Invariant Signatures and Histograms for Object Recognition, Symmetry Detection, and Jigsaw Puzzle Assembly

Peter J. Olver

University of Minnesota

http://www.math.umn.edu/∼olver

University of Chicago, April, 2013

Geometry = Group Theory

Felix Klein's Erlanger Programm (1872):

Each type of geometry is founded on
a transformation group.

Geometry = Group Theory

Felix Klein's Erlanger Programm (1872):

Each type of geometry is founded on a transformation group.

A group G acts on a space M via
$z \longmapsto g \cdot z$, with

- $g \cdot (h \cdot z) = (g \cdot h) \cdot z$
- $e \cdot z = z$

for all $g, h \in G$ and $z \in M$.

Plane Geometries/Groups

Euclidean geometry:

SE(2) — rigid motions (rotations and translations)

$$\begin{pmatrix} \bar{x} \\ \bar{y} \end{pmatrix} = \begin{pmatrix} \cos\theta & -\sin\theta \\ \sin\theta & \cos\theta \end{pmatrix} \begin{pmatrix} x \\ y \end{pmatrix} + \begin{pmatrix} a \\ b \end{pmatrix}$$

E(2) — plus reflections?

Equi-affine geometry:

SA(2) — area-preserving affine transformations:

$$\begin{pmatrix} \bar{x} \\ \bar{y} \end{pmatrix} = \begin{pmatrix} \alpha & \beta \\ \gamma & \delta \end{pmatrix} \begin{pmatrix} x \\ y \end{pmatrix} + \begin{pmatrix} a \\ b \end{pmatrix} \qquad \alpha\,\delta - \beta\,\gamma = 1$$

Projective geometry:

PSL(3) — projective transformations:

$$\bar{x} = \frac{\alpha\,x + \beta\,y + \gamma}{\rho\,x + \sigma\,y + \tau} \qquad \bar{y} = \frac{\lambda\,x + \mu\,y + \nu}{\rho\,x + \sigma\,y + \tau}$$

The Basic Equivalence Problem

The Basic Equivalence Problem

G — transformation group acting on M

The Basic Equivalence Problem

G — transformation group acting on M

Equivalence:

Determine when two subsets

$$N \quad \text{and} \quad \overline{N} \subset M$$

are congruent:

$$\overline{N} = g \cdot N \qquad \text{for} \qquad g \in G$$

The Basic Equivalence Problem

G — transformation group acting on M

Equivalence:

Determine when two subsets

$$N \quad \text{and} \quad \overline{N} \ \subset \ M$$

are congruent:

$$\overline{N} = g \cdot N \qquad \text{for} \qquad g \in G$$

Symmetry:

Find all symmetries,

i.e., self-equivalences or *self-congruences*:

$$N = g \cdot N$$

Tennis, Anyone?

Invariants

The solution to an equivalence problem rests on understanding its invariants.

Definition. If G is a group acting on M, then an invariant is a real-valued function $I : M \to \mathbb{R}$ that does not change under the action of G:

$$I(g \cdot z) = I(z) \qquad \text{for all} \qquad g \in G, \quad z \in M$$

★ If G acts transtively, there are no (non-constant) invariants.

Joint Invariants

A joint invariant is an invariant of the k-fold Cartesian product action of G on $M \times \cdots \times M$:

$$\boxed{I(g \cdot z_1, \ldots, g \cdot z_k) \;=\; I(z_1, \ldots, z_k)}$$

Joint Euclidean Invariants

Theorem. Every joint Euclidean invariant is a function of the interpoint distances

$$d(z_i, z_j) = \| \, z_i - z_j \, \|$$

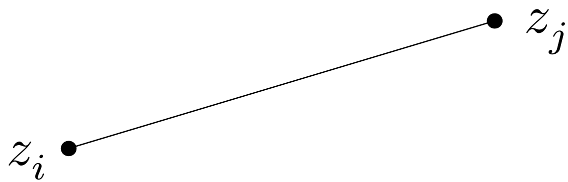

Joint Equi–Affine Invariants

Theorem. Every planar joint equi–affine invariant is a function of the triangular areas

$$A(i,j,k) = \tfrac{1}{2}\,(z_i - z_j) \wedge (z_i - z_k)$$

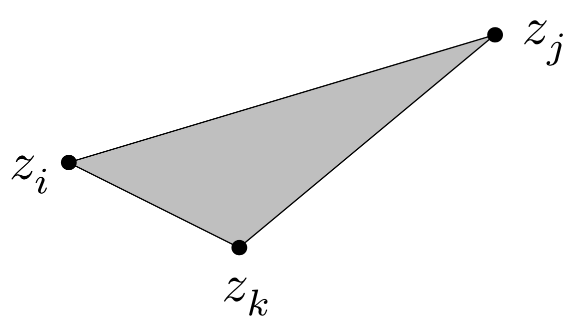

Joint Projective Invariants

Theorem. Every joint projective invariant is a function of the planar cross-ratios

$$[\, z_i, z_j, z_k, z_l, z_m \,] \;=\; \frac{A\,B}{C\,D}$$

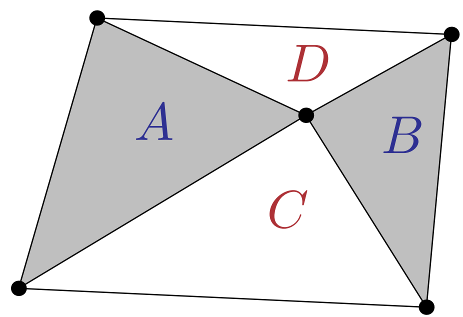

Differential Invariants

Given a submanifold (curve, surface, . . .) $N \subset M$, a differential invariant is an invariant of the action of G on N and its derivatives (jets).

$$\boxed{I(g \cdot z^{(k)}) \;=\; I(z^{(k)})}$$

Euclidean Plane Curves: $G = \mathrm{SE}(2)$

The simplest differential invariant is the curvature, defined as the reciprocal of the radius of the osculating circle:

$$\kappa = \frac{1}{r}$$

Curvature

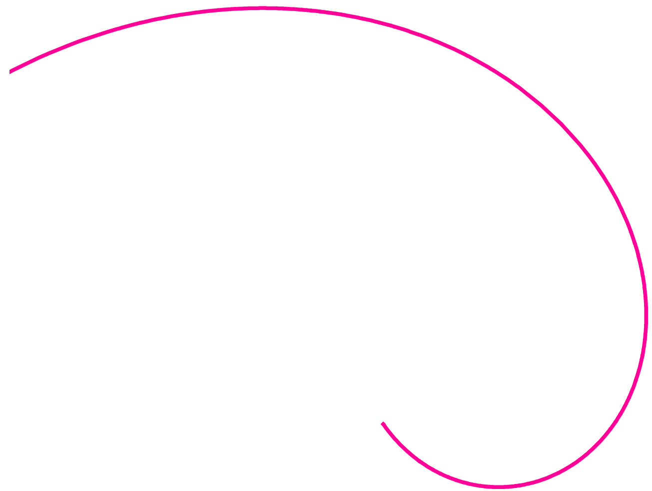

Curvature

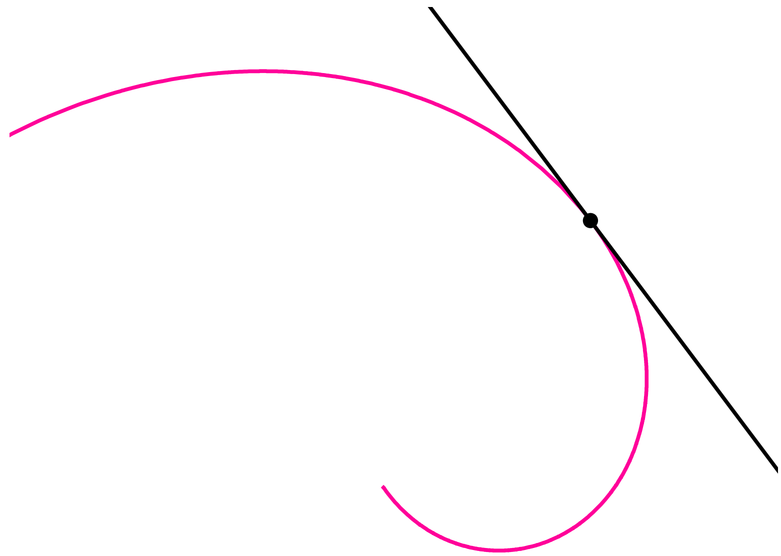

Curvature

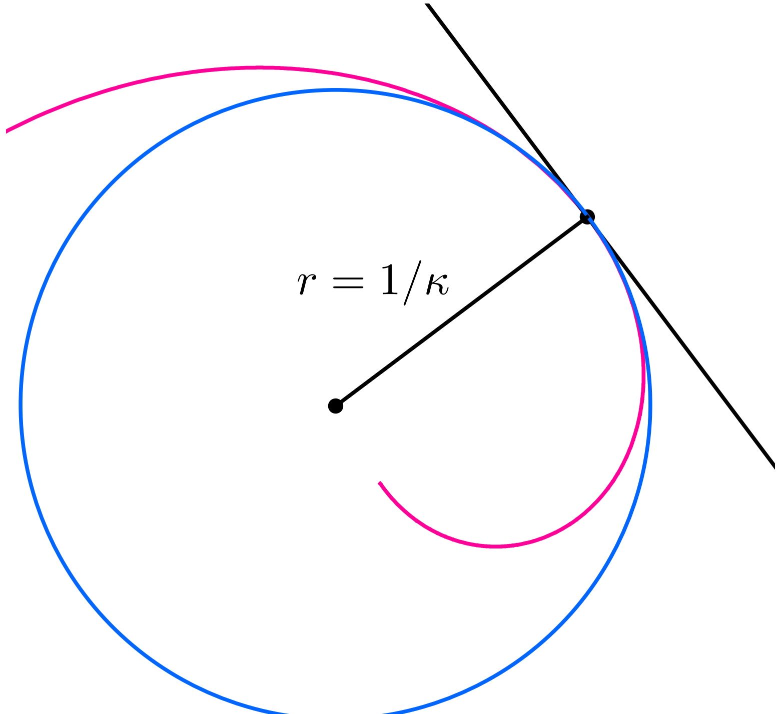

Euclidean Plane Curves: $G = \mathrm{SE}(2) = \mathrm{SO}(2) \ltimes \mathbb{R}^2$

Assume the curve is a graph: $\quad y = u(x)$

Differential invariants:

$$\kappa = \frac{u_{xx}}{(1+u_x^2)^{3/2}}, \qquad \frac{d\kappa}{ds} = \frac{(1+u_x^2)u_{xxx} - 3u_x u_{xx}^2}{(1+u_x^2)^3}, \qquad \frac{d^2\kappa}{ds^2} = \cdots$$

Arc length (invariant one-form):

$$ds = \sqrt{1+u_x^2}\;dx, \qquad \frac{d}{ds} = \frac{1}{\sqrt{1+u_x^2}}\,\frac{d}{dx}$$

Theorem. All Euclidean differential invariants are functions of the derivatives of curvature with respect to arc length: $\kappa, \quad \kappa_s, \quad \kappa_{ss}, \quad \cdots$

Equi-affine Plane Curves: $G = \mathrm{SA}(2) = \mathrm{SL}(2) \ltimes \mathbb{R}^2$

Equi-affine curvature:

$$\kappa = \frac{5\,u_{xx}u_{xxxx} - 3\,u_{xxx}^2}{9\,u_{xx}^{8/3}} \qquad \frac{d\kappa}{ds} = \cdots$$

Equi-affine arc length:

$$ds = \sqrt[3]{u_{xx}}\ dx \qquad \frac{d}{ds} = \frac{1}{\sqrt[3]{u_{xx}}}\,\frac{d}{dx}$$

Theorem. All equi-affine differential invariants are functions of the derivatives of equi-affine curvature with respect to equi-affine arc length:

$$\kappa, \qquad \kappa_s, \qquad \kappa_{ss}, \qquad \cdots$$

Equivalence & Invariants

- Equivalent submanifolds $\quad N \approx \overline{N}$
  must have the same invariants: $\; I = \bar{I}$.

Equivalence & Invariants

- Equivalent submanifolds $\quad N \approx \overline{N}$

  must have the same invariants: $\; I = \bar{I}$.

Constant invariants provide immediate information:

$$\text{e.g.} \qquad \kappa = 2 \quad \Longleftrightarrow \quad \overline{\kappa} = 2$$

Equivalence & Invariants

- Equivalent submanifolds $\quad N \approx \overline{N}$
  must have the same invariants: $\; I = \bar{I}$.

Constant invariants provide immediate information:

$$\text{e.g.} \qquad \kappa = 2 \quad \Longleftrightarrow \quad \overline{\kappa} = 2$$

Non-constant invariants are not useful in isolation, because an equivalence map can drastically alter the dependence on the submanifold parameters:

$$\text{e.g.} \qquad \kappa = x^3 \qquad \text{versus} \qquad \overline{\kappa} = \sinh x$$

However, a functional dependency or syzygy among the invariants *is* intrinsic:

$$\text{e.g.} \qquad \kappa_s = \kappa^3 - 1 \quad \Longleftrightarrow \quad \overline{\kappa}_{\overline{s}} = \overline{\kappa}^3 - 1$$

However, a functional dependency or syzygy among the invariants *is* intrinsic:

$$\text{e.g.} \qquad \kappa_s = \kappa^3 - 1 \quad \Longleftrightarrow \quad \overline{\kappa}_{\overline{s}} = \overline{\kappa}^3 - 1$$

- Universal syzygies — Gauss–Codazzi
- Distinguishing syzygies.

However, a functional dependency or syzygy among the invariants *is* intrinsic:

$$\text{e.g.} \qquad \kappa_s = \kappa^3 - 1 \quad \Longleftrightarrow \quad \overline{\kappa}_{\overline{s}} = \overline{\kappa}^3 - 1$$

- Universal syzygies — Gauss–Codazzi
- Distinguishing syzygies.

Theorem. (Cartan)
Two regular submanifolds are (locally) equivalent if and only if they have identical syzygies among *all* their differential invariants.

Finiteness of Generators and Syzygies

♠ There are, in general, an infinite number of differential invariants and hence an infinite number of syzygies must be compared to establish equivalence.

Finiteness of Generators and Syzygies

♠ There are, in general, an infinite number of differential invariants and hence an infinite number of syzygies must be compared to establish equivalence.

♡ But the higher order differential invariants are always generated by invariant differentiation from a finite collection of basic differential invariants, and the higher order syzygies are all consequences of a finite number of low order syzygies!

Example — Plane Curves

If non-constant, both κ and κ_s depend on a single parameter, and so, locally, are subject to a syzygy:

$$\kappa_s = H(\kappa) \qquad (*)$$

But then

$$\kappa_{ss} = \frac{d}{ds} H(\kappa) = H'(\kappa)\,\kappa_s = H'(\kappa)\,H(\kappa)$$

and similarly for κ_{sss}, etc.

Consequently, all the higher order syzygies are generated by the fundamental first order syzygy $(*)$.

Thus, for Euclidean (or equi-affine or projective or ...) plane curves we need only know a single syzygy between κ and κ_s in order to establish equivalence!

Signature Curves

Definition. The *signature curve* $\mathcal{S} \subset \mathbb{R}^2$ of a plane curve $\mathcal{C} \subset \mathbb{R}^2$ is parametrized by the two lowest order differential invariants

$$\Sigma \ : \ C \ \longrightarrow \ \mathcal{S} = \left\{ \left(\kappa \ , \ \frac{d\kappa}{ds} \right) \right\} \subset \ \mathbb{R}^2$$

$\Longrightarrow$ *Calabi, PJO, Shakiban, Tannenbaum, Haker*

Theorem. Two regular curves $\mathcal{C}$ and $\overline{\mathcal{C}}$ are equivalent:

$$\overline{\mathcal{C}} = g \cdot \mathcal{C}$$

if and only if their signature curves are identical:

$$\overline{\mathcal{S}} = \mathcal{S}$$

$\Longrightarrow$ regular: $(\kappa_s, \kappa_{ss}) \neq 0.$

Symmetry and Signature

Continuous Symmetries

Theorem. The following are equivalent:

- The curve C has a 1-dimensional symmetry group $H \subset G$
- C is the orbit of a 1-dimensional subgroup $H \subset G$
- The signature $\mathcal{S}$ degenerates to a point: $\dim \mathcal{S} = 0$
- All the differential invariants are constant:

$$\kappa = c, \qquad \kappa_s = 0, \qquad \ldots$$

$\Longrightarrow$ Euclidean plane geometry: circles, lines

$\Longrightarrow$ Equi-affine plane geometry: conic sections.

$\Longrightarrow$ Projective plane geometry: W curves (*Lie & Klein*)

Symmetry and Signature

Discrete Symmetries

Definition. The index of a curve C equals the number of points in C which map to a single generic point of its signature:

$$\iota_C = \min \left\{ \, \# \, \Sigma^{-1}\{w\} \, \middle| \, w \in \mathcal{S} \, \right\}$$

Theorem. The number of discrete symmetries of C equals its index ι_C.

The Index

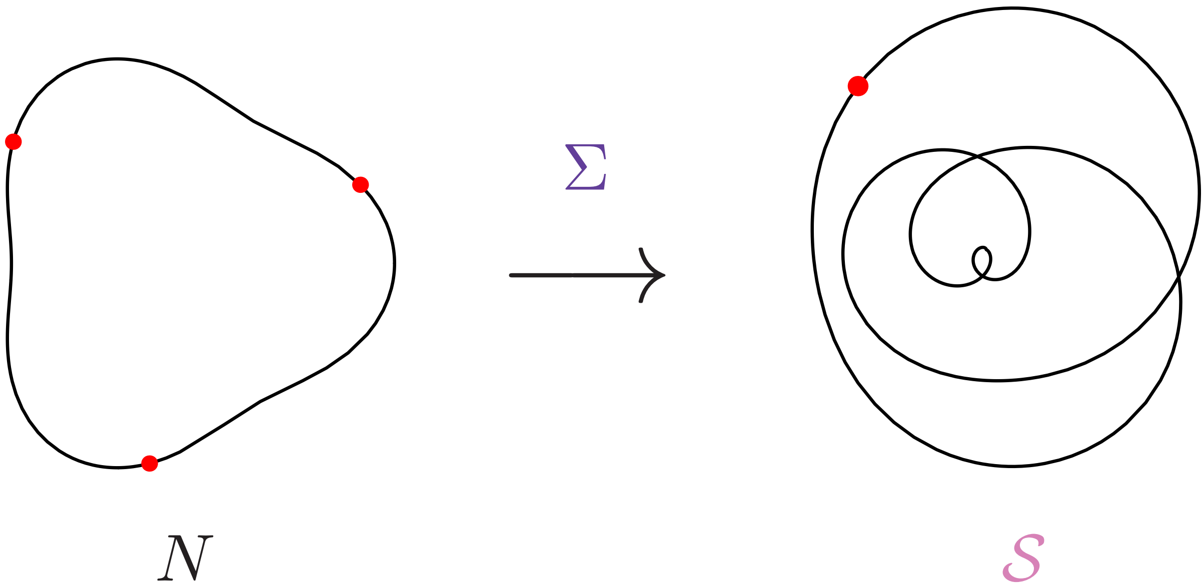

The Curve $x = \cos t + \frac{1}{5}\cos^2 t, \quad y = \sin t + \frac{1}{10}\sin^2 t$

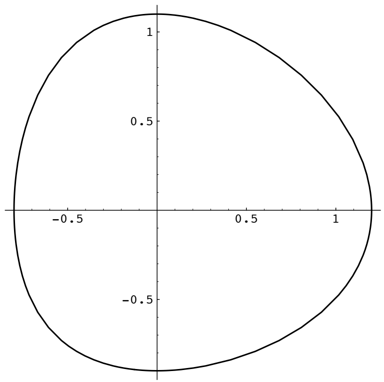

The Original Curve

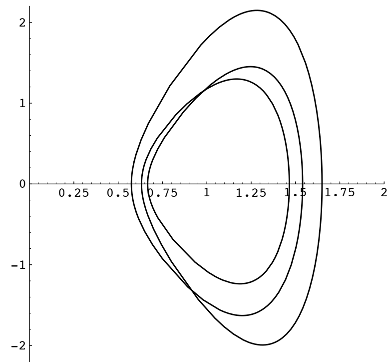

Euclidean Signature

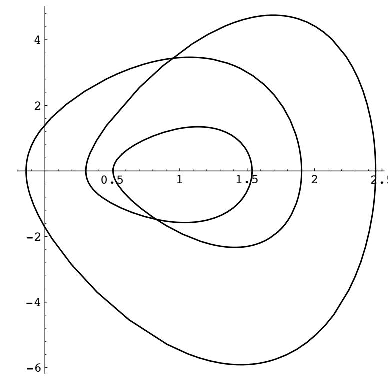

Affine Signature

The Curve $x = \cos t + \frac{1}{5}\cos^2 t, \quad y = \frac{1}{2}x + \sin t + \frac{1}{10}\sin^2 t$

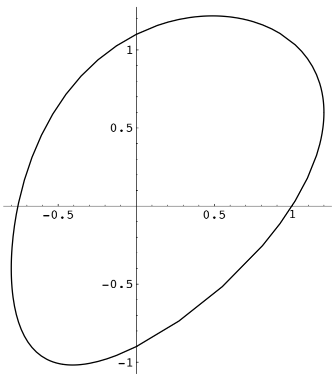

The Original Curve

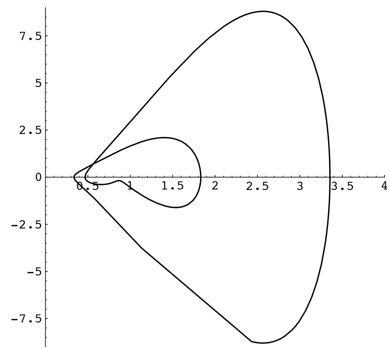

Euclidean Signature

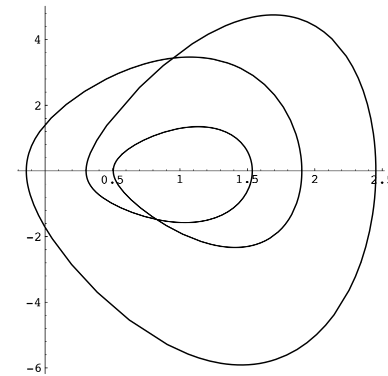

Affine Signature

Canine Left Ventricle Signature

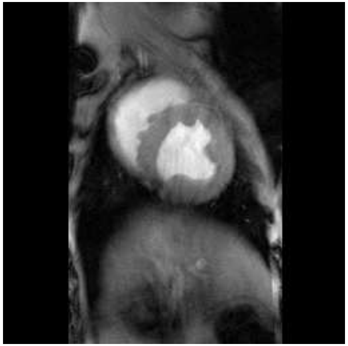

Original Canine Heart
MRI Image

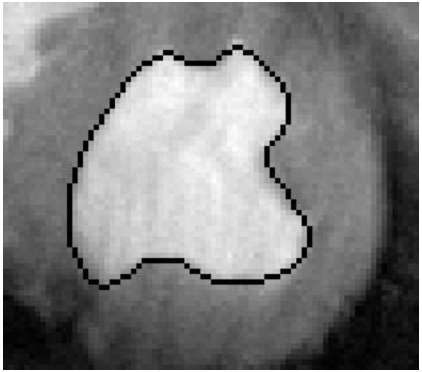

Boundary of Left Ventricle

Smoothed Ventricle Signature

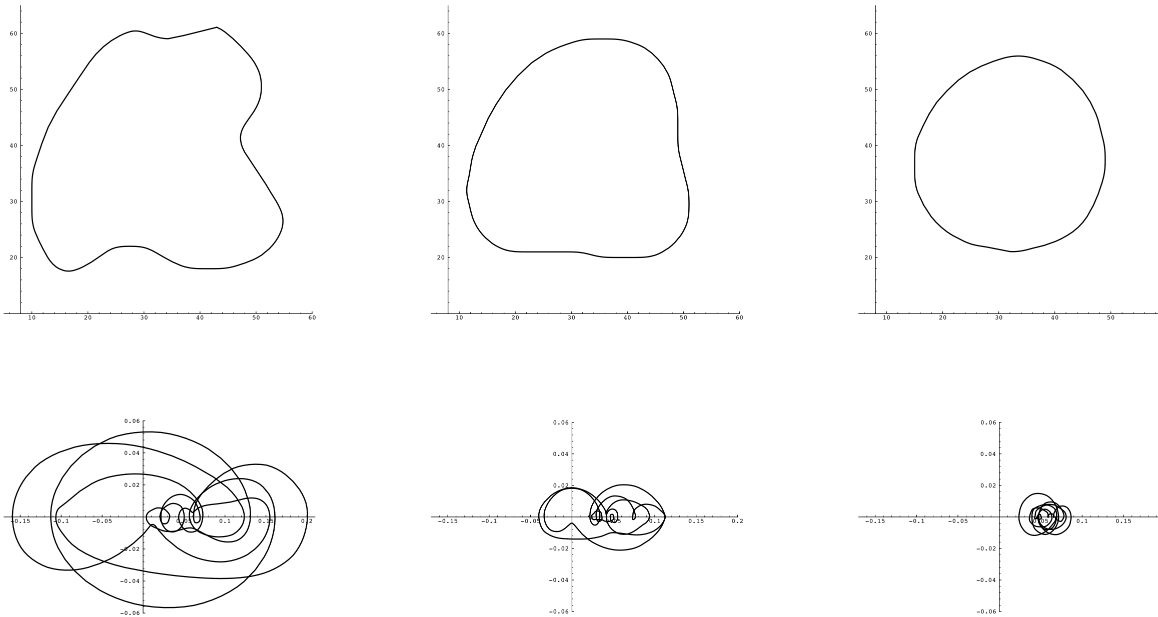

Object Recognition

$\Longrightarrow$ Steve Haker

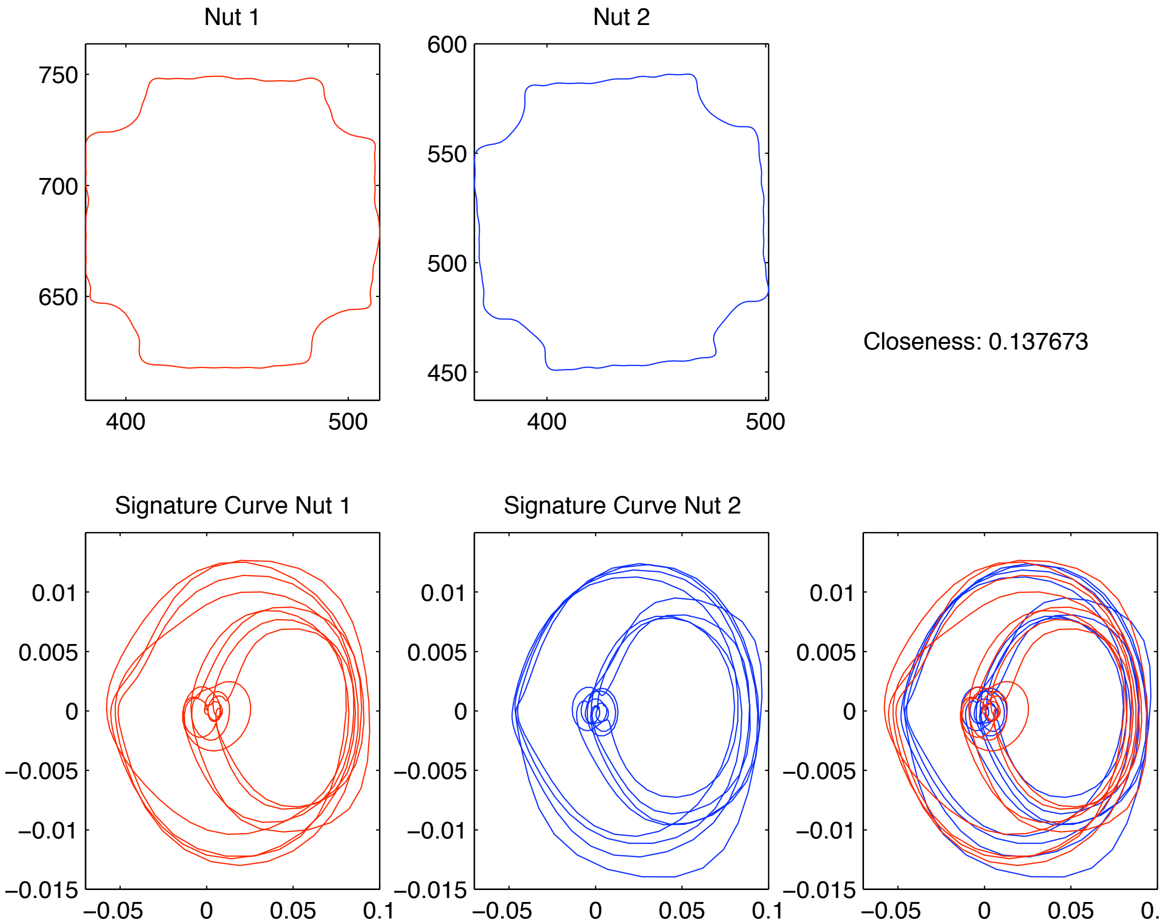
Nut 1
Nut 2
Closeness: 0.137673
Signature Curve Nut 1
Signature Curve Nut 2
750
700
650
600
550
500
450
400
500
0.01
0.005
0
−0.005
−0.01
−0.015
−0.05
0
0.05
0.1

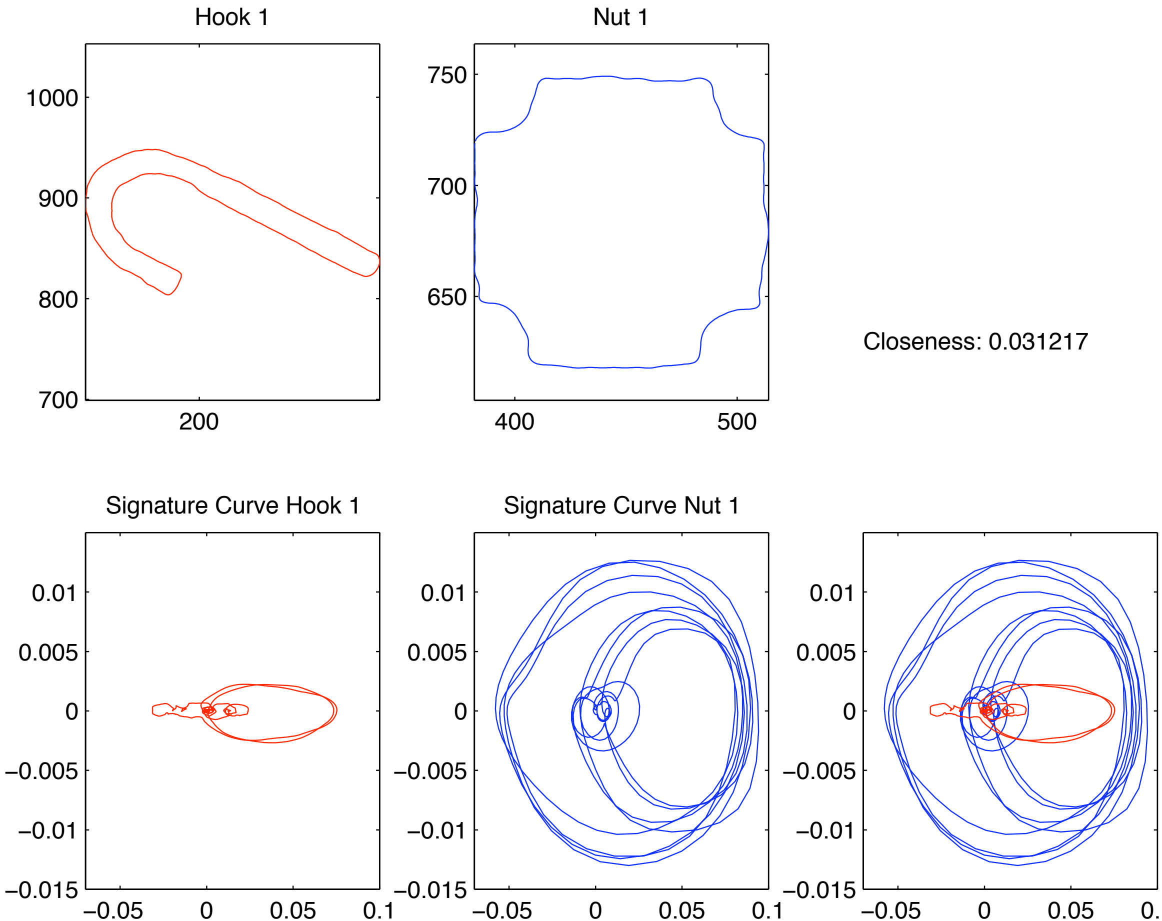
Hook 1
Nut 1
Closeness: 0.031217
Signature Curve Hook 1
Signature Curve Nut 1

Signature Metrics

- Hausdorff
- Monge–Kantorovich transport
- Electrostatic/gravitational attraction
- Latent semantic analysis
- Histograms
- Gromov–Hausdorff & Gromov–Wasserstein

Signatures

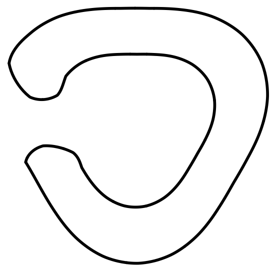

Original curve

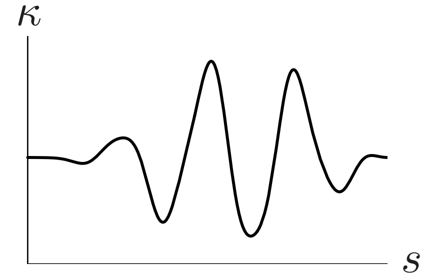

Classical Signature

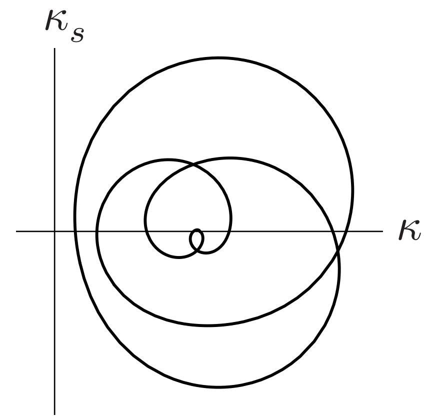

Differential invariant signature

Signatures

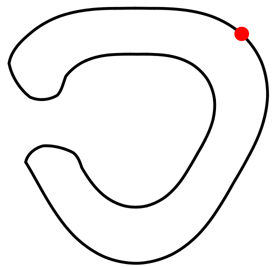

Original curve

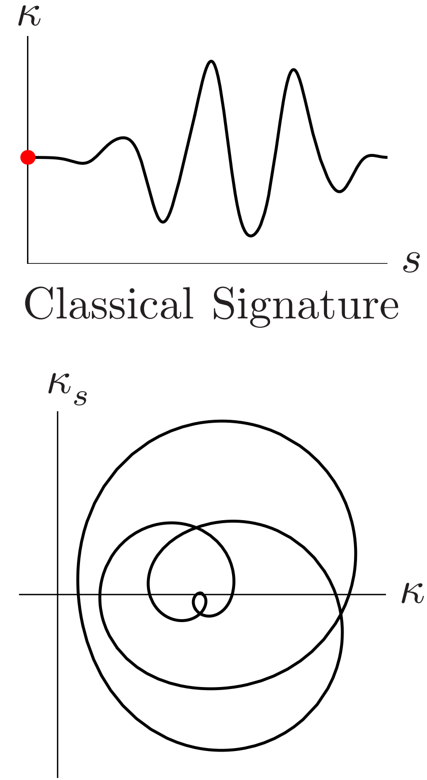

Differential invariant signature

Occlusions

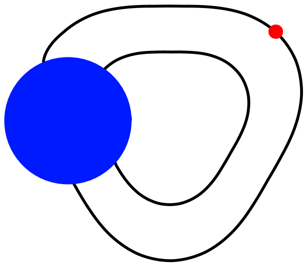

Original curve

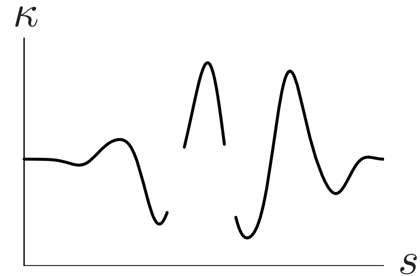

Classical Signature

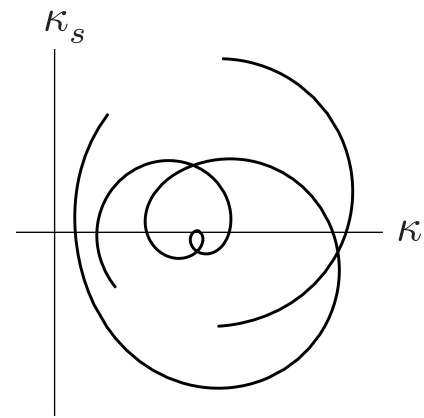

Differential invariant signature

Classical Occlusions

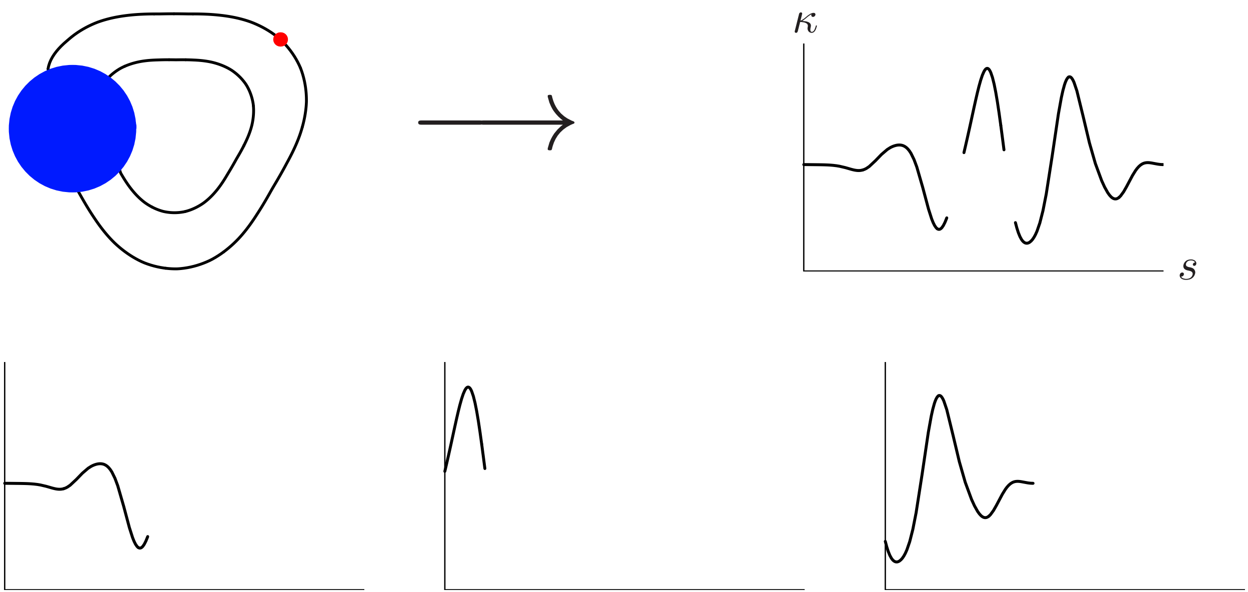

3D Differential Invariant Signatures

Euclidean space curves: $C \subset \mathbb{R}^3$

$$\mathcal{S} = \{ \, (\, \kappa \, , \, \kappa_s \, , \, \tau \,) \, \} \subset \mathbb{R}^3$$

- κ — curvature, τ — torsion

Euclidean surfaces: $S \subset \mathbb{R}^3$ (generic)

$$\mathcal{S} = \left\{ \left(H \, , \, K \, , \, H_{,1} \, , \, H_{,2} \, , \, K_{,1} \, , \, K_{,2} \right) \right\} \subset \mathbb{R}^6$$

or $$\widehat{\mathcal{S}} = \left\{ \left(H \, , \, H_{,1} \, , \, H_{,2} \, , \, H_{,11} \right) \right\} \subset \mathbb{R}^4$$

- H — mean curvature, K — Gauss curvature

Equi–affine surfaces: $S \subset \mathbb{R}^3$ (generic)

$$\mathcal{S} = \left\{ \left(P \, , \, P_{,1} \, , \, P_{,2} , \, P_{,11} \right) \right\} \subset \mathbb{R}^4$$

- P — Pick invariant

Advantages of the Signature Curve

- Purely local — no ambiguities
- Symmetries and approximate symmetries
- Extends to surfaces and higher dimensional submanifolds
- Occlusions and reconstruction
- Partial matching and puzzles

Main disadvantage: Noise sensitivity due to dependence on high order derivatives.

Generalized Vertices

Ordinary vertex: local extremum of curvature

Generalized vertex: $\kappa_s \equiv 0$

- critical point
- circular arc
- straight line segment

Mukhopadhya's Four Vertex Theorem:

A simple closed, non-circular plane curve has $n \geq 4$ generalized vertices.

“Counterexamples”

These degenerate curves all have the same signature:

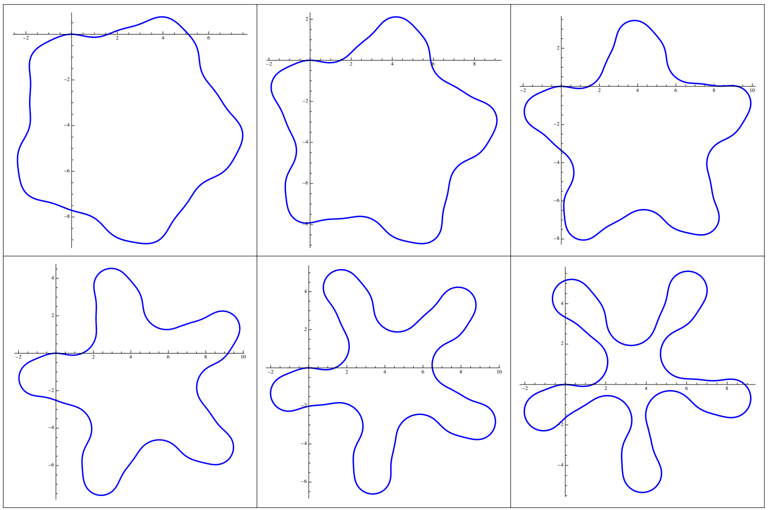

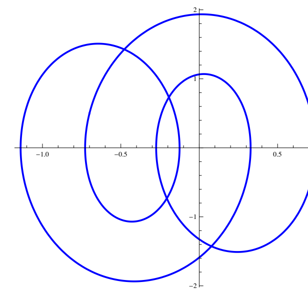

★ Replace vertices with circular arcs: *Musso–Nicoldi*

Bivertex Arcs

Bivertex arc: $\kappa_s \neq 0$ everywhere
except $\kappa_s = 0$ at the two endpoints

The signature S of a bivertex arc is a single arc that starts and ends on the κ–axis.

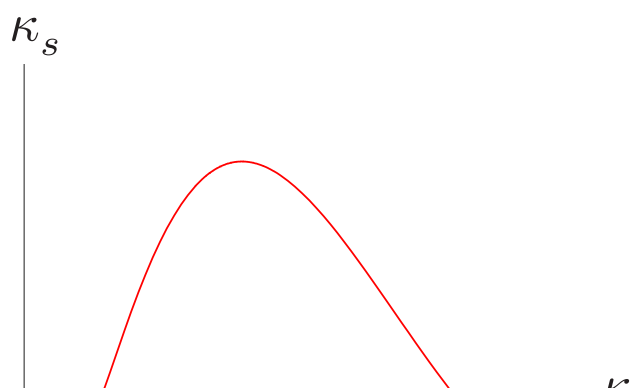

Bivertex Decomposition.

v-regular curve — finitely many generalized vertices

$$\boxed{C = \bigcup_{j=1}^{m} B_j \,\cup\, \bigcup_{k=1}^{n} V_k}$$

$B_1, \ldots, B_m$ — bivertex arcs

$V_1, \ldots, V_n$ — generalized vertices: $n \geq 4$

Main Idea: Compare individual bivertex arcs, and then determine whether the rigid equivalences are (approximately) the same.

D. Hoff & PJO, Extensions of invariant signatures for object recognition, *J. Math. Imaging Vision*, to appear.

Gravitational/Electrostatic Attraction

- Treat the two (signature) curves as masses or as oppositely charged wires. The higher their mutual attraction, the closer they are together.

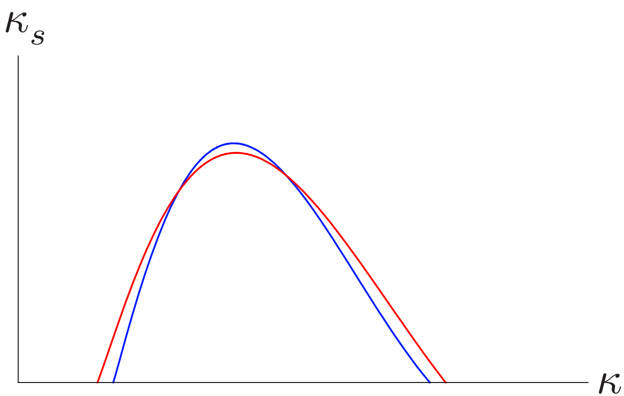

Gravitational/Electrostatic Attraction

- Treat the two (signature) curves as masses or as oppositely charged wires. The higher their mutual attraction, the closer they are together.
- In practice, we are dealing with discrete data (pixels) and so treat the curves and signatures as point masses/charges.

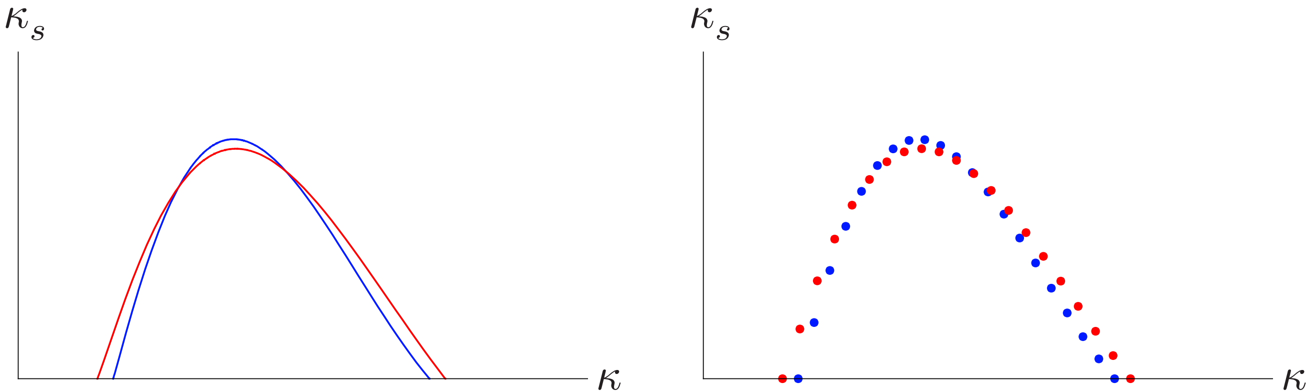

Strength of correspondence:

$$h(\sigma,\tilde{\sigma}) = \begin{cases} \dfrac{1}{d(\sigma,\tilde{\sigma})^{\gamma}+\epsilon}, & d(\sigma,\tilde{\sigma}) < \infty, \\ 0, & d(\sigma,\tilde{\sigma}) = \infty. \end{cases}$$

Separation:

$$d(\sigma,\tilde{\sigma}) = \begin{cases} \dfrac{\|\,\sigma-\tilde{\sigma}\,\|}{D-\|\,\sigma-\tilde{\sigma}\,\|}, & \|\,\sigma-\tilde{\sigma}\,\| < D, \\ \infty, & \|\,\sigma-\tilde{\sigma}\,\| \geq D, \end{cases}$$

Scale of comparison:

$$D(C,\widetilde{C}) = (\, D_{\kappa}(C,\widetilde{C}), D_{\kappa_s}(C,\widetilde{C})\,),$$

$$D_{\kappa}(C,\widetilde{C}) = \max\left\{ \max_{z\in C}(\kappa|_z) - \min_{z\in C}(\kappa|_z),\ \max_{\tilde{z}\in\widetilde{C}}(\kappa|_{\tilde{z}}) - \min_{\tilde{z}\in\widetilde{C}}(\kappa|_{\tilde{z}}) \right\},$$

$$D_{\kappa_s}(C,\widetilde{C}) = \max\left\{ \max_{z\in C}(\kappa_s|_z) - \min_{z\in C}(\kappa_s|_z),\ \max_{\tilde{z}\in\widetilde{C}}(\kappa_s|_{\tilde{z}}) - \min_{\tilde{z}\in\widetilde{C}}(\kappa_s|_{\tilde{z}}) \right\}.$$

Piece Locking

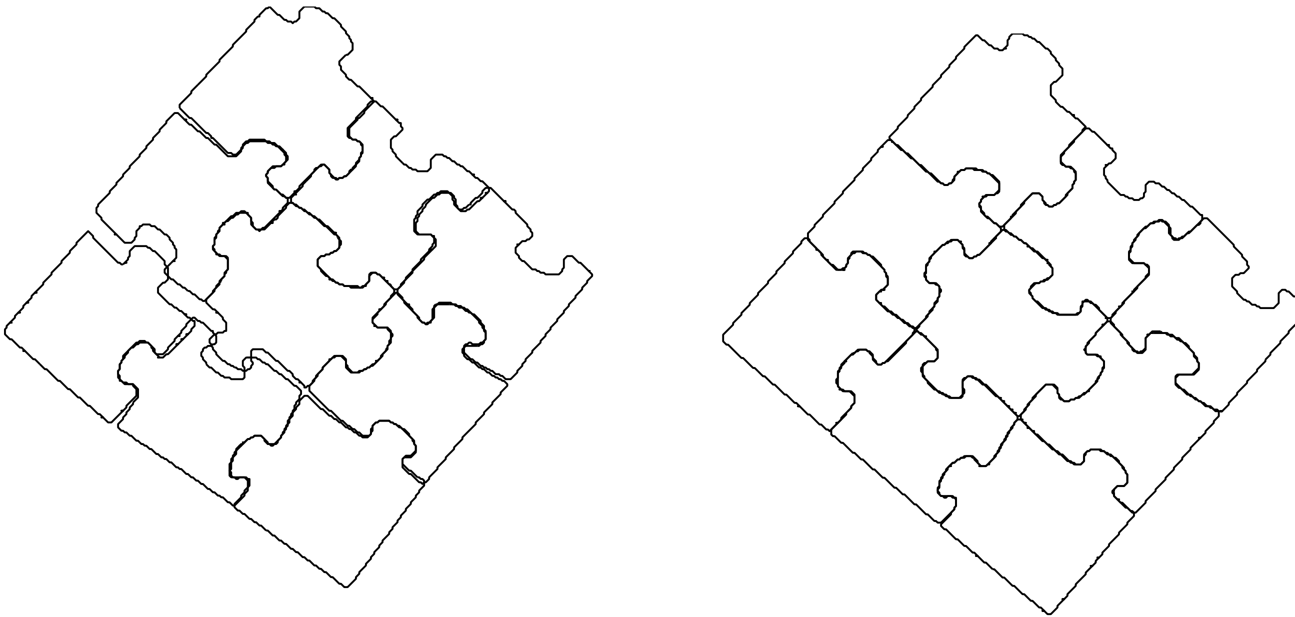

⋆⋆ Minimize force and torque based on gravitational attraction of the two matching edges.

The Baffler Jigsaw Puzzle

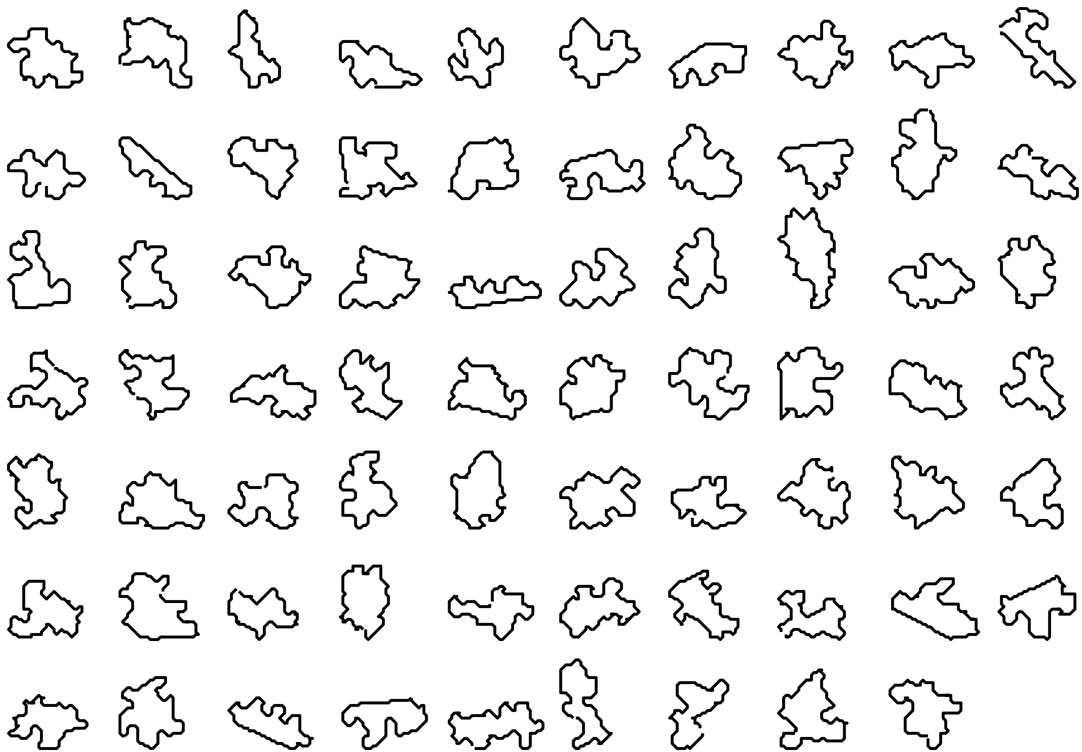

The Baffler Solved

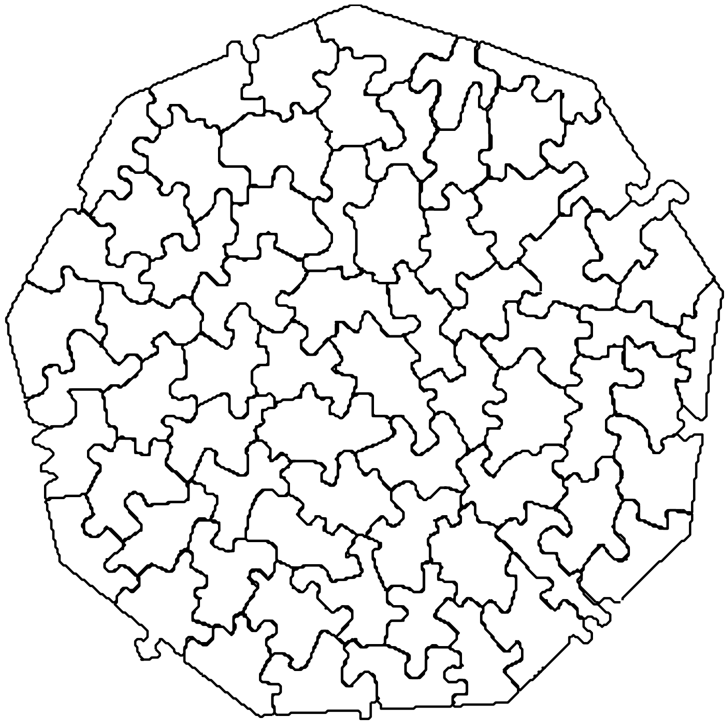

The Rain Forest Giant Floor Puzzle

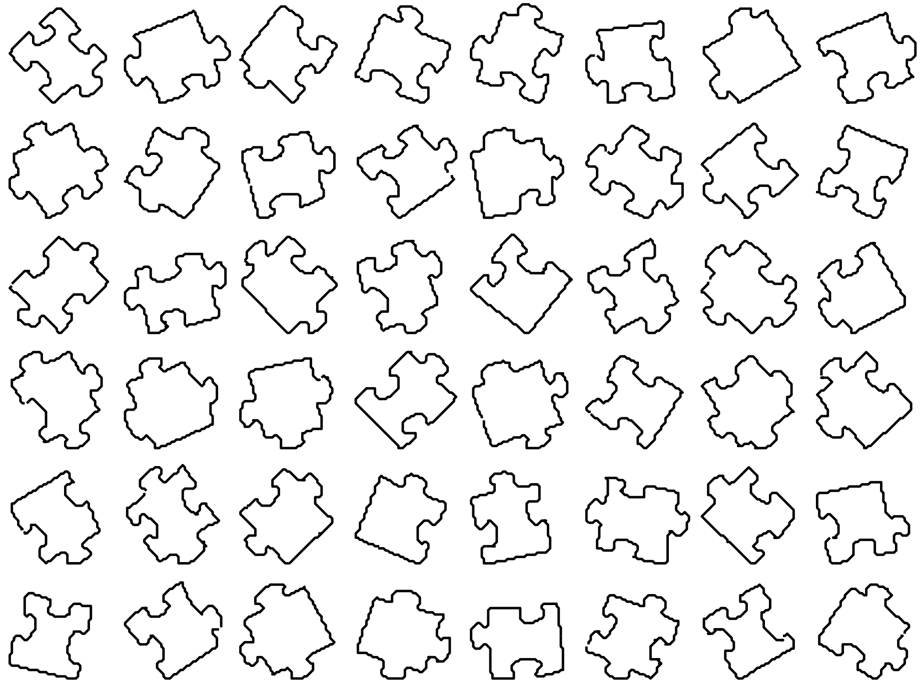

The Rain Forest Puzzle Solved

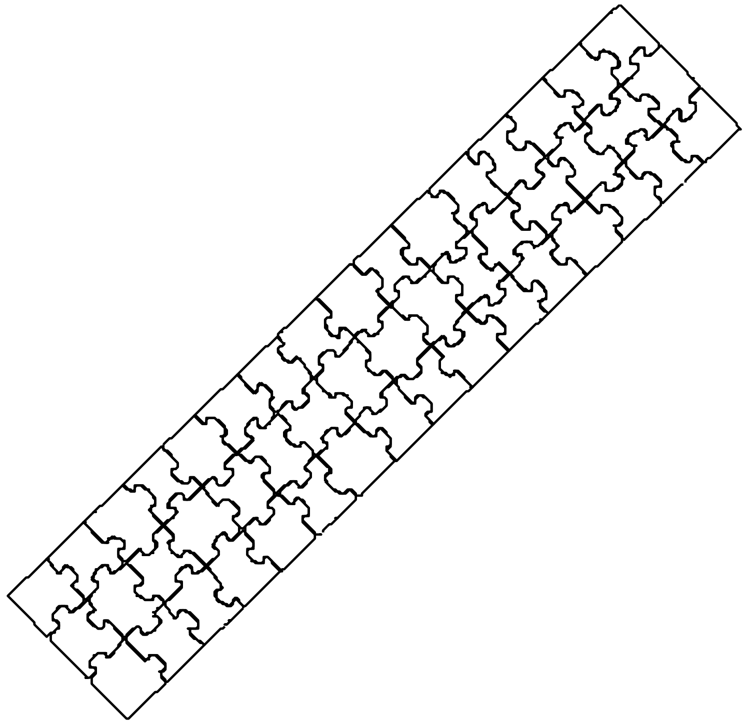

The Rain Forest Puzzle Solved

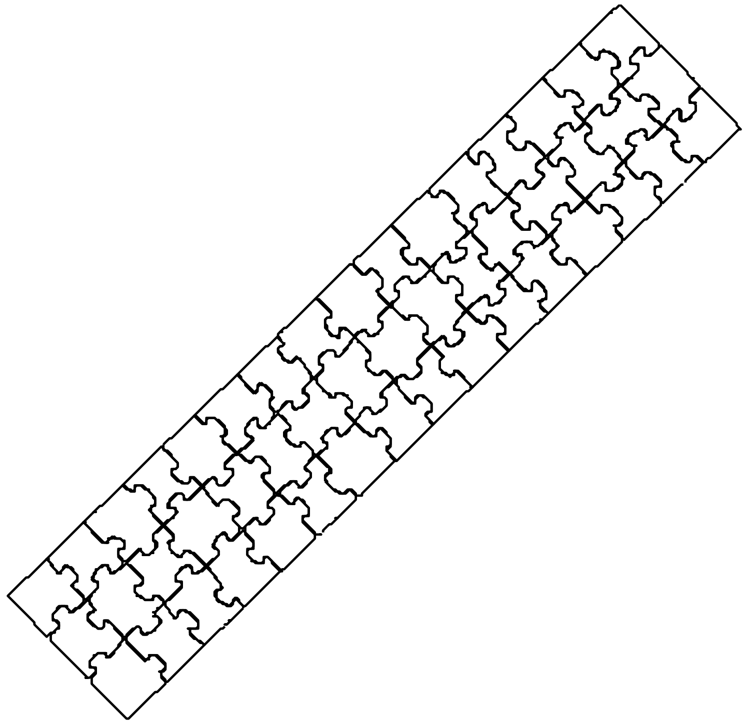

⟹ D. Hoff & PJO, Automatic solution of jigsaw puzzles, preprint, 2012.

The Distance Histogram

Definition. The distance histogram of a finite set of points $P = \{z_1, \ldots, z_n\} \subset V$ is the function

$$\eta_P(r) = \#\left\{ (i, j) \;\middle|\; 1 \le i < j \le n,\ d(z_i, z_j) = r \right\}.$$

The Distance Set

The support of the histogram function,

$$\operatorname{supp} \eta_P = \Delta_P \subset \mathbb{R}^+$$

is the distance set of P.

The Distance Set

The support of the histogram function,

$$\operatorname{supp} \eta_P = \Delta_P \subset \mathbb{R}^+$$

is the distance set of P.

Erdös' distinct distances conjecture (1946):

If $P \subset \mathbb{R}^m$, then $\# \Delta_P \geq c_{m,\varepsilon} (\# P)^{2/m-\varepsilon}$

Characterization of Point Sets

Note: If $\widetilde{P} = g \cdot P$ is obtained from $P \subset \mathbb{R}^m$ by a rigid motion $g \in \mathrm{E}(n)$, then they have the same distance histogram: $\eta_P = \eta_{\widetilde{P}}$.

Characterization of Point Sets

Note: If $\widetilde{P} = g \cdot P$ is obtained from $P \subset \mathbb{R}^m$ by a rigid motion $g \in \mathrm{E}(n)$, then they have the same distance histogram: $\eta_P = \eta_{\widetilde{P}}$.

Question: Can one uniquely characterize, up to rigid motion, a set of points $P\{z_1, \ldots, z_n\} \subset \mathbb{R}^m$ by its distance histogram?

$\Longrightarrow$ Tinkertoy problem.

Yes:

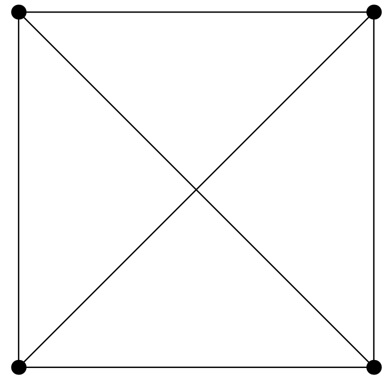

$$\eta = 1, \ 1, \ 1, \ 1, \ \sqrt{2}, \ \sqrt{2}.$$

No:

Kite

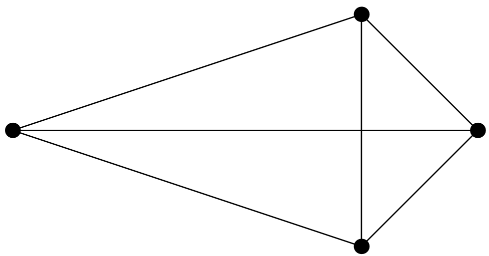

Trapezoid

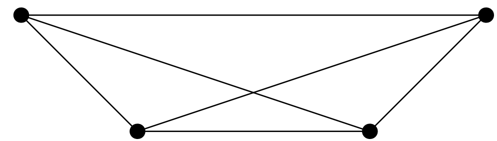

$$\eta = \sqrt{2}, \quad \sqrt{2}, \quad 2, \quad \sqrt{10}, \quad \sqrt{10}, \quad 4.$$

No:

$$\left.\begin{aligned} P &= \{0, 1, 4, 10, 12, 17\} \\ Q &= \{0, 1, 8, 11, 13, 17\} \end{aligned}\right. \subset \mathbb{R}$$

$$\eta = 1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 12, 13, 16, 17$$

$\Longrightarrow$ G. Bloom, *J. Comb. Theory, Ser. A* **22** (1977) 378–379

Characterizing Point Sets by their Distance Histograms

Theorem. Suppose $n \leq 3$ or $n \geq m+2$.
Then there is a Zariski dense open subset in the space of n point configurations in $\mathbb{R}^m$ that are uniquely characterized, up to rigid motion, by their distance histograms.

$\Longrightarrow$ M. Boutin & G. Kemper, *Adv. Appl. Math.* **32** (2004) 709–735

Limiting Curve Histogram

Limiting Curve Histogram

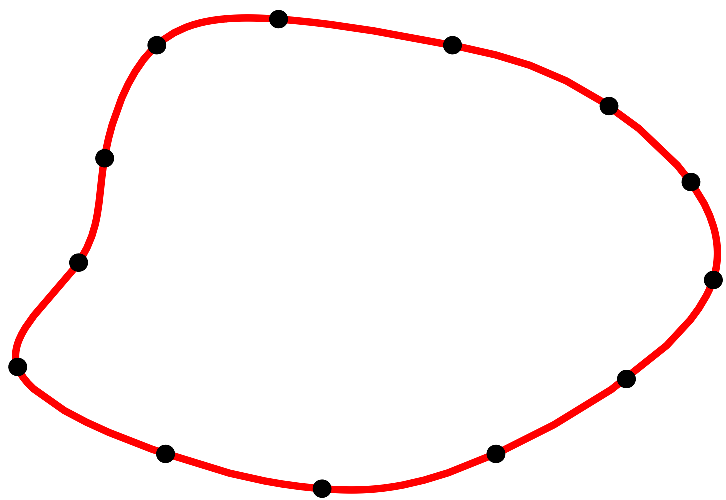

Limiting Curve Histogram

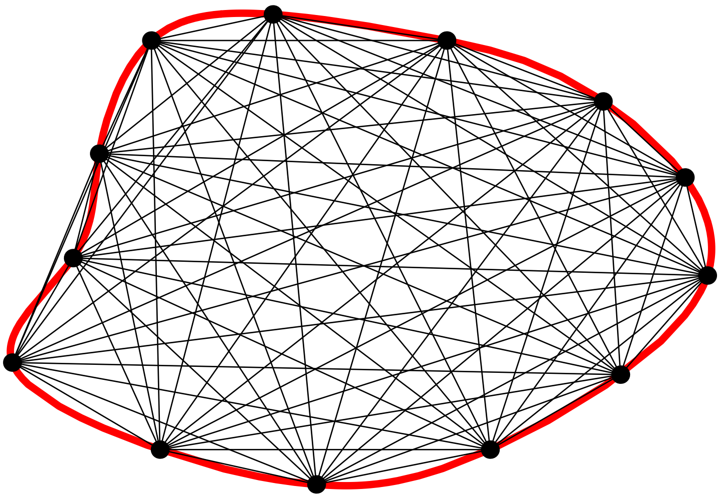

Sample Point Histograms

Cumulative distance histogram: $\quad n = \#P$:

$$\Lambda_P(r) = \frac{1}{n} + \frac{2}{n^2} \sum_{s \leq r} \eta_P(s) = \frac{1}{n^2} \# \left\{ (i,j) \;\middle|\; d(z_i, z_j) \leq r \right\},$$

Note:

$$\eta_P(r) = \tfrac{1}{2} n^2 [\, \Lambda_P(r) - \Lambda_P(r - \delta) \,] \qquad \delta \ll 1.$$

Local cumulative distance histogram:

$$\lambda_P(r, z) = \frac{1}{n} \# \left\{ j \;\middle|\; d(z, z_j) \leq r \right\} = \frac{1}{n} \#(P \cap B_r(z))$$

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

Ball of radius r centered at z:

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

Limiting Curve Histogram Functions

Length of a curve

$$l(C) = \int_C \, ds < \infty$$

Local curve distance histogram function

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

$\Longrightarrow$ The fraction of the curve contained in the ball of radius r centered at z.

Global curve distance histogram function:

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

Convergence of Histograms

Theorem. 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:

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

as the number of sample points goes to infinity.

D. Brinkman & PJO, Invariant histograms, *Amer. Math. Monthly* **118** (2011) 2–24.

Square Curve Histogram with Bounds

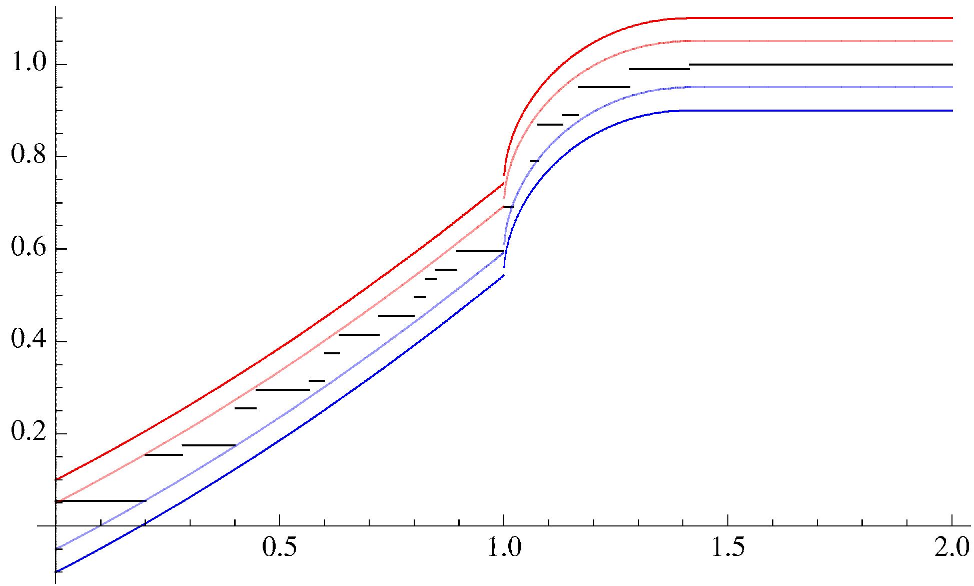

Kite and Trapezoid Curve Histograms

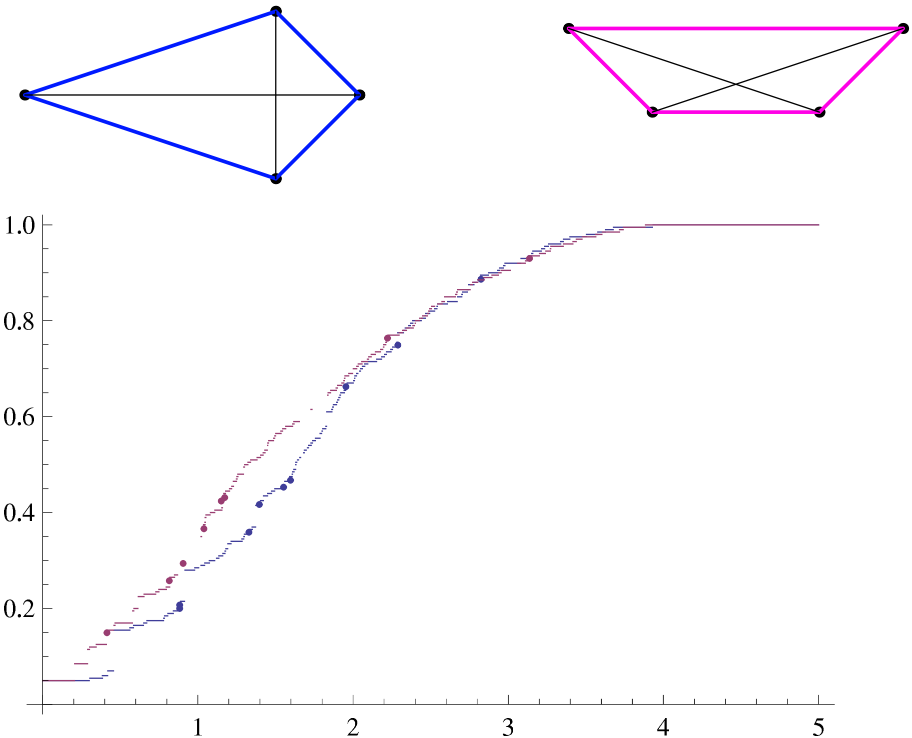

Histogram–Based Shape Recognition

500 sample points

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

Distinguishing Melanomas from Moles

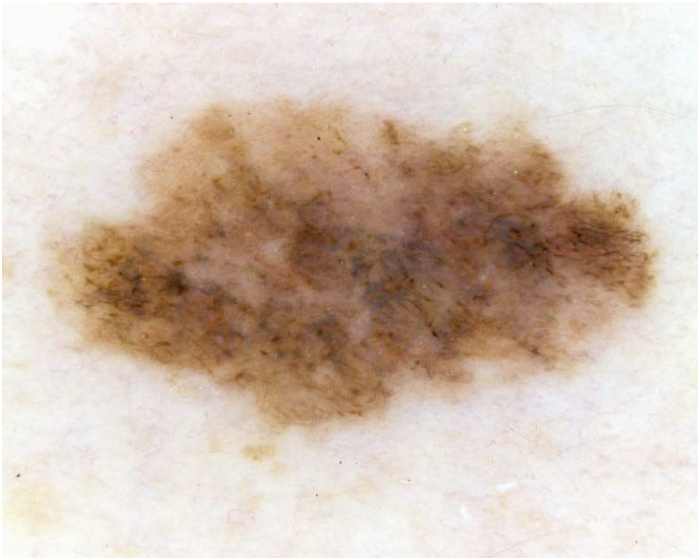

Melanoma

Mole

$\Longrightarrow$ A. Rodriguez, J. Stangl, C. Shakiban

Cumulative Global Histograms

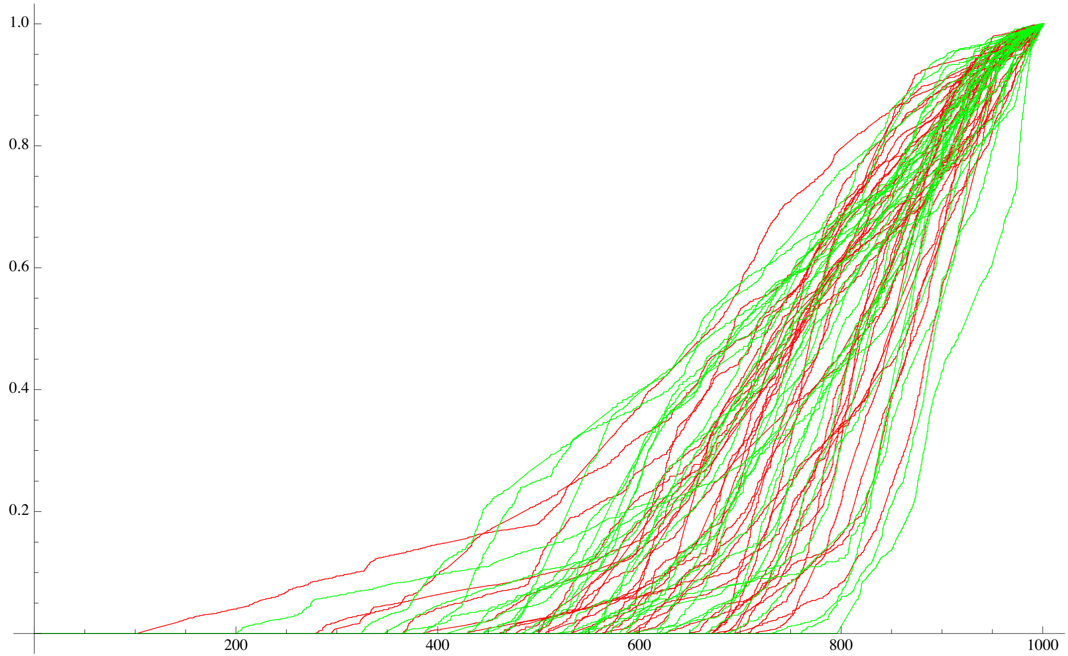

Red: melanoma Green: mole

Logistic Function Fitting

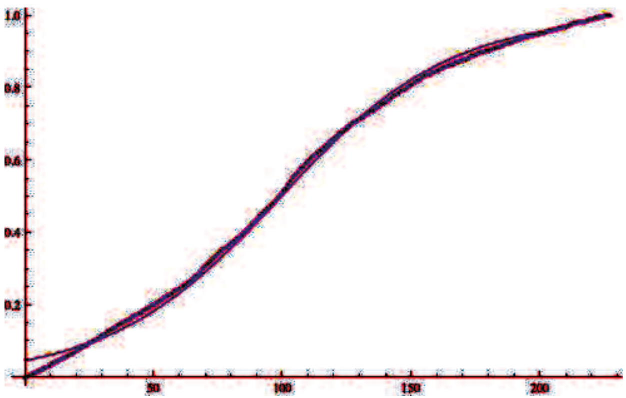

Melanoma

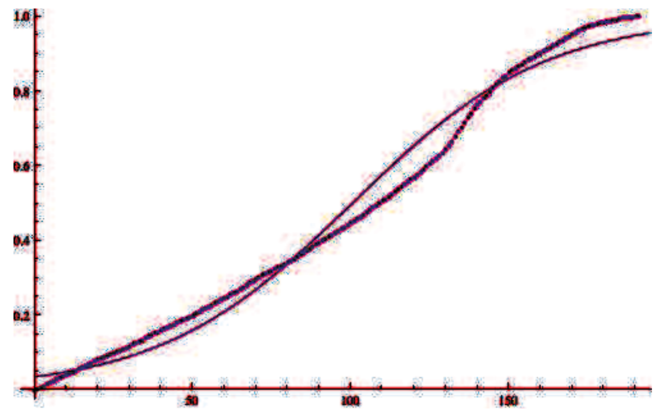

Mole

Logistic Function Fitting — Residuals

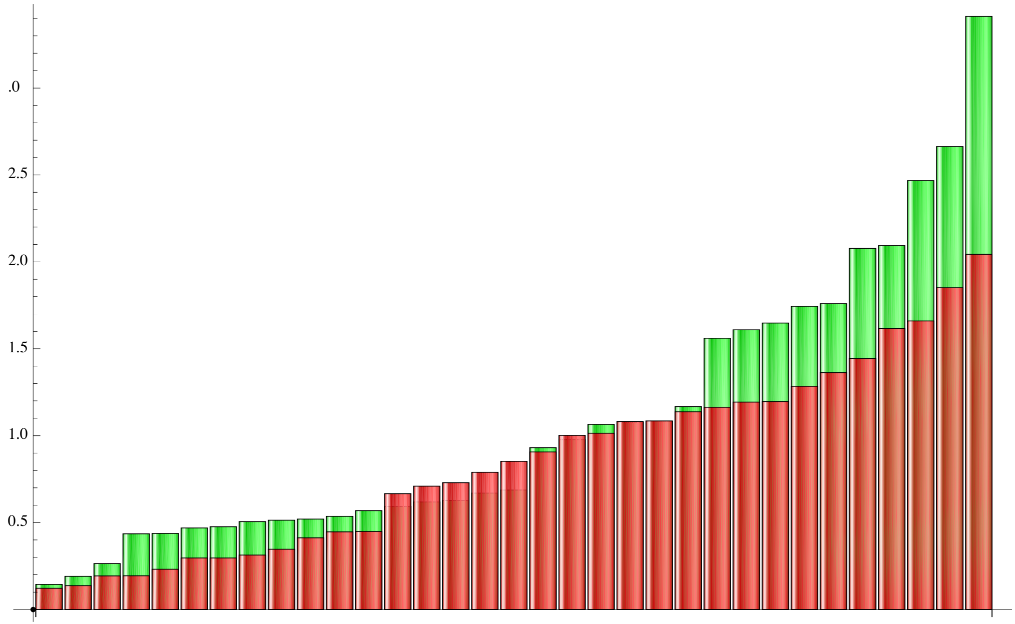

$$\left.\begin{array}{r} \text{Melanoma} = 17.1336 \pm 1.02253 \\ \\ \text{Mole} = 19.5819 \pm 1.42892 \end{array}\right\} \quad 58.7\% \text{ Confidence}$$

Curve Histogram Conjecture

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

Possible Proof Strategies

- Show that any polygon obtained from (densely) discretizing a curve does not lie in the Boutin–Kemper exceptional set.
- Polygons with obtuse angles: taking r small, one can recover (i) the set of angles and (ii) the shortest side length from $H_C(r)$. Further increasing r leads to further geometric information about the polygon ...
- Expand $H_C(r)$ in a Taylor series at $r = 0$ and show that the corresponding integral invariants characterize the curve.

Taylor Expansions

Local distance histogram function:

$$L\, h_C(r,z) = 2r + \tfrac{1}{12}\kappa^2 r^3 + \left(\tfrac{1}{40}\kappa\,\kappa_{ss} + \tfrac{1}{45}\kappa_s^2 + \tfrac{3}{320}\kappa^4\right) r^5 + \cdots \,.$$

Global distance histogram function:

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

Space Curves

Saddle curve:

$$z(t) = (\cos t, \sin t, \cos 2t), \qquad 0 \le t \le 2\pi.$$

Convergence of global curve distance histogram function:

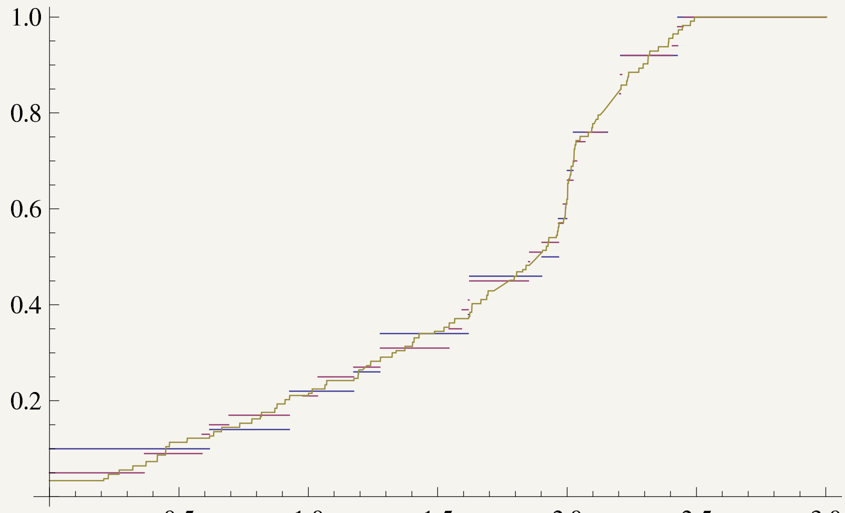

Surfaces

Local and global surface distance histogram functions:

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

Convergence for sphere:

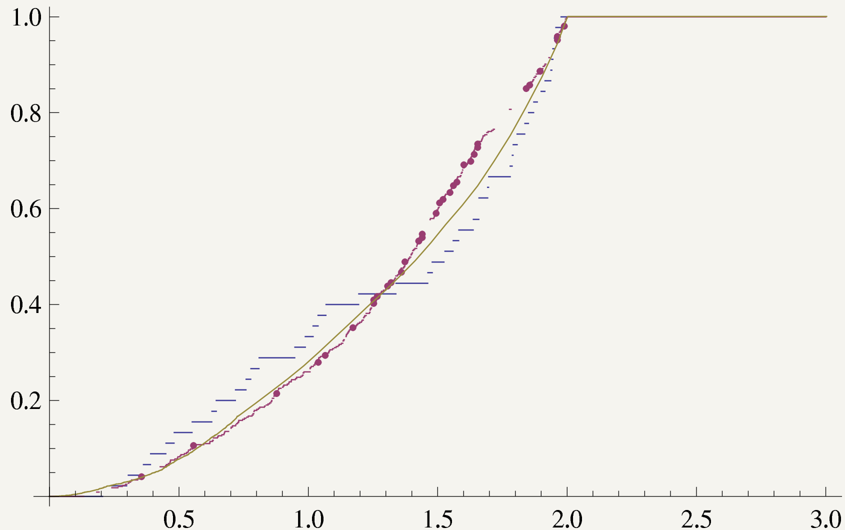

Area Histograms

Rewrite global curve distance histogram function:

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

$$\text{where} \qquad \chi_r(t) = \begin{cases} 1, & t \le r, \\ 0, & t > r, \end{cases}$$

Global curve area histogram function:

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

$$d\hat{s} \text{ — equi-affine arc length element} \qquad L = \int_C d\hat{s}$$

Discrete cumulative area histogram

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

Boutin & Kemper: The area histogram uniquely determines generic point sets $P \subset \mathbb{R}^2$ up to equi-affine motion.

Area Histogram for Circle

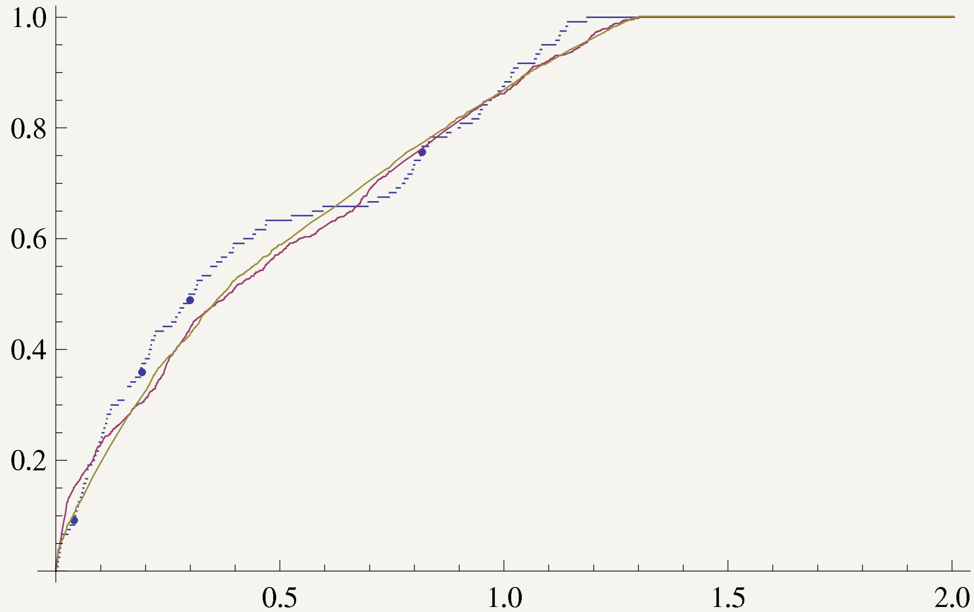

★★ Joint invariant histograms — convergence???

Triangle Distance Histograms

$Z = (\ldots z_i \ldots) \subset M$ —
sample points on a subset $M \subset \mathbb{R}^n$ (curve, surface, etc.)

$T_{i,j,k}$ — triangle with vertices z_i, z_j, z_k.

Side lengths:

$$\sigma(T_{i,j,k}) = (\, d(z_i, z_j), d(z_i, z_k), d(z_j, z_k)\,)$$

Discrete triangle histogram:

$$\mathcal{S} = \sigma(\mathcal{T}) \subset K$$

Triangle inequality cone:

$$K = \{\, (x, y, z) \mid\ x, y, z \geq 0,\ x + y \geq z,\ x + z \geq y,\ y + z \geq x \,\} \subset \mathbb{R}^3.$$

Triangle Histogram Distributions

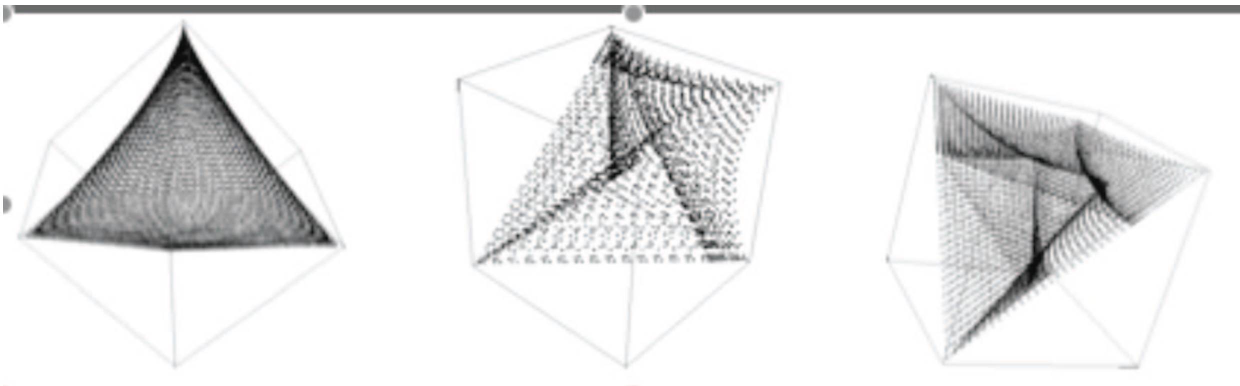

Circle Triangle Square

Convergence to measures . . .

$\Longrightarrow$ Madeleine Kotzagiannidis

Practical Object Recognition

- Scale-invariant feature transform (SIFT) (Lowe)
- Shape contexts (Belongie–Malik–Puzicha)
- Integral invariants (Krim, Kogan, Yezzi, Pottman, . . .)
- Shape distributions (Osada–Funkhouser–Chazelle–Dobkin)
  Surfaces: distances, angles, areas, volumes, etc.
- Gromov–Hausdorff and Gromov-Wasserstein distances (Mémoli)
  $\Longrightarrow$ lower bounds & stability