The shape of objects in two and three dimensions:

Mathematics meets Computer Vision

David Mumford
Brown University

Josiah Willard Gibbs Lecture
Baltimore, January 2003

Outline

1. Introduction: what does it mean to say two shapes are "similar"?

2. Three Riemannian metric methods
   a. Immersed curves in $\mathbb{R}^2$

   $$\mathrm{Diff}(S^1, \mathbb{R}^2) / \mathrm{Diff}(S^1)$$

   b. The full group of diffeomorphisms of $\mathbb{R}^n$

   $$\mathrm{Diff}(\mathbb{R}^n) / \mathrm{Diff}(\mathbb{R}^n, S^{n-1})$$

   c. The Teichmuller approach in $\mathbb{R}^2$

   $$\mathrm{Diff}(S^1) / PSL(2, \mathbb{R})$$

3. Discrete structure in the space of shapes: the medial axis cell decomposition

“Similarity” of shapes in humans and in computer vision

- Human beings have no trouble answering the question – do 2 objects have similar shapes?

- They use this to recognize the same object reappearing or to categorize objects into types (cars, tools, dogs, faces).

- Object recognition programs require measures of shape similarity to recognize, e.g. alpha-numeric characters, parts on an assembly line, faces at an airport.

- People create spontaneously an hierarchical classification of objects into broad categories, subcategories, continuing down to unique objects.

Some shapes and their categories

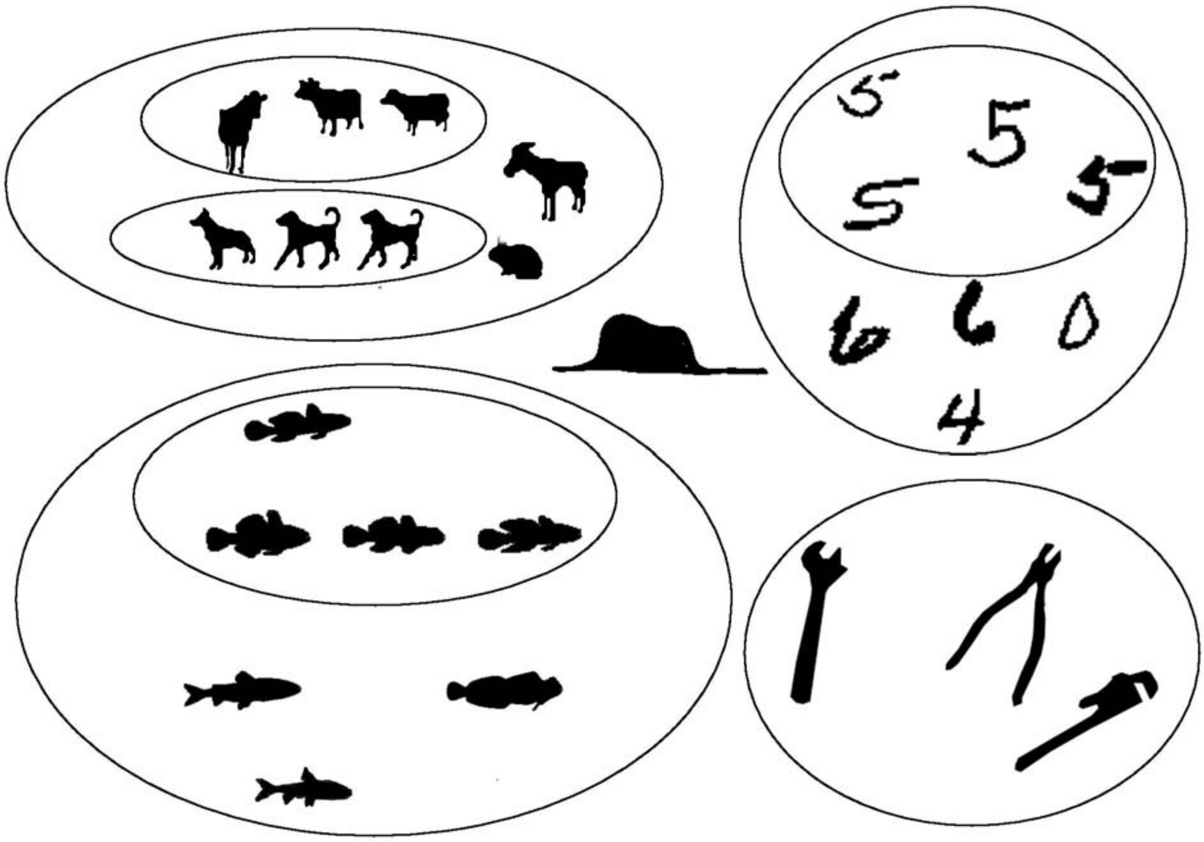

Typical shapes and examples of desired clustering in computer vision experiments. Top right: samples from the NIST handwritten zip code database often used in statistical learning theory; the ‘hat’ is Saint-Exupery’s pattern recognition challenge.

The task of the mathematician

• The natural idea is to define a metric space of "shapes", the distance being inversely proportional to psychophysical similarity.

• What point set? 2D retinal projections of objects or the full 3D object (or even 4D space-time traces of moving objects?

• We will call a ***shape*** any open subset $S \subset \mathbb{R}^n$ with not too convoluted a boundary ∂S, usually restricting to those S which are homeomorphic to a ball (sometimes as is, sometimes modulo Euclidean transformations, sometimes modulo translations and scalings). Usually we take $n = 2$ or 3.

• Call the set of these $\mathcal{S}$. As in Banach space theory, we will have a family of such $\mathcal{S}$ for varying degrees of regularity of ∂S.

An aside on human skills

• People are amazingly bad at remembering new 3D shapes!

• When faced with having to recognize ‘paper-clips’-- p/w-linear 3D curves -- from multiple views, people memorize the multiple 2D views (*Poggio, Bulthoff* et al).

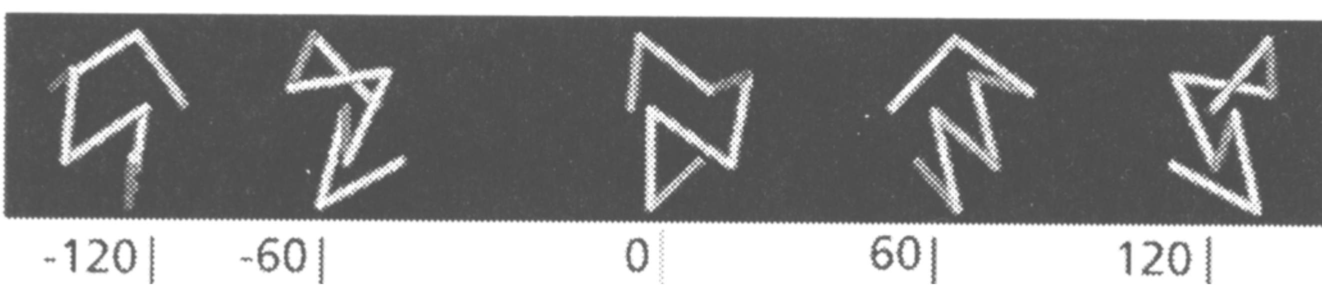

• People have a rich vocabulary and excellent memories for new 2D shapes, much more limited ones for 3D shapes. Perhaps our world is populated with rather special types of 3D shapes.

Shape recognition by man and machine is highly adaptive

- Thwarting the hope that one elegant metric models object recognition in general, we find that both (a) humans and (b) unsupervised object recognition programs tune themselves to differing aspects and features of shapes depending on the task.
- In 1988, *Richard Herrnstein, Steve Kosslyn* and I did a naïve experiment to see if humans and pigeons used a similar metric and whether we could model it mathematically. Pigeons were trained successively to peck for one shape and not any others: # errors define their internal metric. For humans, the inverse of reaction time defines a psychophysical metric.

15 polygons

The experimental stimuli: 15 polygons with varying 'features':

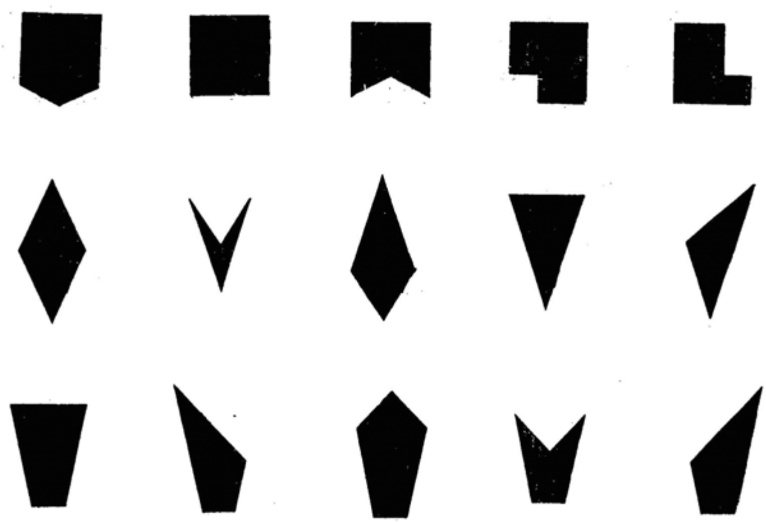

a) When pigeons were trained to discriminate these, their relative learning speeds could be modeled by a simple 2D plot seen here.

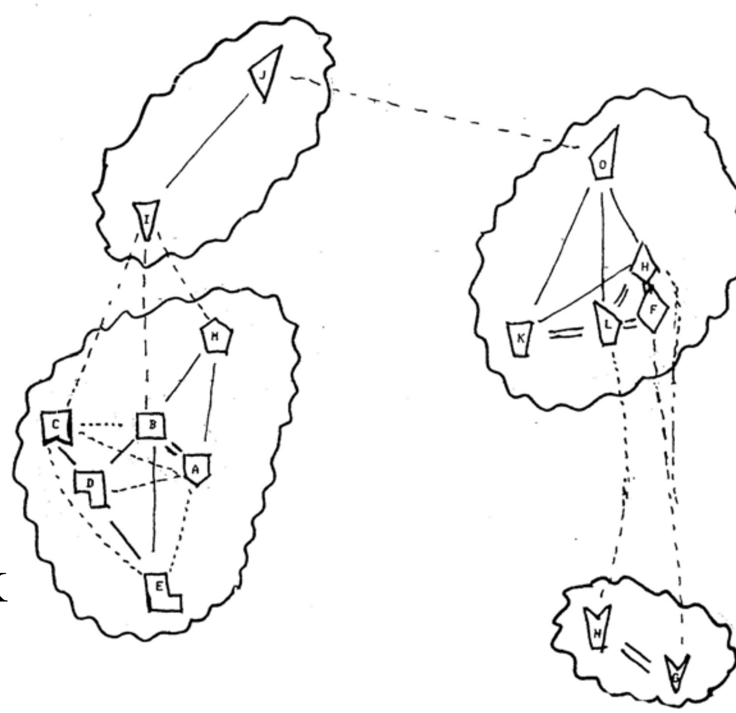

b) People's reaction times in the same task cannot be modeled so easily— because for each task, different features were attended to.

A first look at metrics on shapes $S \subset B$

***L^1*-metrics:**

$$d_1(S,T) = \text{area}(S \Delta T) = \| I_S - I_T \|_{L^1}$$

On $\mathcal{S} = \{S \mid S \text{ measurable, the current } \partial S \text{ is 1-rectifiable with } |\partial S| < \infty\}$.

d_1 is topologically equivalent to *flat metric:*

$$d_{flat}(S,T) = \inf_R \{|R| + |\partial(S - T + R|\}$$

and $\{S \mid |\partial S| \leqslant C\}$ is compact *(Federer-Fleming).*

***L^∞*-metrics:**

$$d_\infty(S,T) = \max\left(\sup_{x \in \partial S} d(x, \partial T),\ \sup_{y \in \partial T} d(y, \partial S) \right)$$

Closely related on $\mathcal{S} = \{\text{finite measures } \mu \text{ on } B\}$ is the *Prohorov metric*:

$$d_{pr}(\mu,\nu) = \sup_{x \in B} \inf_\varepsilon \{\varepsilon \mid |\mu * G_\varepsilon(x) - \nu * G_\varepsilon(x)| < \varepsilon\}$$

in which ball $\{\mu \text{ on } B,\ \|\mu\| \leqslant C\}$ is compact

Why more than one metric is needed

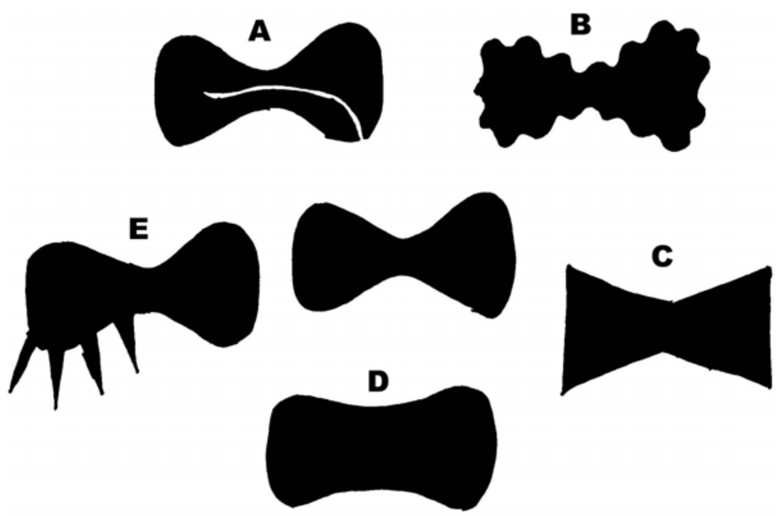

The central shape is similar in various respects to all 5 of the shapes around in – but in different metrics!

In L^1, distances are: A < B,C < D,E

In L^∞, distances are: B < C,D <A,E

In L^∞ with 1-jets: D < B,C < A,E

In L^1 with 2-jets: D < A,B < C,E

To make E close, need 'robust' non-convex metrics that discard outliers.

To make D far, qualitative ideas of 'parts' are needed – as it doesn't break into 2 parts.

D’Arcy Thompson’s idea: related shapes can be deformed to each other

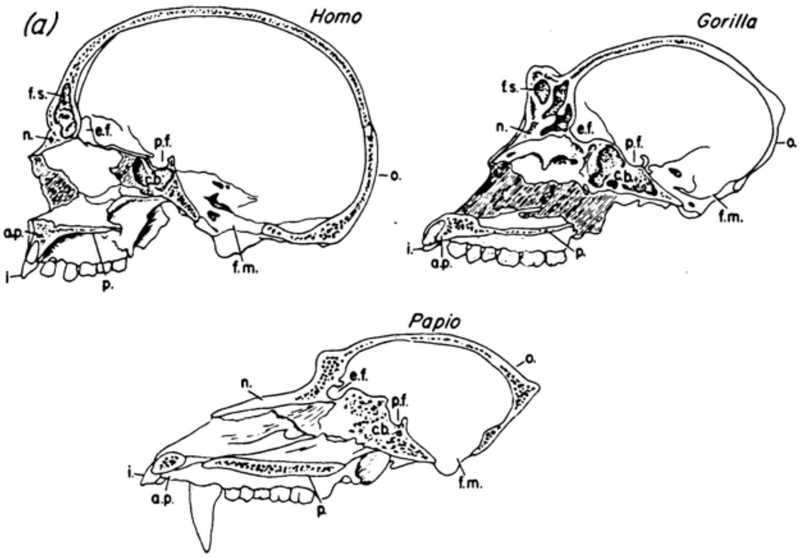

3 primate skulls, with ‘landmark’ points marked on each.

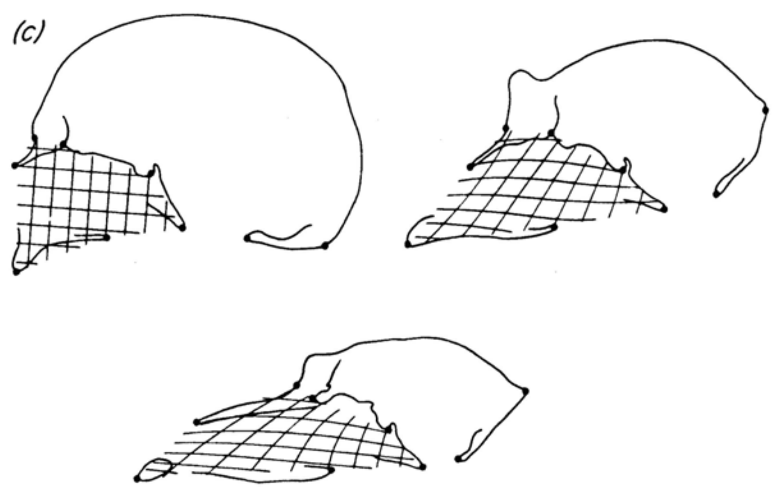

Diffeomorphisms of the human skull onto the two others given by biharmonic interpolation. (from *F. Bookstein*)

Mathematically, this suggests: define the metric as the length of the shortest path, in some Riemannian metric.

Geodesics between shapes and images

(Miller et al)

Make the set of shapes or images into a Riemannian manifold and solve for geodesics

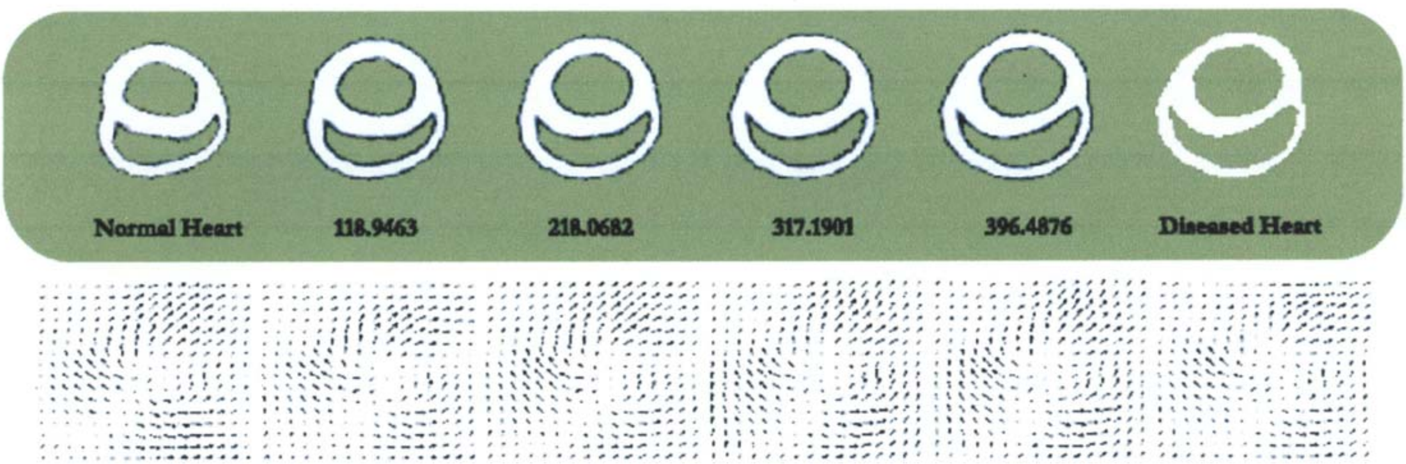

A geodesic between a normal and a diseased heart in 3D: top -- a 2D slice; bottom -- the vector field in the plane

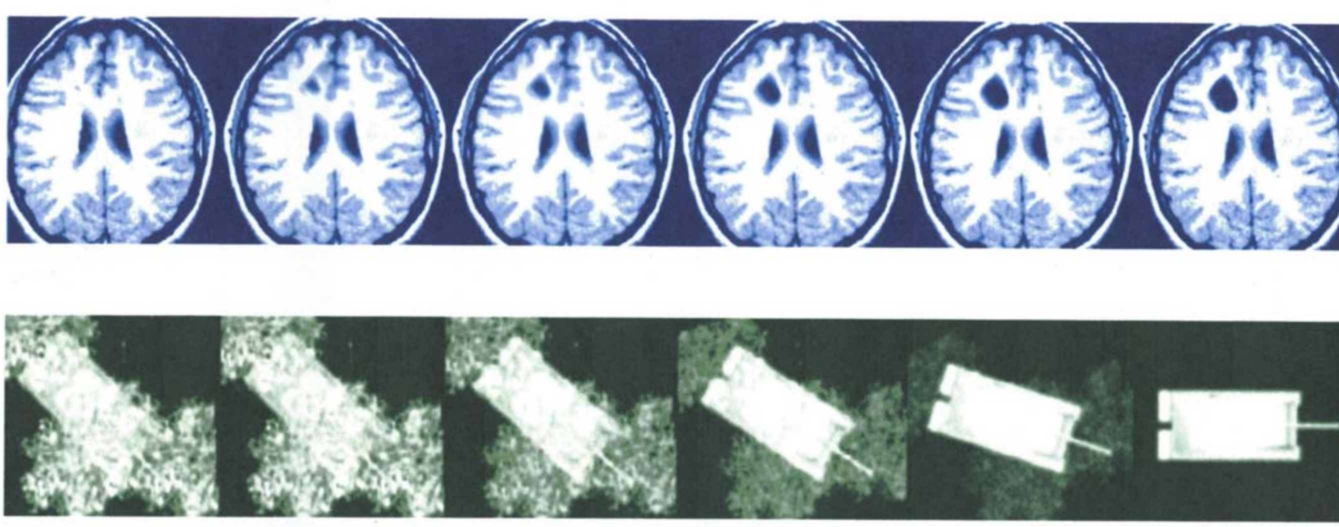

Top: a geodesic between a normal brain and one with a tumor; bottom: a geodesic between a concealed rotated tank and a normalized tank.

NEXT

After this introduction and motivation, let’s study the space of shapes as mathematicians and ask what tools we have for constructing and analyzing Riemannian metrics on these spaces.

1^{st} Riemannian metric: immersed curves (work w. P.Michor)

The nicest infinite dimensional spaces are Hilbert manifolds – but for the space of simple closed curves, things are not so nice!

Local charts: if C_0 has an arc length parametrization $\phi(s)$, let

$$\boxed{C_a = \text{locus}\left(\psi_a(s) = \phi(s) + a(s)\dot{\phi}^{\perp}(s)\right)}$$

where sup$|a|$ < C. BUT $\phi \in C^k \Rightarrow \psi_a \in C^{k-1}$!

Let $\mathcal{U} = \{f \,|\, 1 < |f'| < 2\} \subset C^k([0,1])$, then

$$f \mapsto f^{(-1)}, \mathcal{U} \to \mathcal{U}$$

has no Frechet derivative!

We must expect to start with a smaller space, e.g. the set of C^∞ curves $\mathcal{S}_\infty$, with tangent space C^∞ a's, and complete this to something weaker in a Riemannian metric.

The topology of the space of shapes $\mathcal{S}$

In dimension 2, there is a wonderful deformation retraction of $\mathcal{S}$.

Take the normal vector field *a(s)* to be curvature $\kappa(s)$ and set up the geometric heat equation using the above local chart:

$$\frac{\partial C_t}{\partial t}(s) = \kappa_{C_t}(s).\vec{n}_{C_t}(s)$$

Theorem of Gage-Hamilton-Grayson: This defines a flow on $\mathcal{S}$, carrying every C in finite time to an infinitesimal circle. Adding a pressure and drift term, every C approaches asymptotically the unit circle, hence $\mathcal{S}$ is contractible.

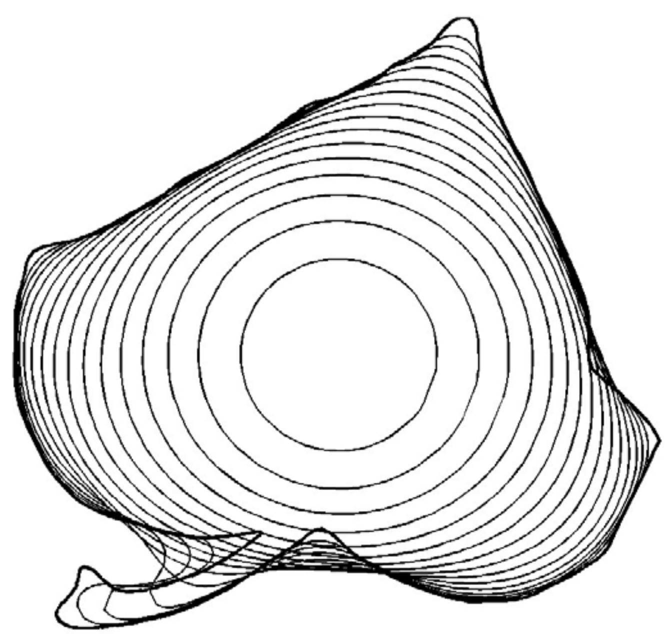

An abortive attempt

Using the charts $\{\psi_a\}$, define the metric by

$$\|a\|^2 = \int_{C_0} a(s)^2\, ds$$

With this definition, $\forall\ C_0, C_1$:

$$\inf(\text{length of path } C_t \text{ from } C_0 \text{ to } C_1)$$

$$= \int_0^1 \left(\int_{C_t} a_t^2\, ds_t \right)^{1/2} dt = 0 \text{ !!}$$

"Shark-skin" effect:

Teeth decrease normal velocity a by ε, increase arc length by ε^{-1}

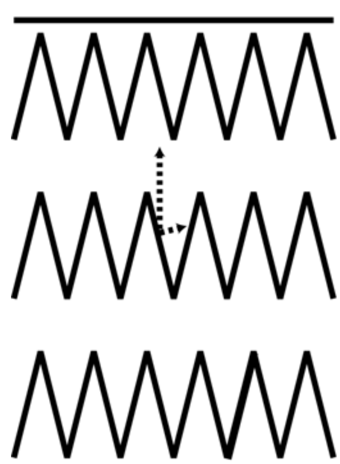

Good metric $d^{(\mathrm{I})}$:

$$\|a\|^2 = \int_{C_0} \left(1 + A\kappa_{C_0}^2(s)\right) a(s)^2\, ds$$

$$d^{(\mathrm{I})} = \inf(\text{path length})$$

Relation to currents

Lemma: $\sqrt{|C|}$ is Lipschitz on $\mathcal{S}_\infty$, cnst $1/(2\sqrt{A})$

Use $\frac{d}{dt}|C_t| = \int a(s)\kappa(s)ds$

Cor.: $|C|$ extends to a continuous function on the completion $\bar{\mathcal{S}}_\infty^{(\mathrm{I})}$ and:

$$d_1(C_0, C_1) \le \left(\sqrt{|C_0|} + d^{(\mathrm{I})}(C_0, C_1)\right) d^{(\mathrm{I})}(C_0, C_1)$$

The metric extends to all *immersions, C^∞* closed curves, not nec. simple, $\mathcal{S}^{\mathrm{imm}}$ and we get a continuous map of the completion:

$$\bar{S}_\infty^{\mathrm{imm}} \to (\text{closed integral currents, flat metric})$$

(or, equivalently, the space of integer-valued measurable fcns f s.t. ∂f is 1-rectifiable, endowed with the L^1-metric.)

Geodesics, curvature of this metric

Let $C_t =$ locus $\phi(s) + a_t(s)\dot{\phi}^{\perp}(s)$,
then $\{C_t\}$ is a geodesic iff:

$$\frac{\partial a_t}{\partial t} = \frac{\kappa a^2}{2} + A\left\{\left(\kappa'' - \kappa^3/2\right)a^2 + 4\kappa' a a' + 2\kappa a'^2\right\}$$

In the plane with orthonormal basis a_1, a_2, the sectional curvature is:

$$R = \int_{C_0} \frac{(1 - A\kappa^2)^2 + 4A^2(2\kappa'^2 - \kappa\kappa'')}{2(1 + A\kappa^2)} W^2 - \int_{C_0} AW'^2$$

$W = a_1 a_2' - a_1' a_2$ (Wronskian)

If A is small compared to W', κ'', sectional curvature is greater than or equal to 0!

Caricature of space $\mathcal{S}$

The origin of positive curvature

Distances shrink in the chart $\{\psi_a\}$:

$$d^{(\mathrm{I})}(C_{a+\varepsilon b}, C_a) \leq d^{(\mathrm{I})}(C_{\varepsilon b}, C_0), \text{ if } A \text{ is small}$$

Just like shark-skin collapse, normal displacement from C_a is less than from C_0 by a factor cos(θ)

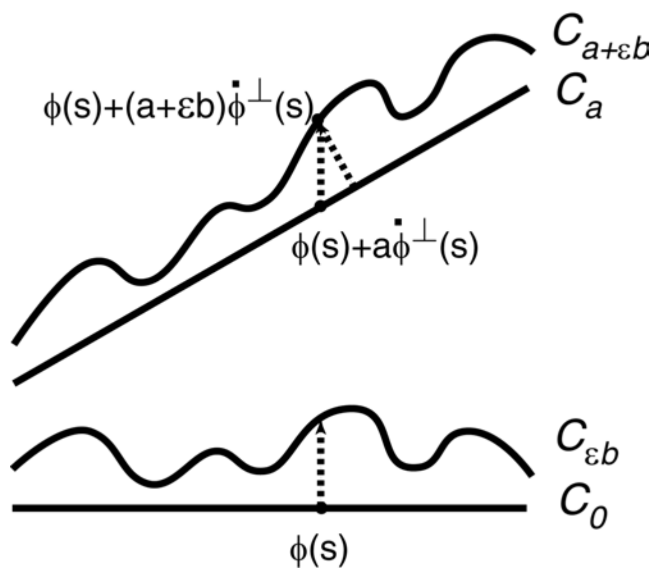

and:

$$\int_{C_a} b_a^2 ds_a = \int (b\cos(\theta))^2 . \frac{ds_0}{\cos(\theta)} \leq \int_{C_0} b(s)^2 ds_0$$

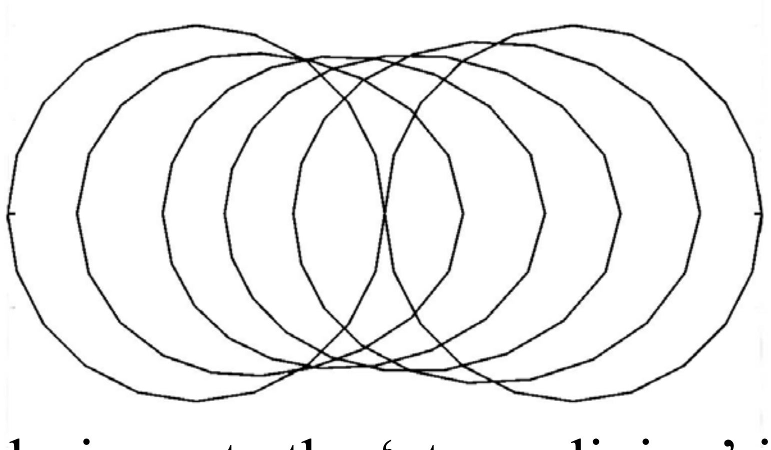

A geodesic: note the ‘streamlining’ in the middle

2^{nd} Riemannian metric: diffeomorphisms of $\mathbb{R}^n$

Work in any dimension n now,

Write the space of shapes $\mathcal{S}$ as a homogeneous space w.r.t. $\mathcal{G} = \mathrm{Diff}(\mathbb{R}^n)$:

$\mathcal{S} \cong \mathcal{G}$/(subgp $\mathcal{H}$ fixing unit sphere)

Put *right* invariant metric on $\mathcal{G}$, i.e.

$$\mathrm{dist}(\psi \circ \phi, \phi) = \mathrm{dist}(\psi, e), \text{ or}$$

$$\mathrm{dist}((I + \varepsilon \vec{v}) \circ \phi, \phi) = \|\vec{v}\|,$$

$$\|\vec{v}\| = \text{some norm in lie algebra}$$

$\mathcal{H}$ acts on right by isometries, so $\mathcal{G} \to \mathcal{S}$ is a Riemannian submersion, geodesics on $\mathcal{S}$ are geodesics on $\mathcal{G}$ starting, and hence continuing, $\perp$ to cosets $\phi\mathcal{H}$

Arnold's result

Introduce a Riemannian metric in
$\mathcal{SG} =$ group of volume preserving diffeos.

For any path $\{\theta_t\}$, let

$$\text{length of path} = \int \left(\sqrt{\int_{\mathbb{R}^k} \left\| v_t(\vec{x}) \right\|^2 d\vec{x}} \right) dt,$$

$$v_t(x) = \frac{\partial \theta_t}{\partial t} \left(\theta_t^{-1}(x) \right)$$

Then he proved geodesics are solutions of Euler's equation of incompressible inviscid fluid flow:

$$\boxed{\frac{\partial v_t}{\partial t} + (v_t . \nabla) v_t = \nabla p}$$

Metrics on the full group
(Christensen, Rabbitt & Miller)

On the full $\mathcal{G}$, need a stronger metric:

$$\|v_t\|_L^2 = \int_{\mathbb{R}^k} \langle Lv_t, v_t \rangle d\vec{x}, \quad L \text{ pos. self-adjoint,}$$

$$\text{e.g. } L = (I-\Delta)^m, \|v_t\|_m^2 = \int \sum_{|\alpha|\le m} \left\| D^\alpha v_t \right\|^2 dx$$

v_t = velocity, $u_t = Lv_t$ = 'momentum' in this metric.

Geodesics now are solutions to a *regularized compressible* form of Euler's equation (*Vishik*):

$$\boxed{\frac{\partial u_t}{\partial t} + (v_t.\nabla)u_t + \text{div}(v_t)u_t = -\sum_i (u_t)_i \vec{\nabla}((v_t)_i)}$$

Treating u as a section of $\Omega^1 \otimes \Omega^n$ (so $\langle u,v \rangle$ makes intrinsic sense), the equation says u is constant along the flow given by v.

The equation is linear in u, so u can be a generalized function!

Inducing a metric on the quotient $\mathcal{S}$

Now n=2. Define a quotient metric on $\mathcal{S}$ by:

$$\|a\|_L = \inf\left\{\|v\|_L \mid v \text{ vector field on } \mathbb{R}^2, a = (v.\vec{n}_C)\right\}$$

This metric is non-degenerate even for L=1- Δ because of:

Lemma: For L=1- Δ, $\sqrt{\text{area}(S)}$ is Lipschitz, and

$$d_{pr}(\mathbb{I}_S, \mathbb{I}_T) \leq C\sqrt{d_L(S,T)}$$

Cor.: The map $S \mapsto \mathbb{I}_S$ extends to a continuous map of the completion of $\mathcal{S}$ into the space of finite measures $\mathcal{M}(\mathbb{R}^2)$.

The infinitesimal metric, on the cotangent space to $\mathcal{S}$ (= 1-forms ω along C, 0 on t_C) is given by the Green's function K_L of L:

$$\|\omega\|_L^2 = \iint_{C\times C} K_L(s,t)\left\langle \omega(s).\omega(t)\right\rangle ds dt$$

On the tangent space, it is a pseudo-differential operator.

Geodesics in the quotient $\mathcal{S}$

For all curves C, define the singular 1-current ω_C:

$$\langle \omega_C, v \rangle = \int_C (v \cdot \vec{n}_C) ds_C$$

Then assume the momentum has the form $u_t = b_t \omega_{C_t}$ for some functions b_t on C_t. This gives geodesics on $\mathcal{S}$:

$$\frac{\partial C}{\partial t} = v \cdot \vec{n}_C, \quad v = K_L * b,$$

$$\frac{\partial b}{\partial t} + (v \cdot \vec{t}_C) \frac{\partial b}{\partial s} + \mathrm{div}(v) b = 0$$

($\partial b / \partial t$ by projecting nearby $C_{t'}$ normally to C_t.)

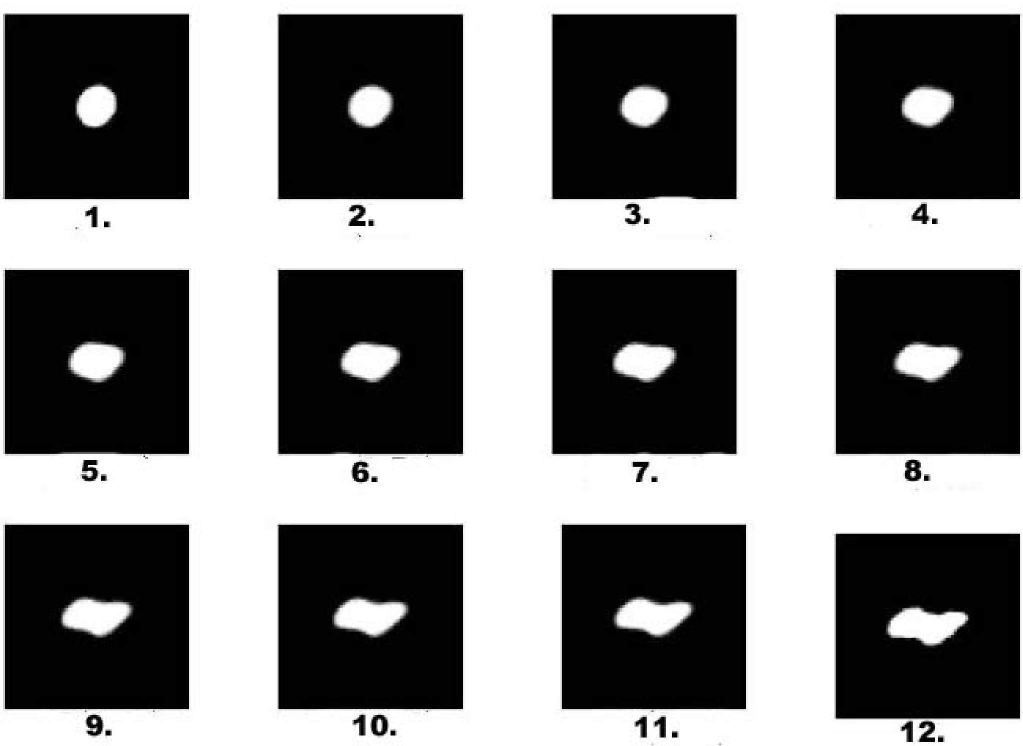

(*F.Beg*)

Diffusion defines a probability measure on $\mathcal{S}$ (Dupuis-Grenander-Miller, Yip)

Diffusion on $\mathcal{G}$ is a random path $\{\phi_t\}$ solving the SDE:

$$\frac{\partial \phi_t}{\partial t} = v(\phi_t(x), t), \quad v(x,t) \text{ Gaussian, cov} = K_L$$

Acting on the unit circle with a random stopping time, we get a measure. Here's a sample path:

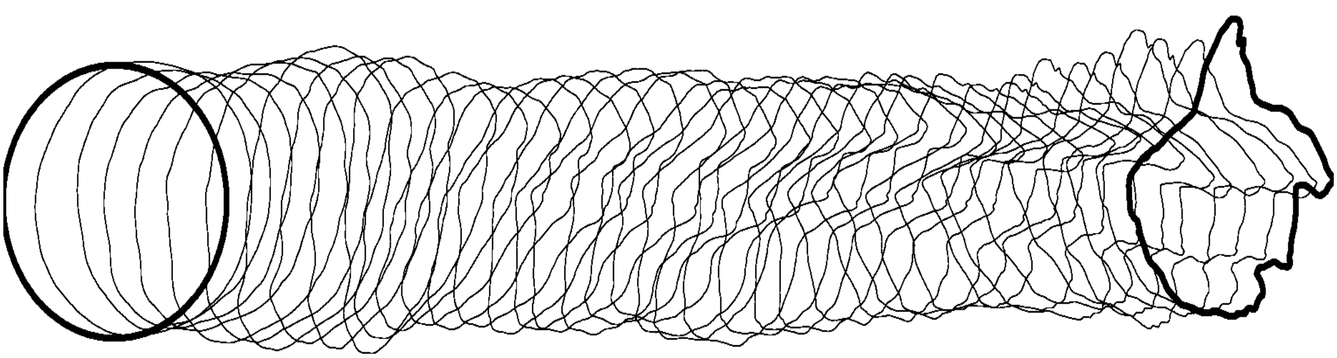

Alternately, we can combine this diffusion with curvature flow and seek the invariant distribution (as in Ornstein-Uhlenbeck motion).

The quotient space of 'landmark points'

(Kendall, Younes)

$\mathcal{G} = \mathrm{Diff}(\mathbb{R}^n)$ also acts transitively on the space of *distinct* m-tuples, $\mathcal{L} \subset (\mathbb{R}^n)^m$, giving a quotient metric:

$$d\left(\{P_i\},\{P_i+\varepsilon v_i\}\right)^2 = \varepsilon^2 \inf_{v(P_i)=v_i} \int \langle Lv, v \rangle$$
$$= \varepsilon^2 \sum_{i,j} G_{i,j} \langle v_i, v_j \rangle$$

where $G=K_L(\|P_i\text{-}P_j\|)^{-1}$. Its geodesics come from those on $\mathcal{G}$ whose momentum has finite support $\sum u_i \delta_{P_i}$. The geodesic equation is an ODE in which particles traveling in the same (resp. opposite) direction attract (resp. repel):

$$\frac{d\vec{P}_i}{dt} = 2\sum_j K_L\left(\|P_i - P_j\|\right)\vec{u}_j$$

$$\frac{d\vec{u}_i}{dt} = -\sum_j \vec{\nabla}_{P_i} K_L\left(\|P_i - P_j\|\right).(\vec{u}_i.\vec{u}_j)$$

Examples of warping to match landmark points (Miller, Younes)

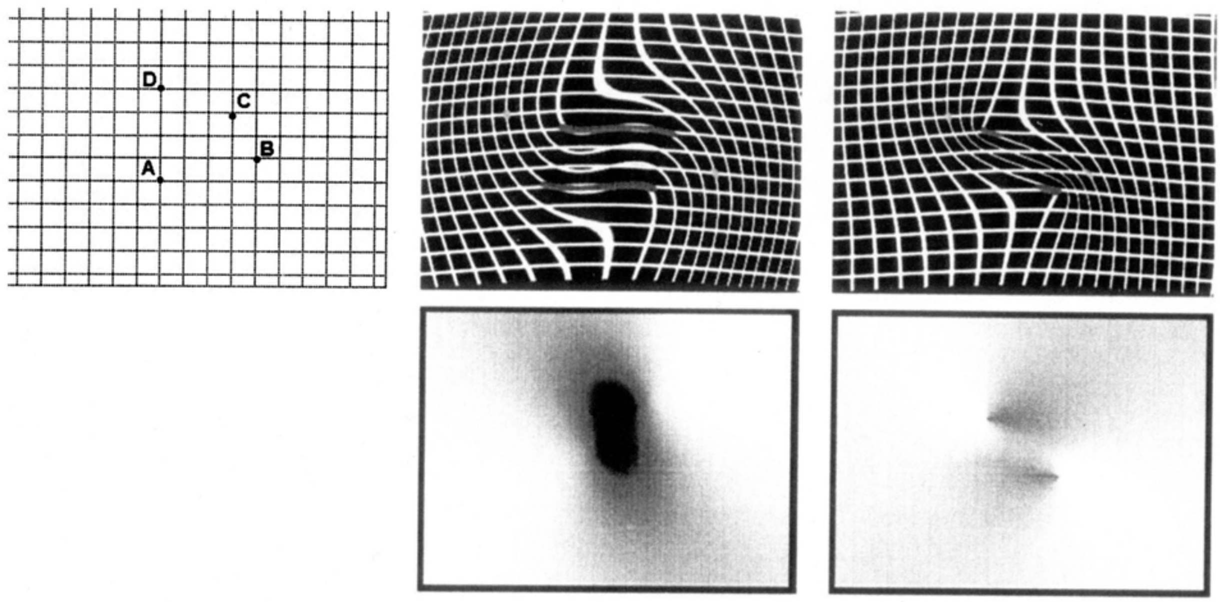

Point *A* is mapped to *B*, and *C* to *D*. On the left, the biharmonic interpolation, which fails to be a diffeomorphism. On the right, the diffeomorphism closest to the identity. On the 2nd row, the determinant (black=negative). Below, another example:

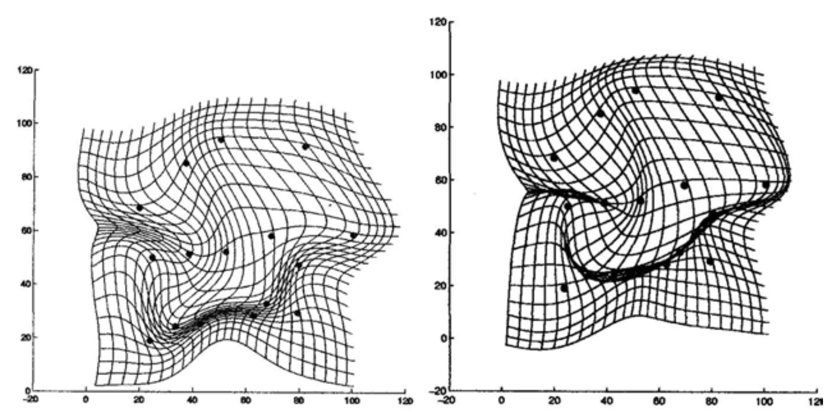

3rd Riemannian metric: the Teichmüller approach

Identify $\mathbb{R}^2$ with $\mathbb{C}$!

Use the Riemann Mapping theorem:

$\forall$ simple closed curves Γ,

$\exists$ conformal map $\Delta \xrightarrow{\phi_0} \mathrm{Int}(\Gamma)$

unique up to

$$\phi_0 \circ A,\ \ A(z) = \frac{\alpha z + \beta}{\bar{\beta} z + \bar{\alpha}},\ \text{with}\ |\alpha|^2 - |\beta|^2 = 1$$

If we embed $\Gamma \subset \mathbb{C} \subset \hat{\mathbb{C}}$, then

$\exists!$ conformal map

$\Delta \xrightarrow{\phi_\infty} \mathrm{Ext}(\Gamma) \cup \{\infty\}$

$\phi_\infty(z) = cz^{-1} + d_0 + d_1 z + \cdots,\ \ c$ pos. real

hence canonical curves

$$|\phi^{-1}(z)| = \text{cnst.},\ \arg \phi^{-1}(z) = \text{cnst.}$$

Thus $\psi = \phi_\infty^{-1} \circ \phi_0 : S^1 \to S^1$ is in $\mathrm{Diff}(S^1)/PSL_2(\mathbb{R})$

$$\boxed{S \longrightarrow \mathrm{Diff}(S^1)/PSL_2(\mathbb{R})}$$

The inverse of this construction is via "sewing"

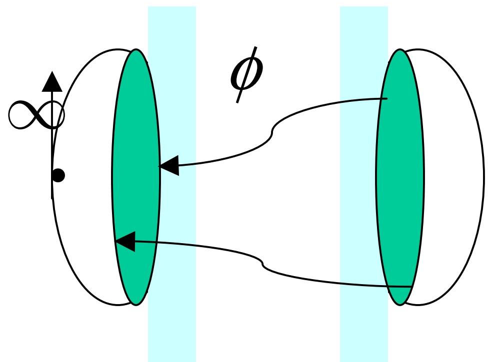

$\forall\phi$, construct abstract Riemann surface by gluing 2 copies of Δ via ϕ

The result is conformally equivalent to

$\hat{\mathbb{C}}$, with $(\infty, t_{\mathbb{R}^+}) \leftrightarrow (\infty, t_{\mathbb{R}^+})$

$$\bar{S} \underset{\text{def}}{=} S/(\text{transl+scalings}) \cong \text{Diff}(S^1)/PSL_2(\mathbb{R})$$

It has a complex structure and a Riemannian (in fact Kähler) metric both invariant under the left action of Diff(S^1).

The complex structure

Put $\overline{S}$ in a complex vector space via the coefficients of $\phi''_\infty / \phi'_\infty$

Identifying the tangent space to $\overline{S}$ at Γ with normal vector fields $a(s)n_\Gamma$ (mod constant/radial fields), the almost cx structure J is given by the Hilbert transform:

a) let θ be the angular coordinate given by ϕ_∞,
b) let $\mathcal{H}=\mathrm{ctn}(\theta/2)/2\pi$
then if $f(\theta)=\Sigma a_n e^{in\theta}$,

$$\mathcal{H} * f(\theta)=\Sigma\, i\, \mathrm{sgn}(n) a_n e^{in\theta},$$

$$\boxed{J(a\cdot\vec{n}_\Gamma)=(\mathcal{H} * a)\cdot\vec{n}_\Gamma}$$

Equivalently, the normal vector field (mod …) also extends uniquely to a holomorphic vector field X_e on ext(Γ), triple zero at ∞, tangential component $J(an_\Gamma)^\perp$.

The Weil-Peterssen metric

Start with the norm on the lie algebra of Diff(S^1):

$$\|v\|^2_{\mathrm{WP}} = \sum_{n\geq 2} (n^3 - n)\,|\,a_n\,|^2,$$

$$\text{where } v = \sum a_n e^{in\theta} \cdot \partial/\partial\theta,\ a_{-n} = \overline{a}_n$$

This is invariant under ad($SL_2(\mathbb{R})$), hence extends to $\overline{S}$ and defines a *homogeneous* Riemannian metric!

It is given at Γ by expanding X_e in terms of angular coordinate θ from int(Γ):

$$\|an_\Gamma\|^2_{\mathrm{WP}} = \sum_{n>1} (n^3 - n)\,|\,a_{-n}\,|^2, \text{ if}$$

$$X_e(\phi_0(e^{i\theta})) = \sum a_n e^{in\theta}\, \partial/\partial z$$

The integrated distance d_{WP} is positive. In fact, there is a continuous map

$$\text{WP-completion}(\overline{S}) \subset \begin{pmatrix} \text{quasi-circles } T(1) \\ \text{Teichmuller metric} \end{pmatrix}$$

I believe all the sectional curvatures of the W-P metric are non-negative: but this does not seem to be in print.

NEXT

So we have 3 metrics. Lots of things are still unknown:

– how about the curvature of the middle metric?

– can we understand the distortion the W-P metric uses relative to the more elementary ones to homogenize this space?

– can we relate the metrics, compare geodesics and find the internal structure of the space?

We go back to human shape perception for some clues.

People see shapes as having parts

There is a universal human tendency to think of shapes as breaking up into parts, imposing a kind of shape grammar:

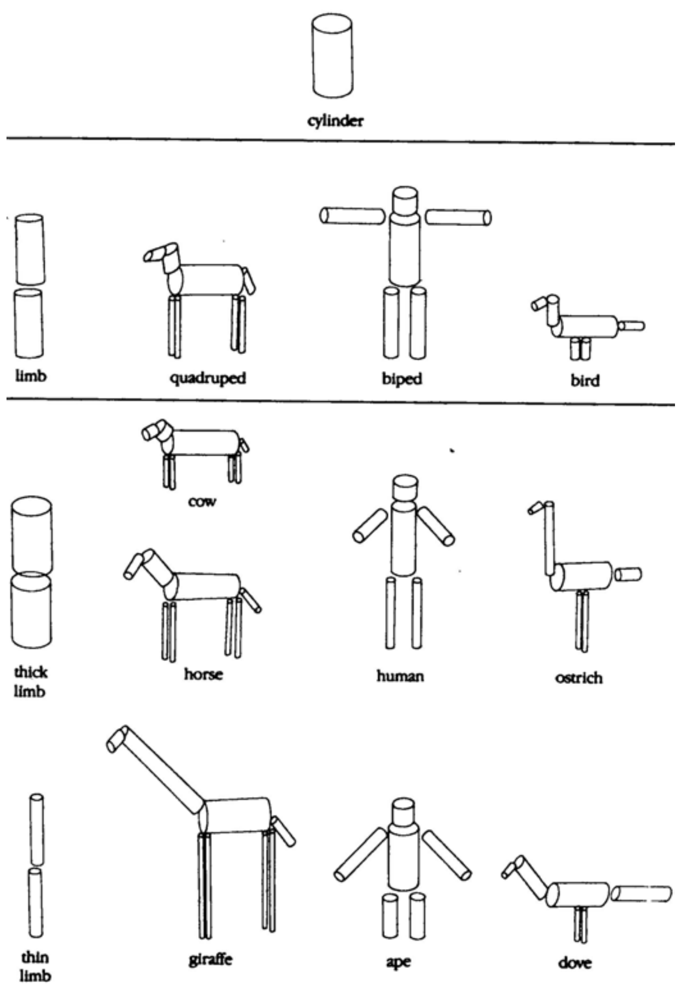

Marr's '3D model', in "Vision", (1980)

Is there some mathematics here?

The medial axis of a 2D shape

In dimension 2, the biologist *Blum* and the topologist *Thurston* invented the same construction to derive a combinatorial description of a shape *S*: equivalently, take the set of bitangent circles inside *S* or compactify $\mathbb{R}^2$ to S^2 and take the convex hull of *S* in D^3. The locus of centers of these circles is called the *medial axis*:

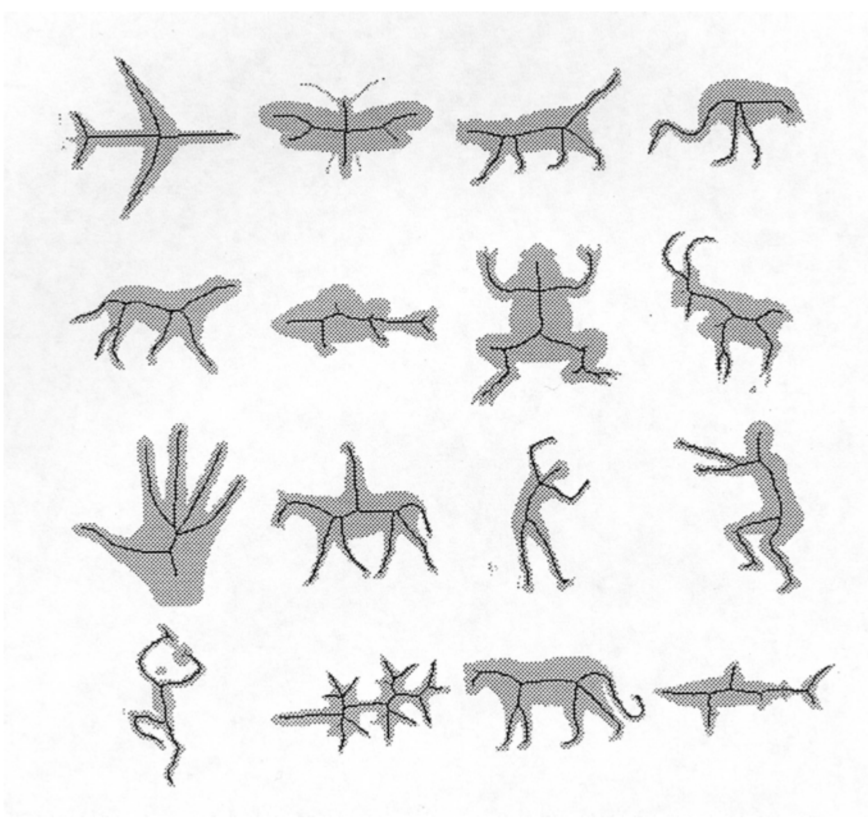

(Examples by *S.-C.Zhu*)

The medial axis seems to be computed in our brains

Psychophysical tests show extra sensitivity at the axis of a shape: the shape here is the cardioid, the arc inside its axis, the plots show contrast-sensitivity along the 2 dotted cross sections. (*I.Kovacs*)

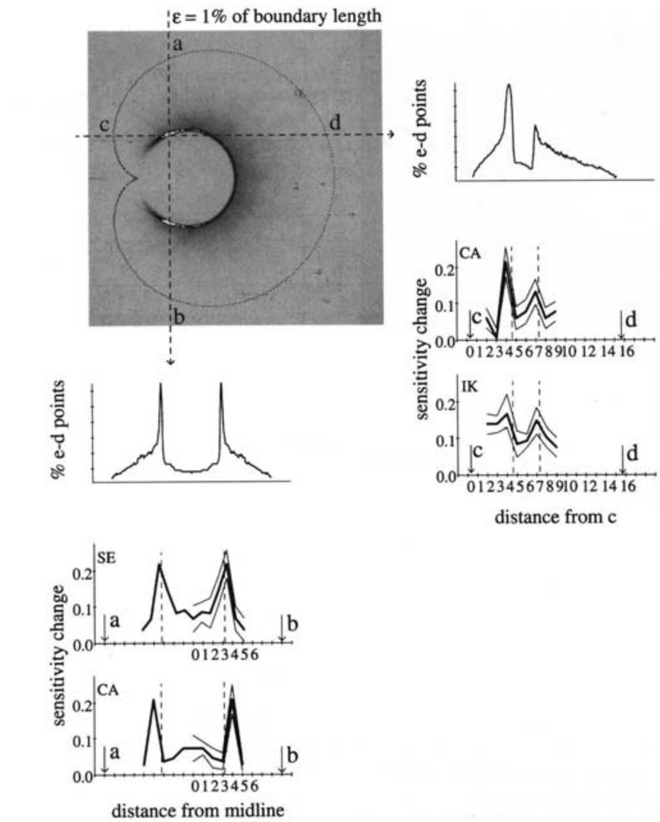

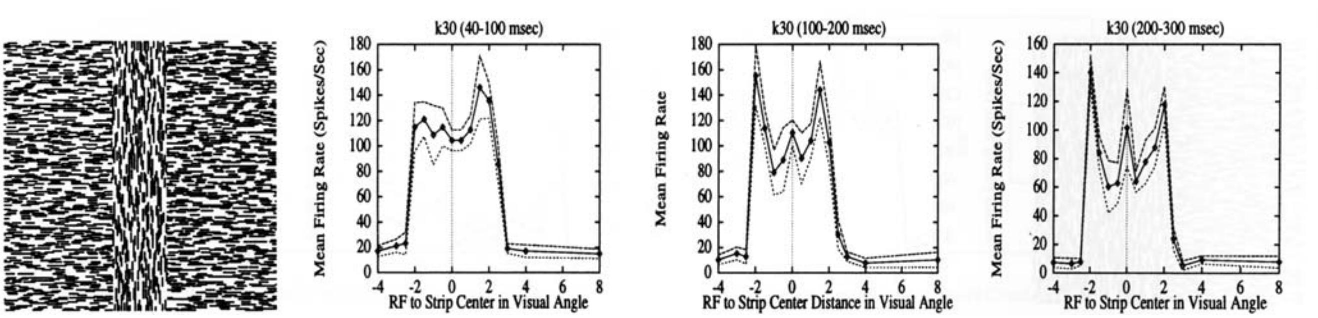

Left: the stimulus. Right: 3 plots of neural responses at varying time lags -- note response on the far right at the axis. (*Tai-Sing Lee*)

A cell decomposition of $\mathcal{S}$: work in progress

In an open dense subset $\mathcal{U}$ of $\mathcal{S}$, the medial axis has a finite number of non-degenerate singularities, either centers of tri-tangent circles or centers of circles osculating to order 4 at a local curvature maximum.

A 1st approximation defines the cells to be the connected components of $\mathcal{U}$, one for each type of tree. Better decompositions arise from a) pruning the axis when the angle between the bitangencies is small and b) adding further combinatorial structure from local minima of the disk radius – *necks* – or the external medial axis – *concavities* (*Kimia*).

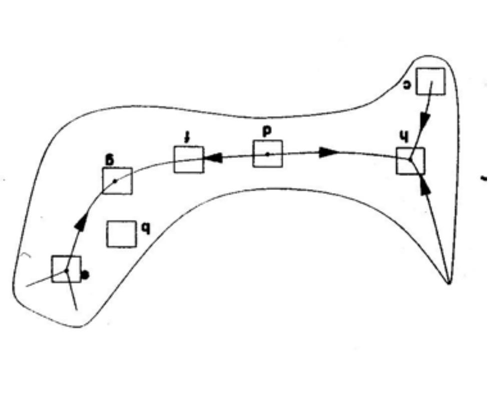

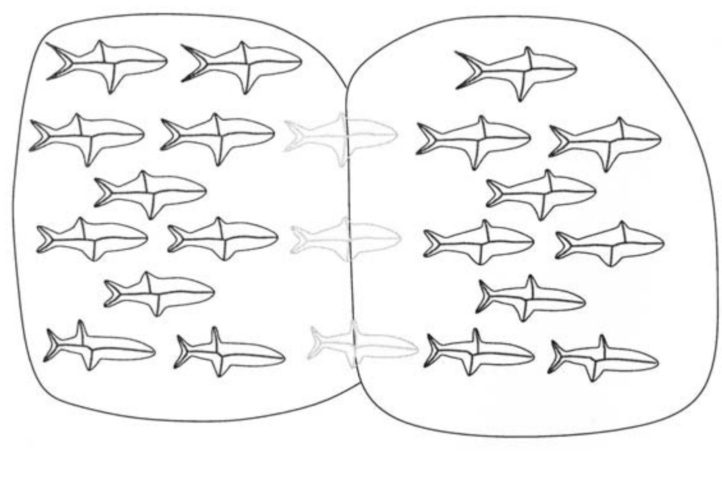

The medial axis can be used to derive better probability measures on $\mathcal{S}$

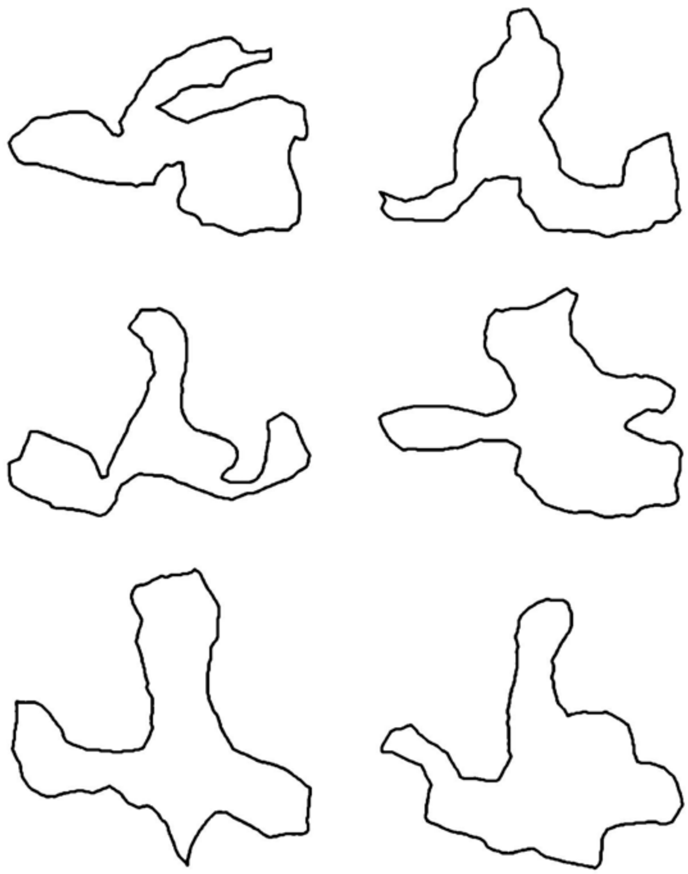

Samples from a probability model on polygons, of exponential type, trained to reproduce marginal distributions on 6 statistics related to curvature and the medial axis (*Song Chun Zhu*)

Dimension 3 is much harder

In dimension 3, $\mathcal{S}$ is still contractible (*Hatcher*) but no simple retraction is known. Mean curvature heat flow and other variants produce singularities. And in high enough dimension, $\mathcal{S}$ is *not* contractible!

In dimension 3, we typically break up objects into 'generalized cylinders', cylinders with irregular cross-sections, which may twist and bend (*Binford*). What is the best mathematical construction of such parts? What is the right grammar for 3D shapes?

Without arc-length parametrization, it is much harder to deal with surfaces, e.g. get local charts on $\mathcal{S}$, construct random shapes via SDE's or polyhedral models, etc.

Would you have guessed this in dimension 3?

In dimension 2, curvature max and min are perceptually obvious. In dimension 3, they generalize to the *ridge curves*, but the joker which makes surface geometry complex are the *umbilics*, which seem to be perceptually invisible!