature of the following simple result concerning the numerical approximation of the integral of a function of bounded variation.

Theorem 10. *If $f(x)$ is of bounded variation on $[a, b]$, then the difference between the integral $\int_a^b f(x) dx$ and any approximating Riemann sum is bounded in absolute value by $\mathcal{V}_a^b[f] \Delta x$, where Δx is the maximal subinterval length appearing in the Riemann sum.*

Proof: Let $a \leq x_1 < x_2 < \dots < x_n \leq b$ be a subdivision, and

$$R = \sum_{i=1}^{n-1} f(y_i) \Delta x_i$$

a corresponding Riemann sum, where $x_i \leq y_i \leq x_{i+1}$ and $\Delta x_i = x_{i+1} - x_i$. On the i -th subinterval $[x_i, x_{i+1}]$, let x_i^+ and x_i^- be the points at which $f(x)$ achieves, respectively, its maximum and minimum values. Then

$$f(x_i^-) \Delta x_i \leq \int_{x_i}^{x_{i+1}} f(x) dx \leq f(x_i^+) \Delta x_i,$$

and hence

$$\left| \int_{x_i}^{x_{i+1}} f(x) dx - f(y_i) \Delta x_i \right| \leq |f(x_i^*) - f(y_i)| \Delta x_i \leq \mathcal{V}_{x_i}^{x_{i+1}}[f] \Delta x_i,$$

where $x_i^* = x_i^+$ or x_i^- is the point that maximizes $|f(x) - f(y_i)|$ on the interval. Therefore, by the cumulative property of the total variation,

$$\left| \int_a^b f(x) dx - \sum_{i=1}^{n-1} f(y_i) \Delta x_i \right| \leq \sum_{i=1}^{n-1} \mathcal{V}_{x_i}^{x_{i+1}}[f] \Delta x_i \leq \mathcal{V}_a^b[f] \Delta x,$$

completing the proof. Q.E.D.

Given a C^1 curve C , the *parallel curve* $\Pi_r(C)$ at *distance* $r > 0$ is defined as the locus of points that are a distance r away from C , as measured along the normal direction, [10]. Since there are two normal directions at each point, $\Pi_r(C)$ consists of two connected, not necessarily simple, curves. If C is only piecewise C^1 , then $\Pi_r(C)$ is defined as the union of parallel curves at distance r to each C^1 segment combined with the circular arcs of radius r centered at each corner that connect the endpoints of the parallel curve segments.

Definition 11. A regular curve C is called *fully regular* if, for each $r > 0$, the intersection $C \cap \Pi_r(C)$ has finitely many connected components.

Not all regular curves are fully regular. For example, it is possible to slightly deform part of a circle of radius $\frac{1}{2}r$ to produce a smooth convex curve that intersects its parallel curve of distance r infinitely often. On the other hand, most regular curves, including all polygons and biarcs, are fully regular.

Theorem 12. *If C is a fully regular curve, then its local histogram function $h(s) = h_C(r, z(s))$ is piecewise continuous and of bounded variation on $[0, L]$.*

This result is a consequence of the following formula for the derivative of the local histogram function.

Proposition 13. *Let $z = z(s) \in C$, and assume that $h(s) = h_C(r, z(s))$ is continuous at s . Let $y_1 = z(s_1), \dots, y_k = z(s_k)$ denote the points of intersection of C with the circle $S_r(z)$ of radius r centered at z . For each $j = 1, \dots, k$, let l_j denote the line through z and y_j . Let θ_j denote the angle between l_j and the tangent vector $t(s) = z'(s)$ to C at z in the direction of increasing s . Let φ_j denote the angle between the line l_j and the tangent vector (not necessarily in the direction of increasing s) to C at y_j that points outside the circle. Then*

$$\frac{dh}{ds} = \frac{d}{ds} h_C(r, z(s)) = \frac{1}{L} \sum_{j=1}^k \frac{\cos \theta_j}{\cos \varphi_j}. \quad (4.4)$$

Proof: Referring to Figure 4, note first that $\cos \varphi_j = 0$ if and only if the curve C is tangent to the circle $S_r(z)$ at the intersection point y_j . In this case, the line l_j is normal to C at y_j , and hence $z \in \Pi_r(C)$, the parallel curve at distance r . We exclude these configurations from the remainder of the argument.

The calculation is simplified by identifying $\mathbb{R}^2 \simeq \mathbb{C}$, so $z = x + iy$. We place $z = z(s) = 0$ at the origin, with the tangent vector to C at z in the direction of the positive x axis. Let $0 < \Delta s \ll 1$. To first order, $z(s + \Delta s) = \Delta s + \dots$. Moreover, as Δs varies, the intersection point $y_j = z(s_j) = r_j e^{i\theta_j}$ moves to a point

$$w_j = z(s_j + \Delta s_j) = y_j + e^{i(\theta_j - \varphi_j)} \Delta s_j + \dots = e^{i\theta_j} (r_j + e^{-i\varphi_j} \Delta s_j) + \dots$$

on the circle $S_r(z(s + \Delta s))$, satisfying the algebraic equation

$$\begin{aligned} r^2 &= |w_j - z(s + \Delta s)|^2 = |e^{i\theta_j} (r_j + e^{-i\varphi_j} \Delta s_j) - \Delta s|^2 + \dots \\ &= |r_j + e^{-i\varphi_j} \Delta s_j - e^{-i\theta_j} \Delta s|^2 + \dots \\ &= (r_j + \Delta s_j \cos \varphi_j - \Delta s \cos \theta_j)^2 + (-\Delta s_j \sin \varphi_j - \Delta s \sin \theta_j)^2 + \dots \\ &= r^2 - 2r \Delta s \cos \theta_j + 2r \Delta s_j \cos \varphi_j + \dots \end{aligned}$$

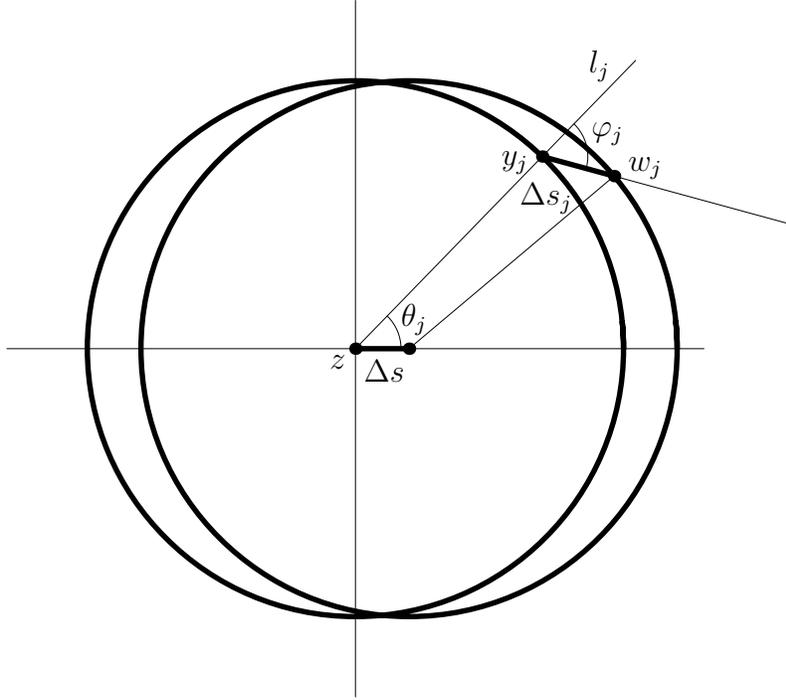


Figure 4: Calculating the Derivative of the Distance Function.

Solving, to leading order,

$$\frac{\Delta s_j}{\Delta s} = \frac{\cos \theta_j}{\cos \varphi_j} + \dots$$

Taking the limit as $\Delta s \rightarrow 0$, summing over the points y_j in the intersection, and then dividing by L completes the proof. *Q.E.D.*

Proof of Theorem 12: As a consequence of (4.4), the derivative $h'(s)$ is defined and nonzero, unless either $z = z(s)$ is a corner of C , or $\cos \varphi_j = 0$, or $\cos \theta_j = 0$. As noted above, the second possibility means that the point $z \in \Pi_r(C)$. The third possibility means that the point $y_j \in C \cap S_r(z)$ lies a distance r away from $z \in C$, as measured along the normal direction at z , and hence $y_j \in \Pi_r(C)$. Thus, by our fully regular hypothesis, $h'(s)$ is defined and nonzero except on a finite number of points and/or connected subintervals. We conclude that $h(s)$ is piecewise continuous, piecewise monotone, and hence of bounded variation. *Q.E.D.*

Let $V = \mathcal{V}_0^L[h]$ denote the total variation of $h(s) = h_C(r, z(s))$. Theorem 10 implies that the error in the Riemann sum approximation (4.1) can be bounded by

$$\left| \frac{1}{n} \sum_{z \in P_n} h_C(r, z) - \frac{1}{L} \int_C h_C(r, z(s)) ds \right| \leq V \Delta l.$$

On the other hand, (3.2) and the triangle inequality imply that

$$\left| \frac{1}{n} \sum_{z \in P_n} \lambda_n(r, z) - \frac{1}{n} \sum_{z \in P_n} h_C(r, z) \right| \leq \frac{m_C \Delta l}{L}. \quad (4.5)$$

Combining these two estimates, we find that the difference between the continuous and discrete global histogram functions can be bounded by

$$|\Lambda_n(r) - H_C(r)| = \left| \frac{1}{n} \sum_{z \in P_n} \lambda_n(r, z) - \frac{1}{L} \int_C h_C(r, z(s)) ds \right| \leq \frac{m_C + V}{L} \Delta l. \quad (4.6)$$

Thus, under our hypotheses, the convergence to the global histogram function is first order in the interpoint spacing Δl .

Example 14. Let C be a unit square, so that $L = l(C) = 4$. Measuring the arc length s along the square starting at a corner, the local histogram function $h_r(s) = h_C(r, z(s))$ can be explicitly constructed using elementary geometry, distinguishing several different configurations. For $0 \leq s \leq \frac{1}{2}$,

$$h_r(s) = \begin{cases} \frac{1}{2}r, & 0 \leq r \leq s, \\ \frac{1}{4}s + \frac{1}{4}r + \frac{1}{4}\sqrt{r^2 - s^2}, & s \leq r \leq 1 - s, \\ \frac{1}{4} + \frac{1}{4}\sqrt{r^2 - s^2} + \frac{1}{4}\sqrt{r^2 - (1 - s)^2}, & 1 - s \leq r \leq 1, \\ \frac{1}{4} + \frac{1}{2}\sqrt{r^2 - 1} + \frac{1}{4}\sqrt{r^2 - s^2} + \frac{1}{4}\sqrt{r^2 - (1 - s)^2}, & 1 \leq r \leq \sqrt{1 + s^2}, \\ \frac{1}{4}s + \frac{1}{2} + \frac{1}{4}\sqrt{r^2 - 1} + \frac{1}{4}\sqrt{r^2 - (1 - s)^2}, & \sqrt{1 + s^2} \leq r \leq \sqrt{1 + (1 - s)^2}, \\ 1, & \sqrt{1 + (1 - s)^2} \leq r, \end{cases} \quad (4.7)$$

while other values follow from the fact that $h_r(s)$ is both 1-periodic and even:

$$h_r(1 - s) = h_r(s) = h_r(1 + s).$$

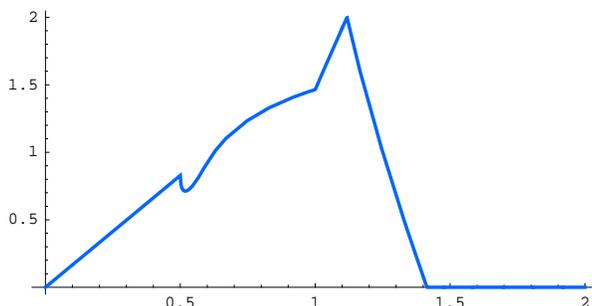


Figure 5: Total Variation of the Local Histogram Function of a Square.

Integration around the square with respect to arc length produces the global histogram function

$$H_C(r) = \begin{cases} \frac{1}{2}r + \left(\frac{1}{8}\pi - \frac{1}{4}\right)r^2, & r < 1, \\ \frac{1}{2} - \frac{1}{4}r^2 + \sqrt{r^2 - 1} + \frac{1}{4}r^2 \left(\sin^{-1} \frac{1}{r} - \cos^{-1} \frac{1}{r} \right), & 1 \leq r < \sqrt{2}, \\ 1, & r \geq \sqrt{2}. \end{cases} \quad (4.8)$$

It is interesting that, while the local histogram function has six intervals with different analytical formulas, the global function has only three.

The total variation of $h_r(s)$ is obtained by integrating the absolute value of its derivative. In Figure 5, we plot the variation as a function of r . (The little dip after $r = .5$ is genuine, and the reader may enjoy seeking a geometrical explanation.) The maximal variation, $V_{max} = 2$, occurs at $r = \sqrt{5}/2$, where $h_{\sqrt{5}/2}(s)$ varies monotonically from $\frac{3}{4}$ to 1 and then back four times as we go around the square. On the other hand, according to Example 7, $m_C = 2$. Thus, an overall bound (4.6) for the error in $\lambda_P(r, z)$, valid for all r , is $(m_C + V_{max}) \Delta l / L = \frac{5}{2} \Delta l$.

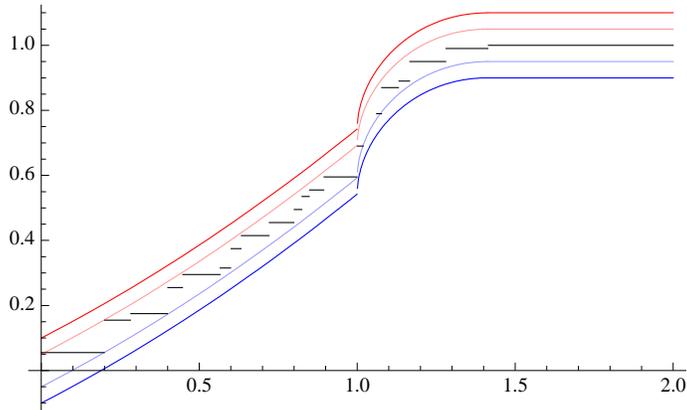


Figure 6: Global Histogram Bounds for a Square.

Figure 6 plots the global cumulative histograms of a square based on $n = 20$ evenly spaced points, along with the bounds $\frac{1}{4} \Delta l$ and $\frac{1}{2} \Delta l$. Observe that the discrete histogram stays within $\frac{1}{4} \Delta l$ of the curve histogram, a tighter bound than we are able to derive analytically. Interestingly, a similarly tight bound appears to hold in all the examples we have looked at so far.

5 Random Point Distributions.

We have thus far proved, under suitable regularity hypotheses, convergence of both the local and global cumulative histograms constructed from uniformly spaced sample points along the curve. However, in practice, it may be difficult to ensure precise uniform spacing of the sample points. For example, if C is an ellipse, then this would require evaluating n elliptic integrals. Hence, for practical shape analysis, we need to examine more general methods of histogram creation. In this section, we analyze the case of sample points $P_n = \{z_1, \dots, z_n\} \subset C$ that are randomly chosen with respect to the uniform arc length distribution.

In this case, we view the cumulative local histogram $\lambda_n(r, z)$ as a random variable representing the fraction of the points z_i that lie within a circle of radius r centered at the point z . Indeed, we can write

$$\lambda_n(r, z) = \frac{1}{n} \sum_{i=1}^n \sigma_i(r, z),$$

where each $\sigma_i(r, z)$ is a random variable that is 1 if $d(z_i, z) \leq r$ and 0 otherwise. Then, for each $i = 1, \dots, n$,

$$E[\sigma_i(r, z)] = \text{Prob}\{d(z_i, z) \leq r\} = \frac{l(C \cap B_r(z))}{L} = h_C(r, z),$$

and hence

$$E[\lambda_n(r, z)] = \frac{1}{n} \sum_{i=1}^n E[\sigma_i(r, z)] = h_C(r, z). \quad (5.1)$$

Similarly, to construct a statistical variable whose expectation approximates the global histogram function $H_C(r)$, consider

$$\Lambda_n(r) = \frac{1}{n^2} \sum_{i=1}^n \#(P \cap B_r(z_i)) = \frac{1}{n} + \frac{1}{n^2} \sum_i \sum_{j \neq i} \sigma_{i,j}(r),$$

where $\sigma_{i,j}(r)$ is a random variable that is 1 if $d(z_i, z) \leq r$ and 0 otherwise. As above, its expected value is

$$\begin{aligned} E[\sigma_{i,j}(r)] &= \text{Prob}\{d(z_i, z_j) \leq r\} \\ &= \frac{1}{L} \int_0^L \text{Prob}\{d(z_i, z(s)) \leq r\} ds = \frac{1}{L} \int_0^L h_C(r, z(s)) ds = H_C(r). \end{aligned}$$

Therefore,

$$E[\Lambda_n(r)] = \frac{1}{n} + \frac{1}{n^2} \sum_i \sum_{j \neq i} E[\sigma_{i,j}(r)] = \frac{1}{n} + \frac{n-1}{n} H_C(r). \quad (5.2)$$

Thus, the expected value of $\Lambda_n(r)$ tends to the global histogram function $H_C(r)$ as $n \rightarrow \infty$.

Next we compute the variances of the local and global histogram functions. First,

$$\text{Var}[\lambda_n(r, z)] = E[\lambda_n(r, z)^2] - E[\lambda_n(r, z)]^2 = \frac{1}{n^2} \sum_{i,j} E_{i,j},$$

where

$$E_{i,j} = E[\sigma_i(r, z) \sigma_j(r, z)] - E[\sigma_i(r, z)] E[\sigma_j(r, z)].$$

On the other hand, if $i \neq j$, then $\sigma_i(r, z)$ and $\sigma_j(r, z)$ are independent random variables, so the expected value of their product is the product of their expected values, and hence $E_{i,j} = 0$. If $i = j$, then

$$E_{i,i} = \text{Var}[\sigma_i(r, z)^2] = E[\sigma_i(r, z)^2] - E[\sigma_i(r, z)]^2 = h_C(r, z) - h_C(r, z)^2,$$

since $\sigma_i(r, z)$ represents an indicator function. We conclude that variance of the local histogram is

$$\text{Var}[\lambda_n(r, z)] = \frac{h_C(r, z) - h_C(r, z)^2}{n}. \quad (5.3)$$

Similarly, to compute the global histogram variance,

$$\text{Var}[\Lambda_n(r)] = E[\Lambda_n(r)^2] - E[\Lambda_n(r)]^2 = \frac{1}{n^4} \left(\sum_{\substack{i,i',j,j' \\ \text{all distinct}}} E_{i,i',j,j'} + \sum_{\substack{i,i',j \neq i',j' \neq i' \\ \text{not all distinct}}} E_{i,i',i,j'} \right),$$

where

$$E_{i,i',j,j'} = E[\sigma_{i,j}(r) \sigma_{i',j'}(r)] - E[\sigma_{i,j}(r)] E[\sigma_{i',j'}(r)].$$

As above, the terms in the first summation are all 0, whereas those in the second are bounded. As there are $O(n^3)$ of the latter, we conclude that

$$\text{Var}[\Lambda_n(r)] = O(n^{-1}). \tag{5.4}$$

Thus, $\Lambda_n(r)$ converges to $H_C(r)$ in the sense that, for any given value of r , the probability of $\Lambda_n(r)$ lying in any interval around $H_C(r)$ approaches 1 as $n \rightarrow \infty$.

Although this probabilistic calculation yields a less stringent error bound in (5.4), it nevertheless provides a practical method for calculating $H_C(r)$ in cases when the extraction of evenly spaced sample points is problematic. However, we still face several important issues. We assumed that P was a set of randomly selected points on the curve, but we did not define how such points should be chosen. Since the calculation of $\Lambda_n(r)$ essentially reduces to evaluating (2.9), we require the points to be sampled uniformly with respect to arc length. One way to accomplish this in practice is to select points using any convenient parameterization of the curve, but weighting the sample in favor of the points with higher curvature using, for instance, rejection sampling, [8].

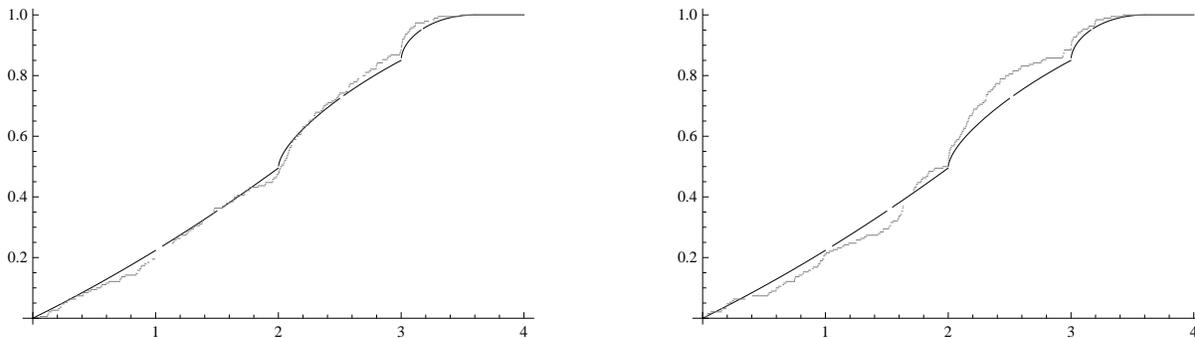


Figure 7: Comparison of approximate histograms of a rectangle.

Example 15. Let C be a 2×3 rectangle. In Figure 7, we graph its global curve histogram function $H_C(r)$ in black and the approximate histograms $\Lambda_n(r)$, based on $n = 20$ sample points, in gray. The first plot is for evenly distributed points, in which the approximation remains within Δl of the continuous histogram function, while the second plot is for randomly generated points, in which the approximation stays within $2 \Delta l$. Thus, both methods work as advertised.

6 Pixelization.

One of our motivating goals has been potential applications of distance histograms to the recognition and classification of objects in digitized images. In such situations, it is often more convenient to use the locations of the pixels that the curve passes through to represent the sample points on the shape boundary. This results in two potential errors in the ensuing

calculations. First, the normalization can change, because different pixels will not contain the same length of curve. As long as the *pixel width* w is small, namely $w < 1/\kappa_{max}$, where κ_{max} denotes the maximum curvature on that part of the curve, and the curve intersects the pixel in only one connected segment, then the longest possible curve length is that of a quarter circle, namely, $\frac{1}{2}\pi w$. The shortest possible length depends on the method used to select pixels. In the extreme case, we include all pixels containing an arbitrarily short curve segment — even if only a single point — and the potential variance in normalization is $\frac{1}{2}\pi w$. However, if the pixelization is very fine then the number of pixels selected to represent the curve will be small compared to the total number of image pixels, and the relative pixel normalizations should not have a noticeable effect.

Second, by using the coordinates of the center of a pixel containing a point, we introduce a possible offset to the location of the sample point of up to a distance

$$\delta = \frac{w}{\sqrt{2}}.$$

Thus, for every distance we calculate, we introduce an additional uncertainty of $2\delta = \sqrt{2}w$. If we approximate $h_C(r, z) \approx \lambda_n(r, z)$, then, even with no statistical error, we still have an uncertainty range of

$$\lambda_n(r - 2\delta, z) < h_C(r, z) < \lambda_n(r + 2\delta, z).$$

Since both h and λ_n are monotonically increasing as functions of r , the resulting contribution to the error can thus be bounded by

$$|\lambda_n(r + 2\delta, z) - \lambda_n(r - 2\delta, z)| \approx 4\delta \left| \frac{\partial \lambda_n}{\partial r}(r, z) \right| = 2\sqrt{2}w \left| \frac{\partial \lambda_n}{\partial r}(r, z) \right| \quad \text{for } \delta \ll 1.$$

Thus, the error in the pixelized approximation will be proportional to the pixel width. This error will not necessarily go to zero as $n \rightarrow \infty$. However, if we assume that the length of the curve is L finite, then, according to Theorem 18 below, we can bound the number of pixels that the curve passes through. Therefore, assuming $h(r)$ is piecewise C^1 and its derivative $\partial h/\partial r$ is bounded in absolute value, then, as $n \rightarrow \infty$, the pixel width $w \rightarrow 0$, and so the local histogram function $\lambda_n(r, z)$ still converges for pixelized images. Similarly, since the global histogram function $H_C(r)$ is just the average of $h_C(r, z)$ along the curve, its pixelized approximation will also converge.

This property is especially important in terms of resistance to noise. For curvature-based techniques of object recognition, as in [7], noise causes large fluctuations in the local curvature-based invariants, which interferes with their use in shape recognition. The distance histogram is largely unaffected by such local noise, and, as we increase the number of points, the effect disappears.

Example 16. Let C be a 2×3 rectangle, set at angle 15° with respect to the coordinate axes. (Results for other angles, even when aligned with the axes, are very similar.) In Figure 7, we compare the graphs of the global histogram $H_C(r)$ (in black) with the discrete approximations $\Lambda_n(r)$ for $n = 20$ points, using evenly distributed sample points (green), random points (blue), and pixelized points (orange). The evenly distributed case provides the closest approximation to the curve, and remains within Δl of the curve histogram function. Both the randomly generated points and the pixelized randomly generated points stay within $2\Delta l$, and so all three methods work as advertised.

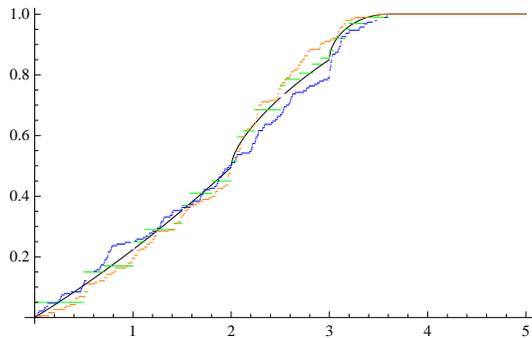


Figure 8: Comparison of Approximate Histograms of a Rectangle.

Example 17. Let us also recheck that the global distance histograms are indeed invariant under Euclidean transformations. Since they preserve distances, the net effect of a rigid motion is to resample the curve. In the random and pixelized cases, this effect is eliminated by selecting the points randomly, but in the uniformly distributed case, selecting different points could affect $\lambda_n(r)$.

In Figure 9, we plot the approximate histograms for a square calculated using different offsets for the points. The first graph uses $n = 20$ sample points, while the second has $n = 50$. The individual graphs are clearly different, but as we increase the number of points, they are converging to the same global histogram function (4.8), in accordance with our general results.

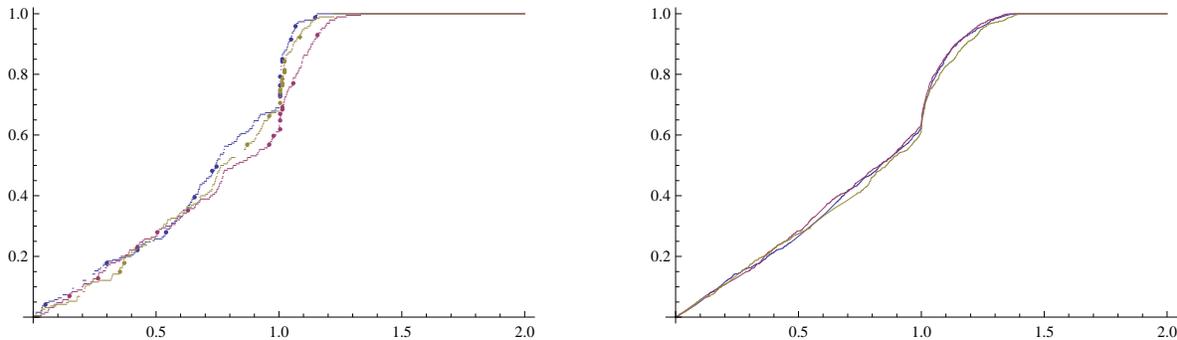


Figure 9: Approximate Histograms of a Square.

The one remaining point is to show how to bound the number of pixels that the curve passes through. Surprisingly, we were unable to find a statement of such a result in the literature, and therefore provide a direct proof. By convention, a pixel is a *closed* square box, and so two neighboring pixels will have a nonempty intersection along a common bounding line segment. As before, a pixel is counted even if the curve only touches one point.

Theorem 18. *Let C be a rectifiable closed curve of length $L = l(C) \geq 3\sqrt{2}w$. Then C passes through at most $n = 3L/(\sqrt{2}w)$ square pixels of pixel width w .*

Proof: By a simple rescaling and translation, we can assume without loss of generality that the pixels have unit width, $w = 1$, with corners on the integer lattice $\mathbb{Z}^2 \subset \mathbb{R}^2$.

We first remark that the bound given in the statement of the theorem is sharp. It is achieved when, for instance, C is a rectangle whose sides are at a 45° angle with the horizontal, and whose corners lie on the pixel lattice. For example, the rectangle in Figure 10 has length $L = 14\sqrt{2}$, and goes through $n = 42 = 3L/\sqrt{2}$ pixels. Even if one disallows pixels that meet the curve in a single point, one can slightly perturb such rectangles to obtain curves that come arbitrarily close to the pixel bound.

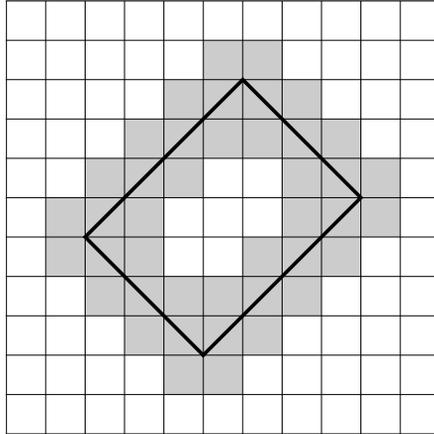


Figure 10: Pixels Containing a Rectangle.

Secondly, if C passes through ≤ 9 pixels, then the bound is trivially satisfied due to the restriction on its length. Indeed, by our convention, an individual pixel boundary, which forms a unit square of length 4, passes through all $9 > 12/\sqrt{2}$ pixels, and hence the stated bound is not valid for curves of shorter length.

Keep in mind that we are dealing with a closed curve. If C is entirely contained within a *horizontal two pixel strip* $\widehat{H}_j = \{(x, y) \mid j - 1 < y < j + 1\}$ for some $j \in \mathbb{Z}$, then it passes through at most $L + 4$ pixels — the maximum being approached by curves that go arbitrarily closely around the horizontal line segment $\{(x, j) \mid i \leq x \leq i + L/2\}$ for L an even integer. Since $L > 4$, the number of pixels the segment passes through is bounded by $L + 4 < 3L/\sqrt{2}$, proving the theorem in this simple case. Thus, for the remainder of the proof, we may assume C passes through at least 10 pixels, and is not contained in a two pixel strip.

Our strategy is to break C up into a finite collection of nonoverlapping segments, prove an appropriate bound on the number of pixels each segment passes through, and then successively merge adjacent segments while maintaining the bound. To each such segment $S \subset C$, let $p(S)$ denote the collection of pixels that it passes through. We split $p(S) = q(S) \cup e(S)$ into two disjoint subsets, where $e(S)$ consists of some of the pixels that contain the endpoints of S , satisfying

$$2 \leq \#e(S) \leq 4, \quad \text{while} \quad \#q(S) \leq \frac{3l(S)}{\sqrt{2}}. \quad (6.1)$$

The initial segments will all be of the form $S = E_0 \cup F \cup E_1$, where F is curve segment that, apart from its endpoints, lies strictly within a *horizontal one pixel strip* $H_j = \{(x, y) \mid j < y < j + 1\}$ for some $j \in \mathbb{Z}$, and such that one of its endpoints lies on one bounding horizontal line $y = j$, while the other endpoint lies on the other line $y = j + 1$,

while E_0 and E_1 are closed straight line segments contained, respectively, in the horizontal lines $y = j$ and $y = j + 1$, each sharing a common endpoint with F . We allow E_0 and/or E_1 to consist of just a single endpoint of S . We define $e(S)$ to be the set of pixels that contain the endpoints of S and lie outside the strip H_j , while $q(S) = p(S) \setminus e(S)$. Note that each endpoint contributes either two or one pixels to $e(S)$, depending on whether it is a lattice point or not, and so $2 \leq \#e(S) \leq 4$. A representative example can be seen in Figure 11: F is the curved segment, E_0 is the right hand endpoint, while E_1 is the line segment on the left. The 14 pixels in $q(S)$ are lightly shaded, while the 3 in $e(S)$ are darker.

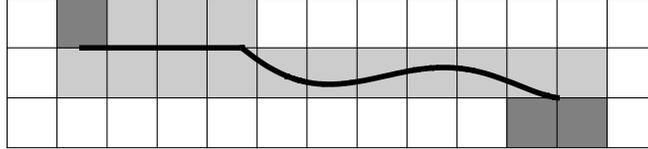


Figure 11: Segment Contained in a Horizontal Strip.

In all situations $l(F) \geq 1$. If F passes through $k \geq 2$ pixels in the strip, then $l(F) \geq \sqrt{1 + (k - 2)^2}$, with the minimal length corresponding to a straight line whose ends are lattice points on opposite sides of the strip. Therefore,

$$k \leq \frac{k l(F)}{\sqrt{1 + (k - 2)^2}} \leq \frac{3 l(F)}{\sqrt{2}},$$

where the maximum is attained when $k = 3$. (Keep in mind that $k \geq 2$ is an integer. In Figure 11, $k = 8$.) The bound trivially holds when $k = 1$. Moreover, it's not hard to see that, by construction, the total number of pixels in $q(C)$ is bounded by

$$\#q(S) \leq k + 2l(E_0) + 2l(E_1) \leq \frac{3}{\sqrt{2}} [l(E_0) + l(E_1) + l(F)] = \frac{3l(S)}{\sqrt{2}}.$$

There are several cases, depending on whether the endpoints of S and F lie on the lattice or not; details are left to the reader.

Now, let's see what happens when we merge two adjacent segments that individually satisfy (6.1), forming a larger segment $S = S_1 \cup S_2 \subsetneq C$. The common endpoint of S_1 and S_2 (which is no longer an endpoint of S) is in either one or two pixels that belong to both $e(S_1)$ and $e(S_2)$, the number depending on whether or not it is a lattice point. If S_1 and S_2 are associated with different horizontal strips, e.g., H_j and H_{j+1} , then these one or two pixels already appear in $q(S) = q(S_1) \cup q(S_2)$, and so the bound (6.1) is an immediate consequence of the bounds for S_1 and S_2 . On the other hand, if S_1 and S_2 are associated with the same horizontal strip, then the one or two pixels in both $e(S_1)$ and $e(S_2)$ that contain the common endpoint do not appear in $q(S_1)$ or $q(S_2)$. However, there are an equal number of pixels that are in both $q(S_1)$ and $q(S_2)$, and so $\#(e(S_1) \cap e(S_2)) = \#(q(S_1) \cap q(S_2))$. Thus, setting $q(S) = q(S_1) \cup q(S_2) \cup (e(S_1) \cap e(S_2))$, we conclude that $\#q(S) \leq \#q(S_1) + \#q(S_2)$, and hence (6.1) also holds for such combined segments.

We continue to merge segments in this fashion. The only detail is in the final merger to form $C = \widehat{S} \cup \widetilde{S}$, when the segments \widehat{S} and \widetilde{S} have two common endpoints. If the

endpoints are more than a single pixel apart, then the merger process can be completed without interference, and the result follows. (This is where the argument breaks down for a pixel boundary square.) Our initial assumptions on C guarantee that we can suitably order the merging process to ensure that this is the case. *Q.E.D.*

Remark: For a nonclosed curve, the same argument implies the number of pixels is bounded by $n = 4 + 3L/(\sqrt{2}w)$. Again, the bound is tight, as can be seen by deleting a suitable part of the rectangle in Figure 10.

7 Histogram–Based Shape Comparison.

In this section, we discuss the question of whether distance histograms can be used, both practically and theoretically, as a means of distinguishing shapes up to rigid motion. We begin with the practical aspects. As we know, if two curves have different global histogram functions, they cannot be rigidly equivalent. For curves arising from digital images, we will approximate the global histogram function by its discrete counterpart based on a reasonably dense sampling of the curve. Since the error in the approximations is proportional to $\Delta l = L/n$, we will calculate the average difference between two histogram plots, normalized with respect to Δl . Our working hypothesis is that differences less than 1 represent histogram approximations that cannot be distinguished.

Tables 1 and 2 show these values for a few elementary shapes. We use random point distributions⁴ to illustrate that identical parameterizations do not necessarily give identical sample histograms. This is also evident from the fact that the matrix is not symmetric — different random sample points were chosen for each trial. However, symmetrically placed entries generally correlate highly, indicating that the comparison is working as intended.

Table 1 is based on discretizing using only $n = 20$ points. As we see, this is too small a sample set to be able to unambiguously distinguish the shapes. Indeed, the 2×3 rectangle and the star appear more similar to each other than they are to a second randomized version of themselves. On the other hand, for the star and the circle, the value of 5.39 is reasonably strong evidence that they are not rigidly equivalent.

Shape	(a)	(b)	(c)	(d)	(e)	(f)
(a) triangle	.35	1.16	1.46	4.20	2.36	3.16
(b) square	1.45	.51	3.63	2.46	1.59	2.89
(c) circle	3.65	4.17	.67	5.87	3.14	5.39
(d) 2×3 rectangle	3.85	1.95	4.82	1.78	1.85	.72
(e) 1×3 rectangle	1.10	1.86	4.02	2.31	1.25	1.93
(f) star	3.90	3.80	5.75	.72	2.55	1.22

Table 1: 20–point Comparison Matrix.

As we increase the number of sample points, the computation time increases (in proportion

⁴More precisely, we first select n uniformly distributed random numbers $s_i \in [0, L]$, $i = 1, \dots, n$, and then take the corresponding n random points $z(s_i) \in C$ based on a given arc length parametrization. In our experiment, the shapes are sufficiently simple that the explicit arc length parametrization is known.

to n^2 for calculating the histograms and n for comparing them), but our ability to differentiate shapes increases as well. In Table 2, based on $n = 500$ sample points, it is now clear that none of the shapes are rigidly equivalent to any of the others. The value of 4 for comparing the 1×3 rectangle to itself is slightly high, but it is still significantly less than any of the values for comparing two different shapes.

Shape	(a)	(b)	(c)	(d)	(e)	(f)
(a) triangle	2.3	20.4	66.9	81.0	28.5	76.8
(b) square	28.2	.5	81.2	73.6	34.8	72.1
(c) circle	66.9	79.6	.5	137.0	89.2	138.0
(d) 2×3 rectangle	85.8	75.9	141.0	2.2	53.4	9.9
(e) 1×3 rectangle	31.8	36.7	83.7	55.7	4.0	46.5
(f) star	81.0	74.3	139.0	9.3	60.5	.9

Table 2: 500-point Comparison Matrix.

Our application of curve histogram functions as a means of classifying shapes up to rigid motion inspires us to ask whether all shapes can be thus distinguished. As we saw, while almost all finite sets of points in Euclidean space can be reconstructed, up to rigid motion, from the distances between them, there are counterexamples, including the kite and trapezoid shown in Figure 1, whose distance histograms are identical. However, the curve histograms $H_C(r)$ based on their outer polygons can easily be distinguished. In Figure 12, we plot the approximate global histograms $\Lambda_n(r)$ based on $n = 20$ uniformly spaced sample points. The kite is shown in blue and the trapezoid is shown in purple.

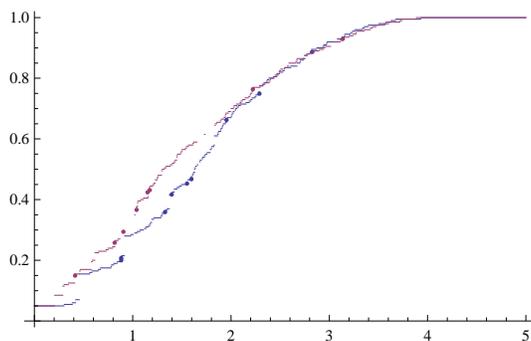


Figure 12: Curve Histograms for the Kite and Trapezoid.

While we have as yet been unable to establish a complete proof, there is a variety of credible evidence in favor of the following:

Conjecture: Two regular plane curves C and \tilde{C} have identical global histogram functions, so $H_C(r) = H_{\tilde{C}}(r)$ for all $r \geq 0$, if and only if they are rigidly equivalent: $C \simeq \tilde{C}$.

One evident proof strategy would be to approximate the histograms by sampling and then apply the convergence result of Theorem 4. If one could prove that the sample points do not, at least when taken sufficiently densely along the curve, lie in the exceptional set of Theorem 2, then our conjecture would follow.

A second strategy is based on our observation that, even when the corners of a polygon lie in the exceptional set, the associated curve histogram still appears to uniquely characterize it. Indeed, if one can prove that the global distance histogram of a simple closed polygon (as opposed to the discrete histogram based on its corners) uniquely characterizes it up to rigid motion, then our conjecture for general curves would follow by suitably approximating them by their interpolating polygons.

To this end, let K be a simple closed polygon of length $L = l(K)$ all of whose angles are *obtuse*, as would be the case with a sufficiently densely sample polygon of a smooth curve. Let l_* be the minimum side length, and d_* be the minimum distance between any two nonadjacent sides. Set $m_* = \min\{l_*, d_*\}$. Then any disk $B_r(z)$ centered at a point $z \in K$ of radius r with $0 < r < \frac{1}{2}m_*$ intersects K in either one or two sides, the latter possibility only occurring when z is within a distance r of the nearest corner. Let z_1, \dots, z_n be the corners of K , and let $\theta_j > \frac{1}{2}\pi$ denote the interior angle at z_j — see Figure 13.

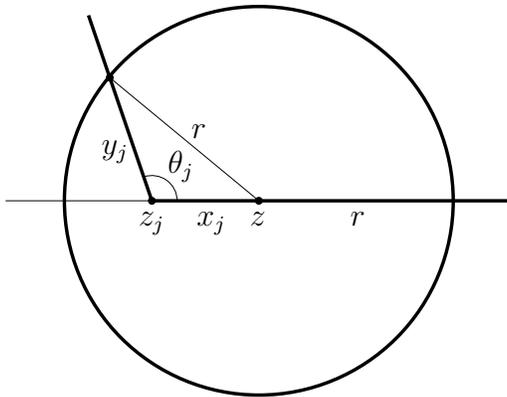


Figure 13: Intersection of a Polygon and a Disk.

Then, for $r > 0$ sufficiently small, and all $z \in K$,

$$L h_K(r, z) = l(K \cap B_r(z)) = \begin{cases} x_j + y_j + r, & x_j = d(z, z_j) < r, \\ 2r, & \text{otherwise,} \end{cases} \quad (7.1)$$

where, by the Law of Cosines, y_j solves the quadratic equation

$$y_j^2 - 2x_j y_j \cos \theta_j + x_j^2 = r^2, \quad \text{with} \quad x_j = d(z, z_j) < r. \quad (7.2)$$

Thus, for small r , the global histogram function (2.9) for such an “obtuse polygon” takes the form

$$H_K(r) = \frac{1}{L} \oint_K h_K(r, z(s)) ds = \frac{2r}{L} - \frac{2nr^2}{L^2} + \frac{2}{L^2} \sum_{j=1}^n \Psi(\theta_j, r), \quad (7.3)$$

where

$$\Psi(\theta_j, r) = \int_0^r [x + y_j(x)] dx, \quad (7.4)$$

with $y_j = y_j(x)$ for $x = x_j$ implicitly defined by (7.2). (There is, in fact, an explicit, but not very enlightening, formula for this integral in terms of elementary functions.)

Observe that (7.3) is a *symmetric function* of the polygonal angles $\theta_1, \dots, \theta_n$, i.e., it is not affected by permutations thereof. Moreover, for distinct angles, the integrals $\Psi(\theta_j, r)$ can be shown to be linearly independent functions of r . This implies that one can recover the set of polygonal angles $\{\theta_1, \dots, \theta_n\}$ from knowledge of the global histogram function $H_K(r)$ for small r . In other words, the polygon's global histogram function does determine its angles up to a permutation.

The strategy for continuing a possible proof would be to gradually increase the size of r . Since, for small r , the histogram function has prescribed the angles, its form is fixed for all $r \leq \frac{1}{2}m_*$. For $r > \frac{1}{2}m_*$, the functional form will change, and this will serve to characterize m_* , the minimal side length or distance between non-adjacent sides. Proceeding in this fashion, as r gradually increases, more and more sides of the polygon can be covered by a disk of that radius, providing more and more geometric information about the polygon from the resulting histogram. This points the way to a proof of our polygonal histogram conjecture, and hence the full curve conjecture. However, the details in such a proof strategy appear to be quite intricate.

Barring a resolution of the histogram conjecture, let us discuss what properties of the curve C can be gleaned from its histogram. First of all, the curve's diameter is equal to the minimal value of r for which $H_C(r) = 1$. Secondly, values where the derivative of the histogram function is very large usually have geometric significance. In the square histogram in Figure 6, this occurs at $r = 1$. In polygons, such values often correspond to distances between parallel sides, because, at such a distance, the disk centered on one of the parallel sides suddenly begins to contain points on the opposite side. For shapes with multiple pairs of parallel sides, we can see this effect at several values of r — such as when $r = 2$ and $r = 3$ in the case of a 2×3 rectangle shown in Figure 7. The magnitude of the effect depends on the overall length of the parallel sides; for instance, the slope at $r = 3$ is larger than that at $r = 2$. However, not every value where the derivative is large is the result of such parallel sides. The histogram function of the Boutin–Kemper kite shown in Figure 12 has two visible jumps, but the kite has no parallel sides.

In a more theoretical direction, let us compute the Taylor expansion of the global histogram function $H_C(r)$ at $r = 0$, assuming that C is sufficiently smooth. The coefficients in the expansion will provide Euclidean-invariant quantities associated with a smooth curve. We begin by constructing the Taylor series of the local histogram function $h_C(r, z)$ based at a point $z \in C$. To expedite the analysis, we apply a suitable rigid motion to move the curve into a “normal form” so that z is at the origin, and the tangent at z is horizontal. Thus, in a neighborhood of $z = (0, 0)$, the curve is the graph of a function $y = y(x)$ with $y(0) = 0$ and $y'(0) = 0$. As a consequence of the moving frame recurrence formulae developed in [9] — or working by direct analysis — we can write down the following Taylor expansion.

Lemma 19. *Under the above assumptions,*

$$y = \frac{1}{2} \kappa x^2 + \frac{1}{6} \kappa_s x^3 + \frac{1}{24} (\kappa_{ss} + 3\kappa^3) x^4 + \frac{1}{120} (\kappa_{sss} + 19\kappa^2 \kappa_s) x^5 + \dots, \quad (7.5)$$

where $\kappa, \kappa_s, \kappa_{ss}, \dots$ denote, respectively, the curvature and its successive arc length derivatives evaluated at $z = (0, 0)$.

We use this formula to find a Taylor expansion for the local histogram function $h_C(r, z)$ at $r = 0$. Assume that r is small. The curve (7.5) will intersect the circle of radius r centered at

the origin at two points $z_{\pm} = (x_{\pm}, y_{\pm}) = (x_{\pm}, y(x_{\pm}))$, which are the solutions to the equation

$$x^2 + y(x)^2 = r^2.$$

Substituting the expansion (7.5) and solving the resulting series equation for x , we find

$$\begin{aligned} x_+ &= r - \frac{1}{8} \kappa^2 r^3 - \frac{1}{12} \kappa \kappa_s r^4 - \left(\frac{1}{48} \kappa \kappa_{ss} + \frac{1}{72} \kappa_s^2 + \frac{1}{128} \kappa^4 \right) r^5 + \dots, \\ x_- &= -r + \frac{1}{8} \kappa^2 r^3 - \frac{1}{12} \kappa \kappa_s r^4 + \left(\frac{1}{48} \kappa \kappa_{ss} + \frac{1}{72} \kappa_s^2 + \frac{1}{128} \kappa^4 \right) r^5 + \dots. \end{aligned} \quad (7.6)$$

Thus, again using (7.5),

$$\begin{aligned} L h_C(r, z) &= \int_{x_-}^{x_+} \sqrt{1 + y'(x)^2} dx \\ &= \int_{x_-}^{x_+} \sqrt{1 + \kappa^2 x^2 + \kappa \kappa_s x^3 + \left(\frac{1}{3} \kappa \kappa_{ss} + \frac{1}{4} \kappa_s^2 + \kappa^4 \right) x^4 + \dots} dx \\ &= \int_{x_-}^{x_+} \left[1 + \frac{1}{2} \kappa^2 x^2 + \frac{1}{2} \kappa \kappa_s x^3 + \left(\frac{1}{6} \kappa \kappa_{ss} + \frac{1}{8} \kappa_s^2 + \frac{3}{8} \kappa^4 \right) x^4 + \dots \right] dx \\ &= \left[x_+ + \frac{1}{6} \kappa^2 x_+^3 + \frac{1}{8} \kappa \kappa_s x_+^4 + \left(\frac{1}{30} \kappa \kappa_{ss} + \frac{1}{40} \kappa_s^2 + \frac{3}{40} \kappa^4 \right) x_+^5 + \dots \right] - \\ &\quad - \left[x_- + \frac{1}{6} \kappa^2 x_-^3 + \frac{1}{8} \kappa \kappa_s x_-^4 + \left(\frac{1}{30} \kappa \kappa_{ss} + \frac{1}{40} \kappa_s^2 + \frac{3}{40} \kappa^4 \right) x_-^5 + \dots \right]. \end{aligned}$$

We now substitute (7.6) to produce

$$\begin{aligned} L h_C(r, z) &= \left(r + \frac{1}{24} \kappa^2 r^3 + \frac{1}{24} \kappa \kappa_s r^4 + \left(\frac{1}{80} \kappa \kappa_{ss} + \frac{1}{90} \kappa_s^2 + \frac{3}{640} \kappa^4 \right) r^5 + \dots \right) - \\ &\quad - \left(-r - \frac{1}{24} \kappa^2 r^3 + \frac{1}{24} \kappa \kappa_s r^4 - \left(\frac{1}{80} \kappa \kappa_{ss} + \frac{1}{90} \kappa_s^2 + \frac{3}{640} \kappa^4 \right) r^5 + \dots \right) \quad (7.7) \\ &= 2r + \frac{1}{12} \kappa^2 r^3 + \left(\frac{1}{40} \kappa \kappa_{ss} + \frac{1}{45} \kappa_s^2 + \frac{3}{320} \kappa^4 \right) r^5 + \dots. \end{aligned}$$

Invariance of both sides of this formula under rigid motions implies that the formula holds as written at any point $z \in C$.

To obtain the Taylor expansion of the global histogram function, we substitute (7.7) back into (2.9), resulting in

$$H_C(r) = \frac{2r}{L} + \frac{r^3}{12L^2} \int_C \kappa^2 ds + \frac{r^5}{5L^2} \int_C \left(\frac{1}{8} \kappa \kappa_{ss} + \frac{1}{9} \kappa_s^2 + \frac{3}{64} \kappa^4 \right) ds + \dots. \quad (7.8)$$

If C is a closed curve, then we can integrate by parts to simplify the final integral:

$$H_C(r) = \frac{2r}{L} + \frac{r^3}{12L^2} \oint_C \kappa^2 ds + \frac{r^5}{40L^2} \oint_C \left(\frac{3}{8} \kappa^4 - \frac{1}{9} \kappa_s^2 \right) ds + \dots. \quad (7.9)$$

Each integral appearing in the Taylor expansion is uniquely determined by the histogram function. An interesting question is whether the resulting collection of integral moments, depending on curvature and its arc length derivatives, uniquely prescribes the curve up to rigid motion. If so, this would establish the validity of our conjecture for smooth curves.

8 Extensions.

There are a number of interesting directions in which this research program can be extended. The most obvious is to apply it to more substantial practical problems in order to gauge whether histogram-based methods can compete with other algorithms for object recognition and classification, particularly in noisy images. In this direction, the method of shape distributions, [19], touted for its invariance, simplicity, and robustness, employs a variety of discrete local and invariant global histograms for distinguishing three-dimensional objects, including distances between points, areas of triangles, volumes of tetrahedra, and angles between segments. An unanswered question is to what extent the corresponding limiting histograms can actually distinguish inequivalent objects, under the appropriate transformation group: Euclidean, equi-affine, conformal, etc.

8.1 Higher Dimensions

Extending our analysis to objects in three or more dimensions requires minimal change to the methodology. For instance, local and global histogram functions of space curves $C \subset \mathbb{R}^3$ are defined by simply replacing the disk of radius r by the solid ball of that radius in the formulas (2.8), (2.9). For example, consider the saddle-like curve parametrized by

$$z(t) = (\cos t, \sin t, \cos 2t), \quad 0 \leq t \leq 2\pi. \quad (8.1)$$

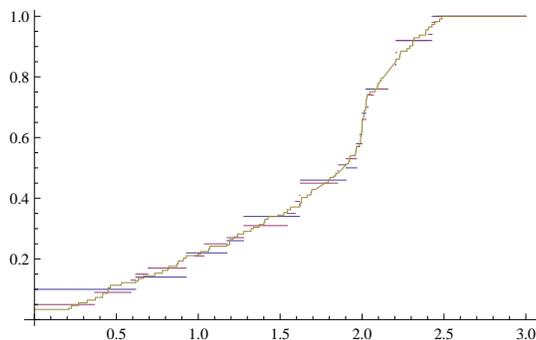


Figure 14: Approximate Distance Histograms for the Three-Dimensional Saddle Curve.

In Figure 14, we plot the discrete approximations $\Lambda_n(r)$ to the curve histogram. The blue plot corresponds to $n = 10$ points, purple to $n = 20$, and yellow $n = 30$. Note that the discrete cumulative histograms appear to converge as $n \rightarrow \infty$.

We can also apply our histogram analysis to two-dimensional surfaces in three-dimensional space. We consider the case of piecewise smooth surfaces $S \subset \mathbb{R}^3$ with finite surface area. Let $P_n \subset S$ be a set of n sample points that are (approximately) uniformly distributed with respect to surface area. We retain the meaning of $\lambda_n(r, z)$ as the proportion of points within a distance r of the point z , (2.4), and $\Lambda_n(r)$ as its average, (2.6). By adapting our proof of Theorem 4 and assuming sufficient regularity of the surface, one can demonstrate that the discrete cumulative histograms $\lambda_n(r, z)$ and $\Lambda_n(r)$ converge, as $n \rightarrow \infty$, to the corresponding

local and global surface histogram functions

$$h_S(r, z) = \frac{\text{area}(S \cap B_r(z))}{\text{area}(S)}, \quad H_S(r) = \frac{1}{\text{area}(S)} \iint_S h_S(r, z) dS. \quad (8.2)$$

The convergence of the discrete histograms is illustrated in Figure 15. Plots of the discrete approximations $\Lambda_n(r)$ for the unit sphere $S^2 = \{\|z\| = 1\} \subset \mathbb{R}^3$ are shown, with $n = 10$ in blue, $n = 30$ in purple, and $n = 100$ in yellow. The global histograms are evidently converging as $n \rightarrow \infty$, albeit at a slower rate than was the case with curves.

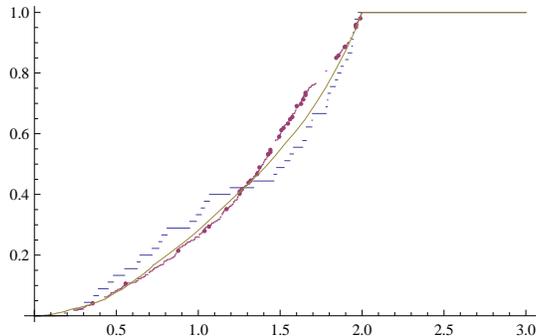


Figure 15: Approximate Distance Histograms of a Sphere.

Future work includes rigorously establishing a convergence theorem for surfaces and higher-dimensional submanifolds of Euclidean space along the lines of Theorem 4. Invariance under rigid motions immediately implies that surfaces with distinct distance histograms cannot be rigidly equivalent. However, it seems unlikely that distance histograms alone suffice to distinguish inequivalent surfaces, and extensions to distance histograms involving more than two points, e.g., that are formed from the side lengths of sampled triangles, are under active investigation. An interesting question is whether distance histograms can be used to distinguish subsets of differing dimensions. Or, to state this another way, can one determine the dimension of a subset from some innate property of its distance histogram?

8.2 Area Histograms

In image processing applications, the invariance of objects under the *equi-affine group*, consisting of all area- or volume-preserving affine transformations of \mathbb{R}^n , is of great importance, [7, 11, 18]. Planar equi-affine (area-preserving) transformations can be viewed as approximations to projective transformations, valid for moderately tilted objects. For example, a round plate viewed at an angle has an elliptical outline, which can be obtained from a circle by an equi-affine transformation. The basic planar equi-affine joint invariant is the area of a triangle, and hence the histogram formed by the areas of triangles formed by all triples in a finite point configuration is invariant under the equi-affine group. Similar to Theorem 2, Boutin and Kemper, [5] also proved that, in most situations, generic planar point configurations are uniquely determined, up to equi-affine transformations, by their area histograms, but there is a lower-dimensional algebraic subvariety of exceptional configurations.

For us, the key question is convergence of the cumulative area histogram based on densely sample points on a plane curve. To define the continuous area histogram, we first note that the distance histogram function (2.9) can be expressed in the alternative form

$$H_C(r) = \frac{1}{L^2} \int_C \int_C \chi_r(d(z(s), z(s'))) ds ds', \quad (8.3)$$

where

$$\chi_r(t) = \begin{cases} 1, & t \leq r, \\ 0, & t > r, \end{cases}$$

denotes the indicator or characteristic function for the disk of radius r . By analogy, we define the area histogram function

$$A_C(r) = \frac{1}{L^3} \oint_C \oint_C \oint_C \chi_r(\text{Area}(z(\hat{s}), z(\hat{s}'), z(\hat{s}''))) d\hat{s} d\hat{s}' d\hat{s}'', \quad (8.4)$$

where $\hat{s}, \hat{s}', \hat{s}''$ now refer to the *equi-affine arc length* of the curve, [10], while $L = \oint_C d\hat{s}$ is its total equi-affine arc length. (In local coordinates, if the curve is the graph of a function $y(x)$ then the equi-affine arc length element is given by $d\hat{s} = \sqrt[3]{y''(x)} dx$.)

The corresponding approximate cumulative area histogram is

$$A_P(r) = \frac{1}{n(n-1)(n-2)} \sum_{z \neq z' \neq z'' \in P} \chi_r(\text{Area}(z, z', z'')), \quad (8.5)$$

which, under suitable equi-affine regularity conditions on the curve, and provided the points are uniformly or randomly distributed with respect to equi-affine arc length, can be shown to converge to the area histogram function (8.4). (Details will appear elsewhere.) We calculate and plot this quantity, and observe that it does converge to the area histogram function (8.4) as $n \rightarrow \infty$. Figure 16 illustrates the convergence for a circle, taking $n = 10$ points in blue, $n = 20$ in purple, and $n = 30$ in yellow.

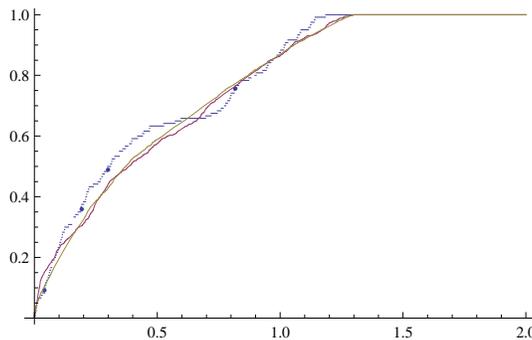


Figure 16: Area Histogram of a Circle.

Let us end by illustrating the equi-affine invariance of the curve area histogram function. Since rectangles of the same area are equivalent under an equi-affine transformation, they have identical area histograms. In Figure 17, we plot area histograms for a 2×2 square in

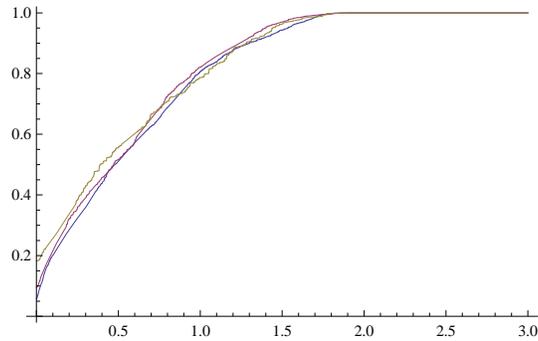


Figure 17: Area Histograms of Affine-Equivalent Rectangles.

blue, a 1×4 rectangle in purple, and a $.5 \times 8$ rectangle in yellow, using $n = 30$ sample points in each case. As expected, the graphs are quite close.

Acknowledgments: The authors would like to thank Facundo Memoli, Igor Pak, Ellen Rice, Guillermo Sapiro, Allen Tannenbaum, and Ofer Zeitouni, as well as the anonymous referees, for helpful comments and advice. The paper is based on the first author’s undergraduate research project (REU) funded by the second author’s NSF Grant DMS 08–07317.

References

- [1] Ankerst, M., Kastenmüller, G., Kriegel, H.–P., and Seidl, T., 3D shape histograms for similarity search and classification in spatial databases, *in: Advances in Spatial Databases*, R.H. Güting, D. Papadias, F. Lochovsky, eds., Lecture Notes in Comp. Sci., vol. 1651, Springer–Verlag, New York, 1999, pp. 207–226.
- [2] Apostol, T.M., *Mathematical Analysis*, Second Edition, Addison–Wesley Publ. Co., Reading, Mass., 1974.
- [3] Belongie, S., Malik, J., and Puzicha, J., Shape matching and object recognition using shape contexts, *IEEE Trans. Pattern Anal. Mach. Intell.* **24** (2002), 509–522.
- [4] Bloom, G.S., A counterexample to a theorem of S. Piccard, *J. Comb. Theory, Ser. A* **22** (1977), 378–379.
- [5] Boutin, M., and Kemper, G., On reconstructing n -point configurations from the distribution of distances or areas, *Adv. Appl. Math.* **32** (2004), 709–735.
- [6] Boutin, M., and Kemper, G., Which point configurations are determined by the distribution of their pairwise distances?, *Internat. J. Comput. Geom. Appl.* **17** (2007), 31–43.
- [7] Calabi, E., Olver, P.J., Shakiban, C., Tannenbaum, A., and Haker, S., Differential and numerically invariant signature curves applied to object recognition, *Int. J. Computer Vision* **26** (1998), 107–135.

- [8] DeGroot, M.H., and Schervish, M.J., *Probability and Statistics*, Addison–Wesley, Redwood City, Calif., 2002.
- [9] Fels, M., and Olver, P.J., Moving coframes. II. Regularization and theoretical foundations, *Acta Appl. Math.* **55** (1999), 127–208.
- [10] Guggenheimer, H.W., *Differential Geometry*, McGraw–Hill, New York, 1963.
- [11] Kanatani, K., *Group–Theoretical Methods in Image Understanding*, Springer–Verlag, New York, 1990.
- [12] Lowe, D.G., Object recognition from local scale-invariant features, in: *Integration of Speech and Image Understanding*, IEEE Computer Society, Los Alamitos, CA, 1999, pp. 1150–1157.
- [13] Manay, S., Cremers, D., Hong, B.-W., Yezzi, A., and Soatto, S., Integral invariants and shape matching, in: *Statistics and Analysis of Shapes*, H. Krim and A. Yezzi, eds., Birkhäuser, Boston, 2006, pp. 137–166.
- [14] Mémoli, F., On the use of Gromov–Hausdorff distances for shape comparison, in: *Symposium on Point Based Graphics*, M. Botsch, R. Pajarola, B. Chen, and M. Zwicker, eds., Eurographics Association, Prague, Czech Republic, 2007, pp. 81–90.
- [15] Mémoli, F., Gromov–Wasserstein distances and the metric approach to object matching, *Found. Comput. Math.* **11** (2011), 417–487.
- [16] Nutbourne, A.W., and Martin, R.R., *Differential Geometry Applied to Curve and Surface Design, Vol. 1: Foundations*, Ellis Horwood, Chichester, 1988.
- [17] Olver, P.J., Joint invariant signatures, *Found. Comput. Math.* **1** (2001), 3–67.
- [18] Olver, P.J., Sapiro, G., and Tannenbaum, A., Differential invariant signatures and flows in computer vision: a symmetry group approach, in: *Geometry–Driven Diffusion in Computer Vision*, B. M. Ter Haar Romeny, ed., Kluwer Acad. Publ., Dordrecht, the Netherlands, 1994, pp. 255–306.
- [19] Osada, R., Funkhouser, T., Chazelle, B., and Dobkin, D., Shape distributions, *ACM Trans. Graphics* **21** (2002), 807–832.
- [20] Pele, O., and Werman, M., A linear time histogram for improved SIFT matching, in: *Computer Vision - ECCV 2008*, part III, D. Forsyth, P. Torr, A. Zisserman, eds., Lecture Notes in Computer Science, vol. 5304, Springer–Verlag, Berlin, 2008, pp. 495–508.
- [21] Pottmann, H., Wallner, J., Huang, Q., and Yang, Y.-L., Integral invariants for robust geometry processing, *Comput. Aided Geom. Design* **26** (2009), 37–60.
- [22] Rustamov, R.M., Laplace–Beltrami eigenfunctions for deformation invariant shape representation, in: *SGP '07: Proceedings of the Fifth Eurographics Symposium on Geometry Processing*, Eurographics Association; Aire-la-Ville, Switzerland, 2007, pp. 225–233.

- [23] Sapiro, G., *Geometric Partial Differential Equations and Image Analysis*, Cambridge University Press, Cambridge, 2001.
- [24] Şaykol, E., Güdükbaya, U., and Ulusoya, Ö., A histogram-based approach for object-based query-by-shape-and-color in image and video databases, *Image Vision Comput.* **23** (2005), 1170–1180.
- [25] Sonka, M., Havlac, V., and Boyle, R., *Image Processing: Analysis and Machine Vision*, Brooks/Cole, Pacific Grove, CA, 1999.
- [26] Vogt, W., Über monotongekrümmte Kurven, *J. Reine Angew. Math.* **144** (1914), 239–248.
- [27] Yale, P.B., *Geometry and Symmetry*, Holden–Day, San Francisco, 1968.