Geometry, Imaging and Computing
Volume 1, Number 2, 173–221, 2014

Metric spaces of shapes and applications: compression, curve matching and low-dimensional representation

Matt Feiszli, Sergey Kushnarev, and Kathryn Leonard*

In this paper we present three metrics on classes of 2D shapes whose outlines are simple closed planar curves. The first, a C^1-type metric on classes of shapes with Lipschitz tangent angle, allows for estimates of massiveness such as ε-entropy. A Sobolev-type metric on piecewise C^2 curves allows for efficient curve matching based on a multiscale wavelet-like analysis. Finally, the Weil-Petersson metric, a Riemannian metric on the class of smooth diffeomorphisms of $S^1 \to \mathbb{R}^2$, allows a low dimensional shape representation, an N-Teichon, whose initial conditions are closely linked to curvature.

Dedication

This paper is dedicated to the 75th birthday of Prof. David Mumford. Prof. Mumford served as a PhD adviser for the authors of this paper, and the work contained here began as three of the last four theses he supervised. His unparalleled knowledge, unmatched scientific intuition and unabating interest in mathematics and applications shaped the authors' academic lives and the field of Pattern Theory itself. The authors are deeply grateful to him for his mentorship and intellectual generosity.

1. Introduction

From a mathematical perspective, spaces of embedded plane curves are among the simplest nonlinear infinite-dimensional spaces, offering the most tractable setting for explicitly computing results on Banach manifolds. From an applications perspective, shape data is incredibly complex, high-dimensional, and nonlinear.

One way to address that complexity is by identifying the data with linear function spaces. As we show in Sections 2 and 3, this approach can

*K. Leonard is partially supported by NSF IIS-0954256.

be quite effective in certain contexts. In order to do standard data analysis, however, the nonlinearity of the space must be accounted for, with shape data parameterized by a shape manifold. Even a simple statistic such as an average shape computed assuming linearity will be quite different from one computed in a way that respects the true underlying geometry. We describe one realization of a shape manifold in Section 4.

There have been many efforts in characterization and quantification of discrete shape spaces, for example, by Kendall [49], Bookstein [50], Dryden and Mardia [51], and Kent and Mardia [52], and of continuous shape spaces using, for example, level sets and active contours [54], or shock graphs [53]. Recently, an effort to understand the underlying geometry of shape space in both discrete and continuous settings has generated several successful frameworks including [15, 19, 23].

In this paper we are adopting a framework of Pattern Theory, in particular, a Riemannian geometric approach to the shape spaces. In this paradigm a shape is considered to be a point in some infinite dimensional manifold. Comparison is performed by the means of introducing a metric on this space and computing a geodesic. This offers a mathematically elegant and consistent framework. Deforming the shape along the geodesic demonstrates that shapes remain in the same class (for example, non-self intersecting C^1 curves) while transforming from one to another. This is an important consideration in some applications. For example, in computational anatomy, transforming the shape along the geodesic means that the template organ is morphing diffeomorphically into a target organ while remaining a plausible organ at every point along the geodesic. The metric measures the cost of this transformation. It will allow us to compute the Karcher mean, an average shape; linearize the shape space; and perform statistics. The pure Riemannian approach comes at a significant computational cost, however, and at the current state of numerical algorithms is not suitable when time is a concern.

The pure Riemannian framework is the most geometrically principled of the three presented here. By thoughtfully weakening the metric, however, the other two frameworks maintain the geometric philosophy of the Riemannian setting while increasing efficiency. The C^1-type in Section 2 metric replaces L^2 with L^∞ and retains tangent angle information in shape similarity. This allows for estimation of the complexity of shape classes, as well as of individual shapes. The $H^{1/2}$-type metric used in Section 3 can be viewed as a weakened linearization of the diffeomorphic approach. It makes similar guarantees about regularity, but exchanges the Riemannian framework for substantial gains in computational complexity for tasks such as matching.

Two of our metrics, the C^1- and Sobolev-type metrics, exploit powerful techniques for function spaces by working in geometrically important function classes associated to shapes. For the C^1-type metric, the function class is the collection of arclength parameterized tangent angle functions $\{\theta(s)\}$. In this setting we are able to obtain estimates on the complexity of shape space as well as a criterion for comparing the relative efficiency of curve representations. For the Sobolev-type metric, we take the class of tangent approximations $\{\beta(s,t)\}$, where $\beta(s,t) = \arg \frac{\gamma(s+1)-\gamma(s)}{\gamma(s)-\gamma(s-t)}$ for an arclength parametrized curve γ. This multiscale analysis provides a wavelet-like analysis of the curve, providing both theorems about regularity and practical algorithms for curve alignment. The third metric, the Riemannian Weil-Petersson metric, induces a hyperbolic geometry on shape space allowing for unique geodesics. We consider particular singular solutions to the geodesic equation, an N-Teichon, which possess tractable evolution equations. Estimates of the curvature of the shape's boundary in Section 4.5 allow us to construct an N-Teichon, that will provide essentially a linear, low-dimensional representation for a shape.

1.1. C^1-type metric

Section 2 introduces standard measures of massiveness: ε-entropy for compact classes of curves with Lipschitz tangent angle, and expected code length in an adaptive coding scheme for the non-compact setting. In the 1950's, Kolmogorov introduced the notion of ε-entropy as a measure of massiveness for sets in spaces where the unit ball was not compact [16]. Given a compact subset K of a metric space (X,ρ), with a minimal ε-covering containing N_ε elements, the ε-entropy of K is given by $\mathcal{H}_\varepsilon(K,\rho) = \log_2 N_\varepsilon$. The logarithm reflects the fact that the number of ε-balls required to cover K in an infinite-dimensional space is on the order of $C^{1/\varepsilon^\lambda}$, for constants C and λ depending on K. The notion of ε-entropy has been useful as an invariant of topological vector spaces [21], a way of measuring sizes of spaces of solutions to PDE's [12], and, in recent years, as a measure of efficiency in tasks of data compression [13, 14, 22]. We are most interested in this last application and its relationship to efficient coverings for recognition purposes. In general, optimal compression rates are obtained using a "true" probability distribution on the data, something we do not have due to the reasons mentioned above. Assuming a uniform probability on shapes, ε-entropy offers a kind of "best in the worst case," or minimax, compression rate.

The idea of adaptive coding arose in the data compression community in the early 1980's, introduced by Rissanen and Langdon in [17] as an alternative to arithmetic coding. Arithmetic coding compresses data using a

known, fixed statistical model, and has near-optimal compression rate when the data comes from the model used to compress it. To accommodate situations where the model is neither known nor fixed, Rissanen and Langdon developed a strategy to adapt the model to data (hence "adaptive" coding), continually updating the model as new data arrives. In this situation, the notion of optimality is not well-defined. Instead, Rissanen proposes a minimum description length (MDL) criterion for evaluating efficiency: the encoding resulting in the shortest expected code length is the most efficient [20].

Our main results in Section 2 are ε-entropy estimates for classes of curves with uniformly bounded arclength and Lipschitz tangent angle function. We obtain a tight estimate for the ε-entropy of general curves, and a slightly weaker one for closed curves, where the requirement of closure impedes the search for a lower bound. We also present the codelength for an adaptive encoding of curves that draws on the techniques developed for the ε-entropy estimates.

1.2. $H^{1/2}$-type metric

In Section 3 we recall some geometric variants of Sobolev spaces introduced in [2], where the authors applied the techniques (using an $H^{3/2}$-type metric) to the problem of denoising curves. Here we use the $H^{1/2}$ variant of the spaces and consider curve matching. The goal of these constructions is to provide a multiscale analysis of a curve which is, in most respects, like having a linear structure on the highly nonlinear space of curves. This representation encodes information similar to the "shape-tree" of Felzenszwalb [7] and "shape contexts" introduced in [1] by Belongie and Malik, while making the space of curves into a linear-like space in a simple and computationally convenient way.

Given a constant-speed curve $\gamma : [0, 2\pi) \to \mathbb{C}$, define the angles

$$\beta(\theta, t) = \arg \frac{\gamma(\theta + t) - \gamma(\theta)}{\gamma(\theta) - \gamma(\theta - t)}$$

Modulo translation, rotation, and scale, these angles provide a lift of the curve γ (a 1-dimensional object) into the strip $[0, 2\pi) \times [0, \infty)$ with coordinate (θ, t). The theorem here is that β angles behave essentially like the Haar wavelet coefficients of γ', and working with them is almost like working with wavelet coefficients. In particular, weighted L^2 norms of the coefficient energies over space and scale give information about the local Sobolev regularity of the curve when it is viewed as a collection of Lipschitz graphs. There

are several possible variants, all with slightly different behaviors in practice, but the same theorems apply in each case. For example, the complex second differences

$$\delta(\theta, t) = \gamma(\theta + t) + \gamma(\theta - t) - 2\gamma(\theta)$$

are the Haar coefficients of γ'. The complex second differences provide a rich and robust shape descriptor for the purpose of curve matching. This multiscale shape descriptor works well in applications, and the mathematics are simple and appealing. Further, it offers a nice way to understand what geometry is being measured when fractional order differential operators are applied to curve evolutions, as is done in the final section with the Weil-Petersson metric.

1.3. Weil-Petersson metric

In Section 4, we consider the Weil-Petersson (WP) metric on the coset space $\mathrm{PSL}_2(\mathbb{R})\backslash\mathbf{Diff}(S^1)$. This coset space (or its completion in the WP metric or in the Teichmüller topology) is known as the universal Teichmüller space and is well-known in many contexts: in the classification of Riemann surfaces [35], conformal and quasi-conformal maps [41], string theory [26] and most recently computer vision [46]. Its completion in the WP metric is an infinite dimensional homogeneous complex Kähler-Hilbert manifold [47].

As we will explain in Section 4.1 below, a particular dense subset of the universal Teichmüller space $T(1)$ is given by $\mathrm{PSL}_2(\mathbb{R})\backslash\mathbf{Diff}(S^1)$, where $\mathbf{Diff}(S^1)$ is the group of orientation preserving C^∞ diffeomorphisms of S^1, and $\mathrm{PSL}_2(\mathbb{R})$ is a subgroup of the Möbius selfmaps of the unit disk, see (8) and the surrounding discussion. This coset space is a Riemannian manifold for the WP metric and has another realization as the space of smooth simple closed curves modulo translations and scalings. We will use the terms 'shape,' 'diffeomorphism,' 'fingerprint,' or 'welding map' to refer to members in this dense subset of $T(1)$.

We have chosen Weil-Petersson metric for two main reasons. First, any two smooth shapes can be connected with a Weil-Petersson geodesic [30]. Second, all sectional curvatures of the metric are negative [47]. Thus geodesics connecting two shapes are unique [36].

Geodesic equations of groups of diffeomorphisms on a general Lie group G were first studied in Arnold's ground-breaking paper [24]. Arnold considered in particular the group of volume preserving diffeomorphisms of Euclidean space in its L^2 metric and found the geodesic equation for the

vector field $\vec{v}(\vec{x},t)$ to be Euler's fluid flow equation (see [25] for a full exposition). Other examples include the periodic Korteweg-deVries (KdV) equation and the periodic Camassa-Holm (C-H) equation [27]. These equations are geodesic equations on the Virasoro group, a central extension by S^1 of the group $\mathbf{Diff}(S^1)$ of the diffeomorphisms of S^1, for the L^2 and H^1 metric respectively. KdV and C-H are two completely integrable partial differential equations and have soliton solutions. Holm and collaborators have found that the geodesic equation on $\mathbf{Diff}(\mathbf{R}^n)$ admits special solutions with many of the properties of solitons: for each fixed time, they are diffeomorphisms which are largely localized in space and retain their general shape as they evolve; furthermore they interact somewhat like KdV solitons [29]. There are not, however, infinitely many conserved quantities so they are not true solitons.

The focus of this paper is a type of singular solution to the EPDiff, which we call Teichons. Singular solutions first arose as peakons (from 'peaked solitons') for a completely integrable Hamiltonian water wave equation, C-H in [27]. The peaks occurred where velocity profiles of the C-H equation had discontinuity in their slops and correspond to Dirac delta distributions of the associated momentum. The EPDiff equation for other metrics was later found independently in [48], and its singular solutions were shown to be important as landmarks in shape analysis [32, 44]. Later they were shown to comprise a singular momentum map for the right action of the diffeomorphisms on embeddings in any dimension [33]. Currently, the use of EPDiff and its landmark solutions is standard in shape analysis [34, 42, 43].

It turns out that considering the Weil-Petersson metric on the coset space $\mathrm{PSL}_2(\mathbb{R})\backslash\mathbf{Diff}(S^1)$ yields another example of a geodesic equation that is similar to KdV and C-H. This equation describing evolution of the velocity field $v(t,\theta)$ is

$$m_t + 2mv_\theta + vm_\theta = 0, \text{ where } m = -\mathcal{H}(v_\theta + v_{\theta\theta\theta}), \tag{1}$$

and $\mathcal{H}$ is the periodic Hilbert transform defined by convolution with $\frac{1}{2\pi}\mathrm{ctn}(\theta/2)$.

It is not known if (1) is completely integrable but it admits a class of soliton-like solutions which we consider in this paper: solutions in which m can be represented as a finite sum of weighted Dirac delta functions. Darryl Holm suggested the portmanteau *Teichons* to describe these soliton-like solutions on Teichmüller space and their corresponding geodesics. We adopt this terminology in this paper.

We use an N-Teichon ansatz (a sum of N Teichons) to reduce the integro-differential equation (1) to a finite-dimensional system of ordinary differen-

Table 1: Summary of properties of the three metrics discussed in this paper

Metric & Curve Class	Properties
C^1-type, formula (2) curves with Lipschitz tangent	– mainly used to determine shape complexity; – linearizes via geometrically relevant function space; – allows for MDL-type modeling comparisons
$H^{1/2}$-type, formulas (6),(7) C^0 curves	– mainly used for matching; – results on par with state of the art; – strong statements about regularity; – rich multiscale descriptor; – trades geodesics for fast computation – produces H^α measures for any α
$H^{3/2}$-type, formula (11) C^∞ curves	– mainly used for computing unique geodesics; – slow computation of geodesics; – allows for low-dimensional Teichon representation

tial equations. The main idea is then to represent any shape by the initial velocity that would send the unit circle to the shape via the geodesic equation. The initial velocity will be constructed from an N-Teichon, thus providing a low-dimensional representation for a given shape. Because of the uniqueness of the geodesic in the WP metric, this representation is well-posed. The algorithm that computes a Teichon geodesic between any two given shapes was described in [40]. Note that in [40] authors did not look for the optimal number of Teichons. In Section 4.5 we derive estimates that suggest a way to choose the Teichons to represent a particular shape.

Section 4 is organized as follows. Section 4.1 introduces the background on universal Teichmüller space $T(1)$, fingerprints (also called welding maps), and the Weil-Petersson metric. With the WP metric, Section 4.2 discusses the geodesic equation on the Teichmüller space (also known as EPDiff), and Section 4.3 discusses Teichon solutions of EPDiff. In Section 4.5 we derive an estimate of the curvature of the shape produced by an N-Teichon.

2. Measures of massiveness with a C^1-type metric

Partially motivated by applications to compression, a subset of the mathematical community has been interested in questions about efficient ε-approximation in Banach spaces. We present the earliest such results for curves, first presented in [18] without complete proof. A lesson from results on Banach spaces is that when ε-balls are rectangular instead of round, they stack more efficiently and therefore give a clear sense of how large an infinite-dimensional metric space might be. The Hausdorff metric, the typical L^∞-type metric on plane curves, is the simplest metric inducing rectangular

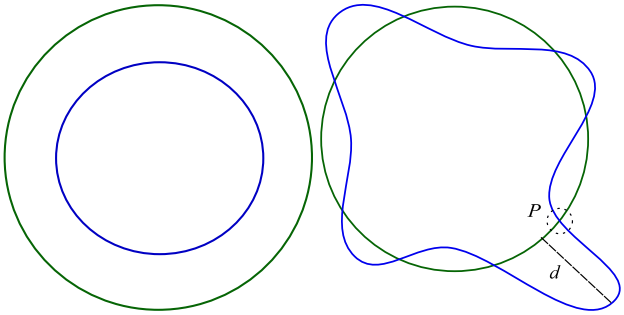

Figure 1: Left: These concentric circles have Hausdorff and ρ_1 distances both equal to the difference in radii of the circles. Right: The Hausdorff and ρ_1 distances may be different for these circles, depending on the value of d. If $d \geq \pi/2$, ρ_1 will equal the Hausdorff distance. If $d < \pi/2$, the Hausdorff distance will equal d while $\rho_1 = \pi/2$, the difference of tangent angles at the circled intersection point P.

ε-balls. For curves, however, orientation is often viewed as more relevant to similarity than location. We therefore introduce a C^1-type metric on curves. See Figure 1.

Consider $\mathcal{E}$, the set of all embeddings of the unit circle into $\mathbb{R}^2$ passing through the origin with horizontal tangent angle. Among these embeddings are curves with Lipschitz tangent angle, denoted $\mathcal{E}^1$. With analogous notation for immersions, we obtain classes $\mathcal{I}$ and $\mathcal{I}^1$. For curves that are not necessarily closed, we have $\mathcal{C}$ and $\mathcal{C}^1$. Naturally, $\mathcal{E}^1 \subset \mathcal{I}^1 \subset \mathcal{C}^1$.

Let $\gamma_1, \gamma_2 \in \mathcal{C}^1$. For $\mathrm{p} \in \gamma_i$, $\theta(\mathrm{p})$ will indicate the tangent angle to γ_i at the point p. A C^1-type metric is quite natural:

$$\rho_1(\gamma_1, \gamma_2) = \max_{\substack{i=1,2 \\ i \neq j}} \left(\max_{\mathrm{q} \in \gamma_i} \min_{\mathrm{p} \in \gamma_j} \left(|\mathrm{p} - \mathrm{q}| + \lambda \, |\theta(\mathrm{p}) - \theta(\mathrm{q})| \right) \right), \tag{2}$$

where λ is a dimension-normalizing constant.

2.1. Totally bounded spaces

Restricting to plane curves with length bounded by L and fixed Lipschitz constant K, we obtain totally bounded spaces $\mathcal{E}_{K,L} \subset \mathcal{I}_{K,L} \subset \mathcal{C}_{K,L}$ (these last two are compact). We can then measure the massiveness of these spaces by counting the minimum number of ε-balls in an ε-covering, taking $\lambda > 1/K$

in Equation 2. Taking logarithms of that minimum number yields the ε-**entropy** of the space [16].

Definition 1. *Let N_ε be the cardinality of a minimal ε-covering for a totally bounded subset $\mathcal{X}$ of a metric space $(\mathcal{M}, \rho)$. The* ε**-entropy** *of $(\mathcal{X}, \rho)$ is then:*

$$\mathcal{H}_\varepsilon(\mathcal{X}, \rho) = \log_2 N_\varepsilon.$$

Our ε-entropy results for curves will draw on the following two results for functions, stated here without proof.

Theorem 1. *[16] For $I = [a, b]$, $C > 0$, define:*

$$\mathcal{F}(C) = \left\{ f : I \to \mathbb{R} \mid f(a) = 0,\ |f(x) - f(x')| \leq C|x - x'|,\ \forall x, x' \in I \right\},$$

and $\rho_\infty(f, g) = \sup_{x \in I} |f(x) - g(x)|$. Then:

$$\mathcal{H}_\varepsilon(\mathcal{F}(C), \rho_\infty) = \begin{cases} \frac{|b-a|C}{\varepsilon} - 1 & \frac{C|b-a|}{\varepsilon} \in \mathbb{Z}^+ \\ \left[\frac{C|b-a|}{\varepsilon}\right] & else. \end{cases}$$

Theorem 2. *[18] Let $I = [a, b]$, and define:*

$$\mathcal{F}_0(C) = \left\{ f \in \mathcal{F}(C) \mid f(b) = 0, \int_I f = 0 \right\}.$$

There exists a 2ε-separated set in $(\mathcal{F}_0, L^\infty)$ with $m_{2\varepsilon}$ elements, where $m_{2\varepsilon} \succeq 2^{\frac{C|b-a|}{\varepsilon}}$.

We now present results for curves. Note that we will pay an estimation tax of $(1 + \frac{K^2\delta}{4})$ for viewing objects as curves instead of functions.

Theorem 3. *With notation as above,*

(a) $\mathcal{H}_\varepsilon(\mathcal{C}^1_{K,L}, \rho_1) \sim \frac{KL}{\varepsilon}$.

(b) $\mathcal{H}_\varepsilon(\mathcal{I}^1_{K,L}, \rho_1) \sim \frac{C}{\varepsilon}$, *for* $\frac{KL-2\pi}{\varepsilon} \leq C \leq KL$.

Proof. We compute the ε-entropy for a class by finding an ε-covering to obtain an upper bound, a 2ε-separated set (a collection of elements in the class that must each belong to a distinct ε-ball) to obtain a lower bound, so that the upper and lower bounds are asymptotically equal as $\varepsilon \to 0$. Note

that an ε-covering for $\mathcal{C}^1_{K,L}$ provides an upper bound for the other curve classes as well.

Denote the function class consisting of associated tangent angle functions to curves in $\mathcal{C}^1_{K,L}$ by $\Theta^1_{K,L} = \{\theta(s) : [0, L] \to \mathbb{R} \mid |\theta(s_1) - \theta(s_2)| \leq K|s_1 - s_2|, \theta(0) = 0\}$. Here we take s to be an arclength parameter with $\gamma(0) = \mathbf{0}$ for each $\gamma \in \mathcal{C}$.

1) **Upper Bound:** Take δ so that $\varepsilon = \delta + \frac{\sqrt{\delta}}{\lambda}$, and set $\xi = 3/2$. Begin by partitioning possible arclengths l, $0 < l \leq L$, into sequential subintervals of width $\frac{\delta^\xi}{4}$ with the first subinterval given by $(0, \frac{\delta^\xi}{4}]$. There will be $[\frac{4L}{\delta^\xi}]$ such subintervals. Parametrize a curve $\gamma \in \mathcal{C}^1_{K,L}$ of length $l \in (l_i - \frac{\delta^\xi}{4}, l_i]$ so that $\gamma(0) = 0$, the tangent angle $\theta(0) = 0$, and $|\frac{d\gamma}{dt}| = \frac{l}{l_i}$. For curves associated to each arclength subinterval, $|\theta(t_1) - \theta_(t_2)| \leq \frac{l}{l_i} K \leq K$.

   Each class Θ^1_{K,l_i} then admits an L^∞ δ-covering with at most $2^{\frac{Kl_i}{\delta}} \leq 2^{\frac{KL}{\delta}}$ elements by Theorem 1. This generates an ε-covering for $(\mathcal{C}^1_{K,l_i}, \rho_1)$. To see this, lift a δ-ball of tangent angle functions in Θ^1_{K,l_i} to its curve primitives in $\mathcal{C}^1_{K,l_i}$. This lifted ball will have ρ_1 radius of $2l_i \sin\frac{\delta}{2} + \lambda\delta \approx (l_i + \lambda)\delta$.

   We must refine the lifted covering to give a ρ_1 ε-covering. To do so, we correct for location (the first term in ρ_1) in three directions every $\Delta = \frac{\delta^{\xi-1}}{4}$ in arclength, introducing discontinuities into the curves at the centers of the refined balls. See Figure 2. The length of the discontinuity is $\frac{\delta^\xi}{4} + \frac{\delta\Delta}{2}$. We claim this procedure produces an ε-covering for $(\mathcal{C}^1_{K,l_i}, \rho_1)$.

   To see this, suppose $|\gamma(0) - \hat{\gamma}(0)| \leq \delta^\xi/2$, $|\theta(t) - \hat{\theta}(t)| \leq \delta$, and suppose that for $k = 0, 1, \ldots, [\frac{l_i}{\Delta}] - \Delta$, $|\gamma(k\Delta) - \hat{\gamma}(k\Delta)| \leq \delta^\xi/2$. Then:

$$\begin{aligned} |\gamma((k+1)\Delta) - \hat{\gamma}((k+1)\Delta)| &\leq \int_{k\Delta}^{(k+1)\Delta} 2\left|\sin\frac{\theta - \hat{\theta}}{2}\right| dt + \frac{\delta^\xi}{2} \\ &\leq \delta\Delta + \frac{\delta^\xi}{2} \\ &= \frac{3\delta^\xi}{4}. \end{aligned}$$

   Once we correct $\hat{\gamma}$ to return to within $\delta^\xi/2$ of any γ at $t = k\Delta$ for each $0 \leq k \leq [\frac{l_i}{\Delta}] - \Delta$, we will have refined the lifted covering of $(\mathcal{C}^1_{K,l_i}, \rho_1)$ to an ε-covering.

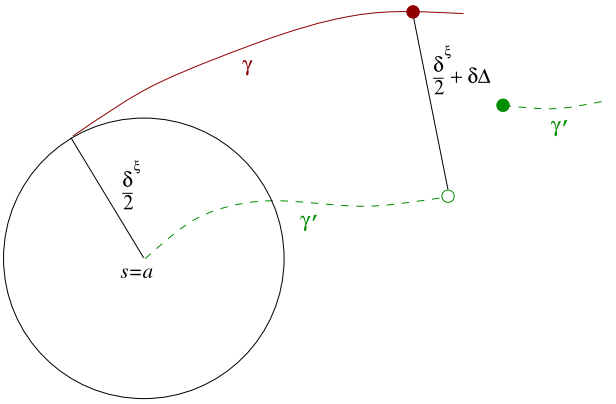

Figure 2: Refining the lifted covering to maintain the required ρ_1 radius by introducing jump discontinuities to the centers of the refined covering.

Let $P = \lim_{t\to k\Delta-} \hat{\gamma}(t)$ for some k. Consider the closed disk of radius $\frac{\delta^\xi}{4}$ centered at P. We set $\hat{\gamma}(k\Delta)$ to be one of the three points at a distance of $\frac{3\delta^\xi}{8}$ away from P along one of the radii at an angle of $0, \pi/3$ or $2\pi/3$ around the disk. If γ is in the refined ball associated to $\hat{\gamma}$, $\gamma(k\Delta)$ will be within the disk and will have distance from $\hat{\gamma}(k\Delta)$ of no more than $\frac{3\sqrt{3}\delta^\xi}{8} < \frac{\delta^\xi}{2}$. The curves $\gamma, \hat{\gamma}$ then satisfy our conditions in the interval $[k\Delta, (k+1)\Delta)$ and are therefore at a distance of no more than $\frac{3\delta^\xi}{4}$ there. As the tangent angles of two curves in a lifted ball are no more than δ apart, we have that each curve γ is within $\varepsilon = \delta^\xi + \lambda\delta$ of at least one center $\hat{\gamma}$ of the refined covering.
Repeating the above process for each k and each l_i produces an ε-covering with at most $[\frac{4L}{\sqrt{\delta^3}}]2^m$ balls, where $m = [\frac{KL}{\delta}] + \log 3[\frac{4L}{\sqrt{\delta}}]$, giving the desired asymptotic result.

2) **Lower Bound for $C^1_{K,L}$:** We construct a 2ε-separated set of curves by generating functions and viewing their realizations as plane curves. This turns out to be more delicate than one might initially suspect. Choose δ so that $\varepsilon = \frac{\delta}{1+\frac{K^2}{4}}$ and set $L' = L/\sqrt{1+K^2\delta}$. Subdivide $[0, L']$ into subintervals $I_k = [k\sqrt{\delta}, (k+1)\sqrt{\delta}]$, $k = 0, 1, \ldots, [\frac{L'}{\sqrt{\delta}}]$. For each I_k, apply Theorem 2 with $C = K$, $\varepsilon = \delta$, $a = k\sqrt{\delta}$, $b = (k+1)\sqrt{\delta}$ to obtain a collection of functions $f_{i,k}$, $i = 1, \ldots, [\frac{K}{\sqrt{\delta}}]$, satisfying for all i, k:

- $||f_{i,k} - f_{j,k}|| \geq 2\delta$ for $i \neq j$

- $f_{i,k}(a) = f_{i,k}(b) = 0$
- $\int_{I_k} f_{i,k} = 0$
- $||F_{i,k}||_\infty \leq \frac{K\sqrt{\delta}}{2}$.

From these, we generate for each i, k a primitive $g_{i,k}$ so that $g'_{i,k} = f_{i,k}$. These primitives have curvature bounded by K and satisfy for all i, k:

- $g_{i,k}(k\sqrt{\delta}) = g_{i,k}((k+1)\sqrt{\delta}) = 0$
- $g_{i,k}$ at x, $\theta_{i,k}(k\delta) = \theta_{i,k}((k+1)\delta)$ (where $\theta_{i,k}(x)$ denotes the tangent angle to the curve $(x, g_{i,k}(x))$)
- $||\theta_{i,k} - \theta_{j,k}|| \geq \frac{2\delta}{1+\frac{K^2\delta}{4}} = 2\varepsilon$.

Create a functions $\{g_i\}$ on $[0, L']$ by concatenating all possible sequences $\{g_{i,k}\}_k$. Since each function $g_{i,k}$ has length given by:

$$\begin{aligned}\int_{I_k} \sqrt{1+f_{i,k}^2} &\leq \int_{I_k} \sqrt{1+(K\frac{\sqrt{\delta}}{2})^2} \\ &= \sqrt{\delta}\cdot\sqrt{1+\frac{K^2\delta}{4}},\end{aligned}$$

the total arclength of any g_i (viewed now as a plane curve) is no more than:

$$\frac{L'}{\delta}\sqrt{\delta}\cdot\sqrt{1+\frac{K^2\delta}{4}} = L.$$

Furthermore, $\rho_1(g_i, g_j) \geq 2\varepsilon$ for $i \neq j$. Recall that $\lambda > 1/K$, and note that if $g_i \neq g_j$, we may assume (possibly by swapping labels i and j) that g'_i has slope K and g'_j has slope $-K$ on some interval $I = [x_0, x_0 + \delta/K]$ which means $|g'_i(x_0 + \delta/K) - g'_j(x_0 + \delta/K)| \geq 2\delta$. Without loss of generality, we may take $x_0 = 0$. Now consider the ρ_1-distance of g_j to the point $(\delta/K, g_i(\delta/K))$:

$$\begin{aligned}\rho_1(g_i, g_j) &\geq \rho_1\left((\delta/K, g_i(\delta/K)), g_j\right) \\ &\geq \min_x \frac{1}{\lambda}\left|\frac{\delta}{K} - x\right| + |\theta_i(\delta/K) - \theta_j(x)| \\ &= \min_x \frac{1}{\lambda}\left|\frac{\delta}{K} - x\right| + \left|\arctan g'_i(\delta/K) - \arctan g'_j(x)\right| \\ &\geq \min_x \frac{1}{\lambda}\left|\frac{\delta}{K} - x\right| + \frac{1}{1+\frac{K^2\delta}{4}}\left|g'_i(\delta/K) - g'_j(x)\right|\end{aligned}$$

$$\begin{aligned} &\geq \min_x \frac{1}{\lambda}\left|\frac{\delta}{K} - x\right| + \frac{1}{1+\frac{K^2\delta}{4}}\left|\delta + Kx\right| \\ &\geq \frac{2\delta}{1+\frac{K^2\delta}{4}} \\ &= 2\varepsilon \end{aligned}$$

The last inequality arises from finding the minimizing $x = \delta/K$.
Hence the curves described by the functions $\{g_i\}$ give a 2ε-separated set for ρ_1. All that remains is accounting: Each subinterval I_k produces m functions $f_{i,k}$ where $m \succeq 2^{\frac{K\sqrt{\delta}}{\delta}}$, and there are $[\frac{L'}{\delta}]$ subintervals, which gives a total number of curves in the 2ε-separated set asymptotically equal to $2^{\frac{K\sqrt{\delta}}{\delta}\cdot\frac{L'}{\delta}} = 2^{\frac{KL}{\varepsilon}}$.

3) **Lower Bound for $I^1_{K,L}$:** The lower bound for closed curves slightly modifies the argument for open curves. Generating a theorem like Theorem 2 that will generate closed curves corresponds to an unsolved question in partition theory. Instead, we sacrifice tightness of our lower bound.
Again, select δ so that $\varepsilon = \frac{\delta}{\sqrt{1+\frac{K^2\delta}{4}}}$. Take L' to satisfy

$$2L' \leq \frac{L}{\sqrt{1+\frac{K^2\delta}{4}}} - \frac{2\pi}{K}$$

and generate functions g_i as above that are ρ_1 2ε-separated. Taking any two such functions (not necessarily distinct), join two halves of a circle of radius $1/K$ to form a closed curve with Lipschitz tangent angle, Lipschitz constant K, and length bounded by L. The number of such curves is 2^m where $m \succeq \frac{2KL'}{\delta} = \frac{KL-2\pi}{\varepsilon}$.

□

2.2. Applications to adaptive coding for unbounded spaces

We replace the finite ε-covering giving an upper bound for the above ε-entropy estimates for the totally bounded spaces $\mathcal{C}_{K,L}$ and $\mathcal{I}_{K,L}$ with a countable ε-covering in the full curve spaces $\mathcal{C}^1$ and $\mathcal{I}^1$. Changing philosophy only slightly, one may view the ε-entropy of a class as the minimum number of bits required in a uniform encoding of the elements of that class with ε error. From this perspective, we consider a countable covering of a non-bounded space as an adaptive (non-uniform) encoding of the elements

of that space. The bits required for encoding a particular element represent the complexity of encoding that element in the given scheme. The number of balls will be infinite, but the number of bits for any given element will be finite. Again, we are interested in asymptotic behavior of the bit rate as $\varepsilon \to 0$.

We have seen that, given a good L^∞ approximation for the tangent angle function for a curve, obtaining a good ρ_1 approximation for the curve itself requires only lower order terms. In other words, the leading term in an adaptive approximation for a curve in $\mathcal{C}$ comes entirely from approximating the corresponding tangent angle. The following theorem for functions therefore tells the whole story for curves.

Theorem 4. *For every $\varepsilon > 0$, there exists a countable codebook $F_\varepsilon = \{f_1, f_2, \ldots\}$, depending only on ε, with the following property. For every Lipschitz function f defined on $[a, b]$ so that $f(a) = 0$ and $f'(x)$ is continuous a.e., there are constants $C(f, \delta)$ such that for all $\delta > 0$, there is a codeword $f_n \in F_\varepsilon$ such that $||f - f_n||_\infty \leq \varepsilon$ and f_n has description length:*

$$L(f_n) \leq \left[\frac{\int |f'| + \delta}{\varepsilon}\right] + C(f, \delta).$$

Proof. Since f' is continuous a.e. and bounded, f' is Riemann integrable. Therefore, for any δ, there exists a step function g taking on rational values, with a finite number of jumps at rational points $\{x'_j\}$, so that $|f'| \leq g$ and $\int g \leq \int |f'| + \delta$. On each subinterval $I_j = [x_j, x_{j+1})$ where g is constant, f is therefore Lipschitz with constant $g(x_j)$. Denote the number of jumps by m.

Using g, determine a variably spaced finite number of points $\{x_k\}$ so that for any k, $\int_{x_k}^{x_k+1} g \leq \varepsilon$. In particular, on each subinterval I_j, select the points spaced $\frac{\varepsilon}{g(x'_j)}$ apart. There will be at most:

$$\left[\frac{g(x'_j)|I_j|}{\varepsilon}\right] + 1 = \left[\frac{\int_{I_j} g(x)}{\varepsilon}\right] + 1$$

such points. Take $\{x_k\}$ to be the collection of $\{x'_j\}$ together with these equally spaced points.

Construct an approximation f_n for f. We claim there exists a piecewise linear function ϕ_n, with slope $\pm g(x_k)$ on the interval $J_k = [x_k, x_{k+1})$, vanishing at a, so that $f \subset K(\phi_n)$, where $K(\phi_n)$ is the corridor of width 2ε with ϕ_n as its top boundary. On J_1, take $\phi_n(x) = g(0)x$. Certainly, since f

is Lipschitz with constant $g(0)$ on I_1, $f \subset K(\phi_n)$. Inductively, assume ϕ_n has been constructed so that $f \subset K(\phi_n)$ for $x \leq x_k$. We wish to define ϕ_n for I_k so that f remains in $K(\phi_n)$. Since f is Lipschitz with constant $g(x_k)$ on I_k and $f \subset K(\phi_n)$ for $x \leq x_k$, one of the following is true:

(a) $f(x_{k+1}) \in [\phi_n(x_k) - g(x_k)(x - x_k), \phi_n(x_k) + g(x_k)(x - x_k)] \subset [\phi_n(x_k) - \varepsilon, \phi_n(x_k) + \varepsilon]$,

(b) $f(x_{k+1}) \in [\phi_n(x_k) - g(x_k)(x - x_k) - 2\varepsilon, \phi_n(x_k) + g(x_k)(x - x_k) - 2\varepsilon] \subset [\phi_n(x_k) - 3\varepsilon, \phi_n(x_k) - \varepsilon]$.

If (a) is true, then defining ϕ_n to have positive slope on I_k gives $f \subset K(\phi_n)$. If (b) is true, then defining ϕ_n to have negative slope on I_k gives the desired result. And so we have constructed a ϕ_n so that $f \subset K(\phi_n)$. Taking f_n to be the center of the corridor $K(\phi_n)$, we have $||f - f_n||_\infty \leq \varepsilon$.

Encode f by encoding f_n, or equivalently, ϕ_n. To do so requires describing g, which in turn requires describing the collection of points $\{x_k\}$ as outlined above. We must also describe the sequence of signs $\pm$ to assign to the slopes $g(x'_j)$ at each of the points $\{x_k\}$. Since g has rational jumps at rational values, encoding g requires a fixed and finite number of bits depending only on f and δ, yielding the constant $C(f, \delta)$. Describing a sign requires a single bit. As this must be done at each of the $\{x_k\}$, we see encoding the sequence of signs requires at most:

$$\sum_j \left(\left[\frac{\int_{I_j} g}{\varepsilon} \right] + 1 \right) + m \leq \left[\frac{\int_{[a,b]} g}{\varepsilon} \right] + 2m$$

bits. Then, absorbing $2m$ into $C(f, \delta)$, the total number of bits required to describe f_n satisfies:

$$L(f_n) \leq \frac{\int g}{\varepsilon} + C(f, \delta) \leq \frac{\int |f'| + \delta}{\varepsilon} + C(f, \delta),$$

as claimed. □

Corollary 1. *Suppose $\gamma \in \mathcal{C}$ has curvature $\kappa(s)$, s an arclength parameter. Then the leading term in the number of bits for an adaptive encoding of γ with ε-accuracy in the metric ρ_1 is at most:*

$$\left[\frac{\int |\kappa(s)|\, ds + \delta}{\varepsilon} \right].$$

Corollary 1 gives a measure of the complexity of a particular shape. It also allows now for an intrinsic evaluation of curve-based shape representations for any given shape. Fix $\delta > 0$. Presented with two candidate shape representations, the one requiring fewer bits for the given shape is the one that more efficiently captures that shape's structure (this is Rissanen's minimum description length [20]). For a comparison of boundary representation to the medial axis representation based on Corollary 1, see [18] where we obtain a precise criterion for when the boundary is more efficient.

3. Curve matching with a Sobolev-type metric

One of the most important problems in shape is curve matching – finding similar points on a pair of curves, for example in an object recognition problem. One typically wishes to introduce a notion of shape similarity (which may or may not be a metric) on a space of curves in order to make statements about which curves most closely resemble each other. We here introduce a multiscale analysis on curves which can be used to provide a local or global measure of similarity between a pair of curves. These geometric Sobolev-type metrics have their roots in a theorem from harmonic analysis which relates L^2 norms of fractional derivatives to averages of finite differences, and they provide both effective algorithms and strong statements about regularity.

3.1. Some classical results

Our wavelet-like constructions are motivated by the following theorem (see Stein [5]):

Theorem 5. *Let $f : \mathbb{R}^n \to \mathbb{R}$ be in $L^2 \cap C^1$. Then*

$$\int_{\mathbb{R}^n} \int_{\mathbb{R}^n} \frac{|f(x+t) + f(x-t) - 2f(x)|^2}{|t|^2} \frac{dt}{|t|^n} dx = c_n \int_{\mathbb{R}^n} \sum_{k=1}^{n} \left| \frac{\partial f}{\partial x_k} \right|^2 dx$$

where c_n depends only on the dimension.

Changing the power of $|t|$ reweights the various scales and will increase or decrease the number of derivatives being taken. For $0 < \alpha < 2$ we have that

$$\|f^{(\alpha)}\|_{L^2}^2 = c_{n,\alpha} \int_{\mathbb{R}^n} \int_{\mathbb{R}^n} \frac{|f(x+t) + f(x-t) - 2f(x)|^2}{|t|^{2\alpha}} \frac{dt}{|t|^n} dx.$$

While it is tempting to look at the second differences and think of the second derivative, this is not particularly useful here. The second differences are in fact the Haar wavelet coefficients of f', so the theorem is really looking at the first derivative. Note that the integral does not even converge at $\alpha = 2$.

There is also a local statement:

Theorem 6. *For almost every $x \in \mathbb{R}^n$, the function $f : \mathbb{R}^n \to \mathbb{R}$ has a weak derivative in the L^2 sense at x iff*

$$\int_{\mathbb{R}^n} \frac{|f(x+t) + f(x-t) - 2f(x)|^2}{|t|^2} \frac{dt}{|t|^n} < \infty.$$

Again, if we introduce the parameter α we get corresponding statements for the order-α derivative.

Remark 1. *The function*

$$F(x,t) = f(x+t) + f(x-t) - 2f(x)$$

defines, modulo affine functions, a lift of $f : \mathbb{R}^n \to \mathbb{R}$ to the space of functions $\mathbb{R}^n \times \mathbb{R}^n \to \mathbb{R}$. One may define the square function SF as

$$SF(x)^2 \equiv \int_{\mathbb{R}^n} \frac{|f(x+t) + f(x-t) - 2f(x)|^2}{|t|^2} \frac{dt}{|t|^n}.$$

One now sees that Theorems 5 and 6 both relate regularity of f to the square function SF. We will use very similar square functions later, when we lift a curve $\Gamma \subset \mathbb{R}^2$ to a function $\beta : \mathbb{R} \times \mathbb{R} \to \mathbb{R}$ which has a similar square function associated with it.

Let $\gamma(s)$ be an arclength[1] parametrization of a smooth curve $\Gamma \subset \mathbb{R}^2$ and let $\{z_n\}_{i=1}^N$ be a collection of sample points taken at equal arclength along Γ (we are not yet assuming any noise in the samples). Given the parametrization $\gamma(s) = x(s) + iy(s)$, we define the angle $\beta(s,t)$ with respect to γ by

$$\beta(s,t) \equiv \arg \frac{\gamma(s+t) - \gamma(s)}{\gamma(s) - \gamma(s-t)}.$$

[1] Arclength sampling is not at all necessary; we assume it mostly for a concise presentation. The results can easily be modified if, for example, the sampling intervals are bounded above and below.

We introduce the α "norm"[2] for curves, defined for closed curves as

$$\|\gamma\|_\alpha^2 = \int_0^L \int_0^{L/2} \beta(s,t)^2 \frac{dt}{|t|^{2\alpha-1}} ds \tag{3}$$

where L is the length of Γ, and for open curves as

$$\|\gamma\|_\alpha^2 = \int_0^L \int_0^{\min(s,L-s)} \beta(s,t)^2 \frac{dt}{|t|^{2\alpha-1}} ds. \tag{4}$$

In making computations on discretized closed curves[3] we will use the equivalent dyadic variant

$$\|\gamma\|_\alpha^2 = \sum_{n=1}^{N} \sum_{k=1}^{K} \beta(n,k)^2 2^{-2k(1-\alpha)} \tag{5}$$

where

$$\beta(n,k) = \arg \frac{\gamma_{n+2^{k-1}} - \gamma_n}{\gamma_n - \gamma_{n-2^{k-1}}}.$$

These constructions give geometric variants of theorem (5); the β's are to curves what the second differences are to functions. The next subsection gives these results in the dyadic case.

3.2. The mapping from β's to wavelet coefficients

Consider $\Gamma \subset \mathbb{C}$ which is the graph of a Lipschitz function; i.e. $\Gamma = \{x + iA(x)\} \subset \mathbb{C}$ where $A : \mathbb{R} \to \mathbb{R}$ satisfies $\|A'\|_\infty < M$. If $\gamma : \mathbb{R} \to \mathbb{C}$ is an arclength parametrization of Γ we define $\phi : \mathbb{R} \to [-\arctan M, \arctan M]$ a.e. by writing

$$\gamma'(s) = e^{i\phi(s)}.$$

[2]It is not correct to refer to this as a norm on curves; we have not imposed a linear structure on the space of curves. However, these integrals of β angles play a role analogous to the Sobolev norms on function spaces and the β's themselves share many features with wavelet coefficients from the point of view of analysis. We hope the reader will forgive this abuse of terminology.

[3]Open curves can be handled as well; the easiest method is simply truncating the sum near the endpoints.

Let $\mathcal{D}$ be the set of dyadic intervals of $\mathbb{R}$. Given $I \in \mathcal{D}$, we write x_I for the midpoint of I and x_I^+, x_I^- for the right and left endpoints, respectively. Denote by h_I the Haar wavelet associated with I:

$$h_I(x) = \begin{cases} |I|^{-1/2}, & x \in [x_I, x_I^+] \\ -|I|^{-1/2}, & x \in [x_I^-, x_I). \end{cases}$$

We will generally write x_I for intervals in the domain of a function and s_I for intervals in the domain of a curve. Given our curve $\gamma(s)$, we may consider the dyadic β angles

$$\beta_I \gamma = \arg \frac{\gamma(s_I^+) - \gamma(s_I)}{\gamma(s_I) - \gamma(s_I^-)}$$

as well as the Haar coefficients of γ'

$$a_I = \langle \gamma', h_I \rangle.$$

We see that in the sense of distributions, the Haar coefficients of γ' in fact agree with the second differences of γ, up to a rescaling:

$$\begin{aligned} a_I &= \langle \gamma', h_I \rangle \\ &= |I|^{-1/2} \left(\gamma(s_I^+) + \gamma(s_I^-) - 2\gamma(s_I) \right). \end{aligned}$$

Geometrically, a_I is a vector joining $\gamma(s_I)$ to the opposing vertex of the parallelogram defined by the three points $(\gamma(s_I^+), \gamma(s_I), \gamma(s_I^-))$. The next theorem shows that on Lipschitz graphs, the β's are almost the wavelet coefficients a_I. While this is pointwise false, it becomes true after integrating. We first state a lemma. Decompose a_I into "tangential" and "normal" components

$$a_I = a_I^t + a_I^n$$

where

$$\begin{aligned} a_I^t &= \mathrm{proj}_{\ell_I} a_I \\ a_I^n &= \mathrm{proj}_{\ell_I^\perp} a_I \end{aligned}$$

and ℓ is the line joining $\gamma(s_I^-)$ to $\gamma(s_I^+)$ (see figure 3). We think of ℓ as a rough tangent line, thus the terminology. We will use the fact that $|a_I|^2 = |a_I^t|^2 + |a_I^n|^2$. Our next lemma says that the β's and the normal components a_I^n agree up to a rescaling; the proof is a simple geometric argument.

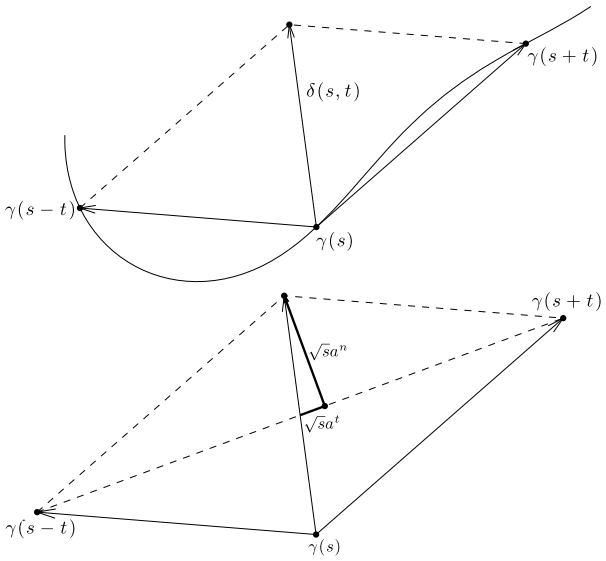

Figure 3: Wavelet coefficients of complex-valued γ' in the continuous setting. Left: The wavelet coefficients $a(s,t)$ are, up to the L^2 rescaling, given by the vector $\delta(s,t)$ joining $\gamma(s)$ to the opposite vertex of the parallelgram. Right: Again up to rescaling, the "normal" and "tangential" parts a^n, a^t are given by the projections of δ onto the line ℓ joining $\gamma(s-t)$ to $\gamma(s+t)$.

Lemma 1 (β is a_I^n)**.** *With β, γ, ϕ, a_I defined as above,*

$$|a_I^n| \sim |\beta_I||I|^{1/2}$$

with constant that depends only on M.

In fact, under certain conditions, we can use β in place of the entire coefficient a_I. Pointwise, this is false, but on Lipschitz graphs, a_I^n controls a_I^t after integrating, when α is large enough. The next theorem makes this precise. Let $\mathcal{I}$ be a dyadic partition of the domain of A.

Theorem 7 (β's are almost wavelet coefficients)**.** *For $0 < \alpha < 1$,*

$$\sum_{I \in \mathcal{I}} |a_I|^2 |I|^{-2\alpha} \sim \sum_{I \in \mathcal{I}} \beta_I^2 |I|^{1-2\alpha}$$

with constant that depends only on the Lipschitz constant M.

Remark 2. *The condition $\alpha > 0$ in Theorem 7 is necessary, and the counterexample for other α contains the gist of the proof. Consider the curve*

Figure 4: The condition $\alpha > 0$ is necessary in Theorem 7. For $\alpha < 0$, the cost of the wiggly part vanishes as the number of wiggles increases.

of length 2 obtained by joining the line segment $[-1, 0] \subset \mathbb{R}$ *to* 2^k *triangle functions with sides of slope* $\pm M$*, as shown in figure 4. Let* γ *be an arclength parametrization of this curve. Note that the rightmost endpoint of the curve lies at* $(\frac{1}{\sqrt{1+M^2}}, 0)$*, independent of* k*, so the tangential part of the largest* a_I *is non-zero. If* $\mathcal{I}$ *is a dyadic partition of* $[-1, 1]$ *we see there are* 2^k *non-zero* β_I*'s which share some common value* β*. We have*

$$\begin{aligned}\sum_{I\in\mathcal{I}} \beta_I^2 |I|^{1-2\alpha} &= \beta^2 2^k \left(2^{-k}\right)^{1-2\alpha} \\ &= \beta^2 2^{2\alpha k}.\end{aligned}$$

If $\alpha < 0$ *we can make this arbitrarily small by increasing* k*, so no variant of Theorem 7 can hold.*

In light of Theorem 7 it follows that statements about local regularity of $\gamma(s)$ which can be obtained via decay of wavelet coefficients can be translated immediately into an equivalent statement about the β's. For example, Theorem 7 has an immediate corollary:

Corollary 2 (Characterization of Sobolev spaces via β)**.** *For* $0 < s < 1$*,* $\gamma \in H^{1+s}$ *iff*

$$\sum_{I\in\mathcal{I}} \beta_I^2 |I|^{1-2\alpha} \leq \infty$$

Proof. By Theorem 7, the fractional derivative seminorms induced by the a_I and the β's are comparable, and γ is always in L^2 by construction, hence this is simply a restatement of the definition of H^{1+s}. □

In summary, studying the β angles and their decay at small scales is equivalent to studying the wavelet coefficients of the derivative γ'. Any statement about regularity obtained using one construction translates immediately into a corresponding statement about the other. In the next sections we will construct matching algorithms based on square functions of β and other related quantities.

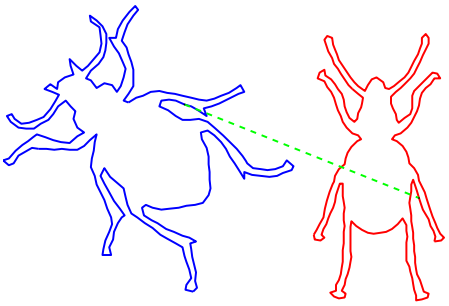

Figure 5: Curve alignment. These are two images from the MPEG7 dataset. The point $\gamma_1(\theta)$ on the blue curve Γ_1 is matched to $\gamma_2 \circ \sigma(\theta)$ on Γ_2.

3.3. Curve matching

We now apply the theoretical results of the previous section to design a family of curve matching algorithms. Let $\gamma_1, \gamma_2 : [0, 2\pi) \to \mathbb{C}$ be unit-speed parametrizations of some pair of plane curves Γ_1, Γ_2. For simplicity we will assume these are Jordan curves; the main reason is that the periodic versions of all our formulas are simpler to state than the versions which allow for endpoints of arcs. We normalize both curves to have length 2π.

A standard technique to define a distance between two curves is to consider all aligning homeomorphisms $\sigma : [0, 2\pi] \to [0, 2\pi]$ which define correspondences between points $\gamma_1(\theta)$ on Γ_1 and $\gamma_2 \circ \sigma(\theta)$ on Γ_2 (see Figure 5). One then uses geometric criteria to associate a cost with the map σ. For example, using the β angles described above, one could match curves as follows. For $i = 1, 2$, define the angles on the i-the curve at scale t:

$$\beta_i(\theta, t) = \arg \frac{\gamma_i(\theta + t)}{\gamma_i(\theta - t)}.$$

In analogy with square functions in analysis (see [6]), we now define a type of geometric square function associated with the angles β. In fact, more generally we can allow β to be any lift of a curve to $\mathbb{R}^2$ (or some subset of $\mathbb{R}^2$ and make the following definition.

Definition 2 (Geometric square function)**.** *Let $\mathcal{S}$ be the space of simple closed curves of length 2π and let $P : \mathcal{S} \to [0, 2\pi] \times \mathbb{R}$ be any lift. Define the square function associated with P to be*

$$S_\alpha P(\theta, t)^2 = \int_0^\infty P(\theta, t)^2 \frac{dt}{t^{2\alpha - 1}}.$$

In the case where $P = \beta$, for example, we recover the geometric Sobolev norms introduced in the previous section. If γ is an arclength parametrization of a curve, then

$$\|\gamma\|_\alpha^2 = \int_0^{2\pi} S_\alpha \beta(\theta)^2 d\theta.$$

As we discuss below, we will consider lifts other than β.

Now given a map $\sigma : [0,1] \to [0,1]$ and a choice of $\alpha \in (0,2)$ we define the matching cost

$$K(\sigma) = \int_0^{2\pi} \int_0^{\infty} \left(\beta_1(\theta,t) - \beta_2(\sigma(\theta),t)\right)^2 \frac{dt}{t^{2\alpha}} d\theta. \tag{6}$$

This is the natural norm on the lift β which is compatible with the square function $S\beta$. To define a distance, one chooses a class Σ of maps (say, homeomorphisms, or Lipschitz mappings), and defines

$$d^2(\Gamma_1, \Gamma_2) = \inf_{\sigma \in \Sigma} K(\sigma)$$

or the symmetric variant

$$d^2_{sym}(\Gamma_1, \Gamma_2) = d^2(\Gamma_1, \Gamma_2) + d^2(\Gamma_2, \Gamma_1).$$

We have the following properties for d_{sym}.

Proposition 1. *Let Γ_1, Γ_2 be continuous closed curves and let Σ be the set of homeomorphisms $[0,1] \to [0,1]$. The distance d_{sym} defined by*

$$d^2_{sym}(\Gamma_1, \Gamma_2) = d^2(\Gamma_1, \Gamma_2) + d^2(\Gamma_2, \Gamma_1)$$

satisfies

1) $d_{sym}(\Gamma_1, \Gamma_2) \geq 0$
2) $d_{sym}(\Gamma_1, \Gamma_2) = d_{sym}(\Gamma_2, \Gamma_1)$
3) $d_{sym}(\Gamma_1, \Gamma_2) = 0$ *iff* $\Gamma_1 = \Gamma_2$.

Proof. Non-negativity and symmetry are obvious. We claim that $d(\Gamma_1, \Gamma_2) = 0$ iff $\Gamma_2 = \Gamma_1$ essentially follows from the fact that β is a lift. Indeed, assume we have two curves $\Gamma_1 \neq \Gamma_2$, but $d(\Gamma_1, \Gamma_2) = 0$. The second condition implies $\beta_1 = \beta_2$ a.e., and since β is continuous in θ and t we have $\beta_1 \equiv \beta_2$, hence $\Gamma_1 = \Gamma_2$. □

3.4. Discretization and computation

In order to achieve greater robustness, we wish to weaken the matching cost (6) significantly. In the discrete setting, we proceed as follows. Say we have N sample points taken at arclength along each curve. We wish to map some $M \leq N$ points on Γ_1 to similar points on Γ_2. The aligning transformation σ then becomes a map

$$\sigma : \{1, 2, ..., M\} \to \{1, 2, ..., N\} \tag{7}$$

Once we fix a maximum scale of interest K, our distance can be approximated as

$$\mathrm{d}(\Gamma_1, \Gamma_2) = \min_{\sigma \in \Sigma_N} \sum_{n=1}^{M} \sum_{k=1}^{K} k^{-(1+2\alpha)} |\beta_1(n,k) - \beta_2(\sigma(n),k)|^2$$

where Σ_N is some class of admissible discrete alignments and the $\beta_i(n,k)$ are the β's computed about the n-th point on curve i at scale k. By scale k, we mean the k-th dyadic scale; i.e. if $z(n)$ are equally-spaced points along a curve, then set

$$\beta(n,k) = \arg\Big(z(n+2^k) - z(n)\Big) - \arg(z(n) - z(n-2^k))$$

for $k = 0, 1, 2, ..., K$. In our experiments below, we have made the following choices:

1) Each curve is resampled to 256 uniformly-spaced points. We match 128 randomly-chosen points on Γ_1 to their best matches on Γ_2 and vice-versa to compute d_{sym}.
2) Σ_N is the collection of all injective maps $\sigma : \{1, 2, ..., 128\} \to \{1, 2, ..., 256\}$, without any regard for topology.
3) The β's are computed at 7 dyadic scales. We weight each scale corresponding to the choice of $\alpha = 1/2$ to produce a $H^{(1/2)}$-type norm.
4) The optimal match may be computed with, for example, the Hungarian algorithm; in practice we use a fast greedy algorithm which loops through a random permutation of the 128 points on Γ_1, pairing each point to the best match on Γ_2 that has not already been paired. We have not observed any significant difference between the classification performance of this naive greedy algorithm compared to the optimal match.

5) The symmetric distance d_{sym} was used in all experiments.

Remark 3. *The choice of $H^{1/2}$ is natural in the following way: if one assumes that the boundaries in the dataset have been corrupted by a white noise of a fixed (but unknown) variance, the $H^{1/2}$ norm weights each scale according to the signal-to-noise ratio in that octave. Indeed, the empirical results confirm that this is a reasonable choice, as the best results are obtained very close to $H^{1/2}$. $H^{1/2}$ is also the Sobolev critical exponent at which the curve becomes continuous.*

3.4.1. Other multiscale geometric quantities In addition to computing the angles β, we also implement the same algorithm, but replacing the β's with the log ratio of point separation to arclength[4]. We denote this by

$$\ell(n,k) = \log \frac{|z(\theta+t) - z(\theta-t)|}{2t}.$$

Since Theorem 5 also holds in the case of complex-valued functions, we consider the complex second differences; these are not naturally scale-invariant so we normalize by the length of the underlying arc, which is 2^{-k} given our normalizations described above.

$$\delta(n,k) = 2^{-k}\left(z(n+2^k) + z(n-2^k) - 2z(n)\right)$$

Rotation-invariance was achieved for the complex second differences in two ways. Best performance was obtained by rotating so the tangent to Γ_1 pointed along the real axis before computing the δ's about a given point. Rotating so the "coarse tangents", i.e. the line joining the endpoints of the arc at scale k, pointed along the real axis gave only slightly lesser performance, with the loss almost certainly owing to the fact that this allows inconsistent matching to take place across scales.

Summarizing, we have not one method, but a family of curve matching algorithms based on the distance d_{sym}. The matching algorithm is the same in all cases; the only difference is the geometry under consideration. We have tried several multiscale measures, swapping the angles β for the complex differences δ or the length ratios ℓ. In each case we use dyadic scales and make the scale-invariant choice of placing equal weight on each scale (equivalent to looking at the 3/2 derivative). We list the choices here and summarize the discussion above; experimental results appear in the next section.

[4]This is asymptotically the angle $\beta(s,t)^2$, so one can use ℓ to study local regularity. We use it here to capture distortion of inter-point distances which is not captured by β.

1) Angles β
   These are the β angles described above.
2) Complex second differences δ, tangent normalization
   The complex second differences correspond to thinking of the curve as a complex-valued function γ on the unit circle with $|\gamma'| = 1$ and applying Theorem 5. We normalize the second differences by arclength, as described above. "Tangent normalization" means that we achieve rotation invariance by rotating the curve so the tangent line coincides with the real axis before computing δ's about each point.
3) Complex second differences δ, "coarse tangent" normalization
   This differs from the previous only in how we achieve rotation invariance: instead of rotating so the tangent to the curve lies along the real axis, we rotate differently at each point and each scale. We rotate so the "coarse tangent", i.e. the line joining the endpoints of the arc at scale k, lies along the real axis.
4) Distance ratios ℓ
   These are the ratios of point separation to arclength, as described above.
5) Blend: β, ℓ
   This is a blended distance obtained by normalizing β and ℓ so the mean distances are equal and then summing the two, equally weighted.

3.5. Experimental results: MPEG7 CE Shape-1 Part B

The MPEG7 CE Shape-1 B test dataset is generally felt to be of the hardest datasets in the literature for testing shape classification algorithms. The dataset contains 20 binary images in each of 70 categories. Some are hand-drawn, some were extracted by thresholding digital images, some are cartoons, etc. The dataset has essentially every possible major source of difficulty, including occluders, noise, and widely-varying resolutions. On the other hand, it does not have many examples of shapes with articulated parts. In any event, the standard test benchmark is the following.

1) For each image I_k in the dataset, $1 \leq k \leq 1400$
   a) Choose 40 other images from the dataset
   b) Count how many of the 40 are in the same category as I_k (max possible is 20)
2) Final score is called the "bull's-eye percentage":

$$BEP = \frac{\text{total images in correct class}}{28000}$$

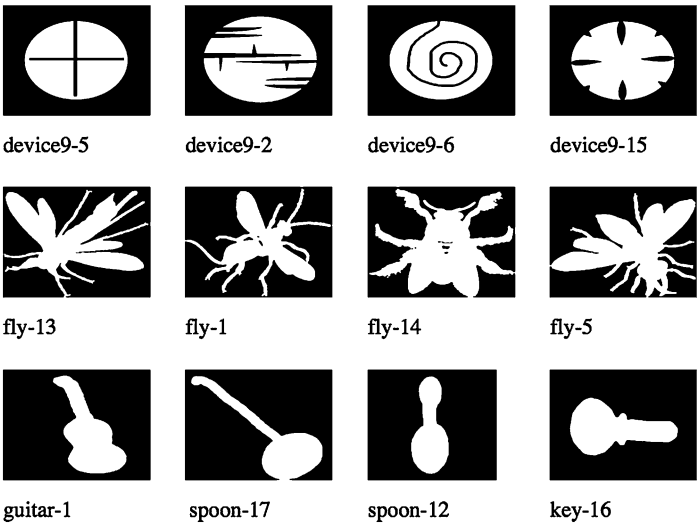

Figure 6: Some sample images from the MPEG7 test dataset. In certain categories, like the spoons, guitar, and key images in the bottom row, it is not clear that any shape-based algorithm (nor human subjects) would be able to correctly classify the images. Image from Veltkamp/Latecki '00.

$(28000 = 20 * 1400$ correct is the best possible.)

Some representative results from existing methods. Many, many more approaches exist, some with very good scores.

Method	*BEP*
Shape-tree [7]	87.7
Inner-distance Shape Context [3]	85.4
Shape Context [1]	76.51
Curvature Scale Space [4]	75.44

We report the following results using the methods described above:

Method	*BEP*
Complex second differences δ, tangent normalization	85.64
Complex second differences δ, coarse tangent normalization	84.57
Blend: β, ℓ	78.59
Angles β	75.43
Distance ratios ℓ	74.21

The complex second differences yield performance nearly equal to the shape-tree method even using a much weaker matching algorithm. The β angles and log-distance ratios ℓ both perform reasonably well alone, but conspire to give better performance when blended. The blend was created by normalizing both distances by the mean distance and summing their squares with equal weight on each. The point seems to be that measuring angle and distance alone is good; measuring them both together is better.

3.6. Remarks

The $H^{1/2}$-type methods described here use a rich multiscale curve descriptor, similar to the shape-tree and shape context methods mentioned above. The descriptor is rich enough that a fast greedy algorithm produces very good results. If K is the number of scales at which β is computed, and N is the number of sample points, then the alignment is $O(KN^2)$ where the bulk of the time is spent computing the fragment-to-fragment distance matrix. Most important, the constant is small and an entire pass through the MPEG7 dataset using MATLAB on a 2009 Mac Pro workstation can be completed in less than 20 minutes. Repeating the same experiments using the $O(N^3)$ Hungarian matching algorithm, rather than the greedy algorithm, resulted in negligible increases in performance (on the order of .5%) while increasing the runtime to approximately 12 hours, again in MATLAB on the same workstation.

The multiscale descriptor used here can also be used as a local descriptor to perform optimal alignments of one curve to another using a dynamic programming algorithm which preserves fragment order along the curve. The resulting algorithm is very similar to the dynamic time warping (DTW) methods in common use in applications such as handwriting recognition, with a rich local descriptor (see [11] for related geometric algorithms). In the case of the MPEG7 dataset, the dynamic programming alignments actually produced worse results than the greedy nontopological algorithm described above. However, this should not be terribly surprising when one looks at the dataset, since many curves within a given class exhibit different topologies, or strikingly different placement of salient features. Fragment-based approaches are more relevant when feature locations change significantly. In informal experiments, the $H^{1/2}$ multiscale descriptor appears to be a promising candidate for a local descriptor in handwriting recognition algorithms.

4. Weil-Petersson metric: conformal mappings, fingerprints, and Teichons

4.1. Shapes as diffeomorphisms of the circle S^1

4.1.1. Fingerprints Let $\mathbb{D}_{\text{int}}$ be the open unit disk in the complex plane $\mathbb{C}$, i.e. $\mathbb{D}_{\text{int}} = \{z \in \mathbb{C} \mid |z| < 1\}$, and let $\mathbb{D}_{\text{ext}} = \{z \in \mathbb{C} \mid |z| > 1\}$ be its exterior.

For every simple closed curve Γ in $\mathbb{C}$ denote by Γ_{int} its union with the region enclosed by it, and denote by Γ_{ext} its union with the infinite region outside of Γ (including ∞).

Then by the Riemann mapping theorem, for all Γ there exist two conformal maps

$$f_{\text{int}} : \mathbb{D}_{\text{int}} \to \Gamma_{\text{int}}, \qquad f_{\text{ext}} : \mathbb{D}_{\text{ext}} \to \Gamma_{\text{ext}}.$$

The interior map f_{int} is unique up to replacing f_{int} by $f_{\text{int}} \circ A$ for any Möbius transformation $A : \mathbb{D}_{\text{int}} \to \mathbb{D}_{\text{int}}$, where A is defined as

$$A(z) = \frac{az + b}{\bar{b}z + \bar{a}}, \ |a|^2 - |b|^2 = 1. \tag{8}$$

This subgroup of Möbius group of selfmaps of the circle is denoted $\text{PSL}_2(\mathbb{R})$.

The map f_{ext} is normalizated such that f_{ext} maps ∞ to ∞, and that its differential carries the real positive axis of the $\mathbb{D}$-plane at infinity to the real positive axis of the Γ-plane at infinity.

Thus we define the map ψ which is called the 'fingerprint' (in Teichmüller theory this is known as a 'welding map') of the shape

$$\psi = f_{\text{int}}^{-1} \circ f_{\text{ext}} \in \text{PSL}_2(\mathbb{R}) \backslash \mathbf{Diff}(S^1). \tag{9}$$

Note, that $f_{\text{ext}}(S^1) = \Gamma$, $f_{\text{int}}^{-1}(\Gamma) = S^1$. The fingerprint $\psi : S^1 \to S^1$ is a real-valued orientation-preserving diffeomorphism, and it uniquely identifies the shape Γ (modulo scaling and rigid translations). Due to the Möbius transformation ambiguity in the choice of f_{int}, we see by construction that ψ is a member of the right coset space $\text{PSL}_2(\mathbb{R}) \backslash \mathbf{Diff}(S^1)$. An example of a shape along with four realizations of its fingerprint is given in Figure 7.

The inverse map from diffeomorphisms to shapes is called *welding*, and is defined as follows: starting with ψ, construct an abstract Riemann surface by 'welding' the boundaries of $\mathbb{D}_{\text{int}}$ and $\mathbb{D}_{\text{ext}}$ via ψ. The resulting Riemann surface must be conformally equivalent to the Riemann sphere. Choose a conformal map f from the welded surface to the sphere taking $\infty \in \mathbb{D}_{\text{ext}}$ to itself and having real positive derivative there. Let $\Gamma = f(S^1)$ (for details and the numerical implementation see [46]).

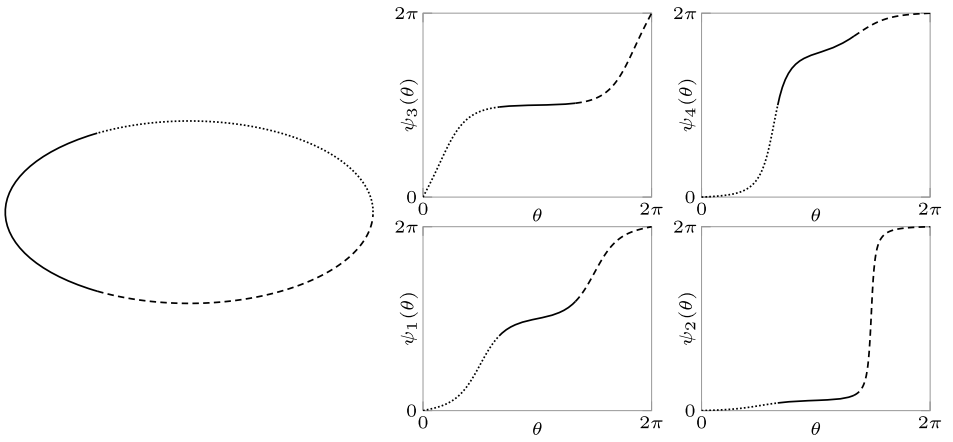

Figure 7: Left: an ellipse. Right: various welding map representatives of the ellipse shape corresponding to the same equivalence class in $\mathrm{PSL}_2(\mathbb{R})\backslash\mathbf{Diff}(S^1)$.

4.1.2. Weil-Petersson norm The Lie algebra of the group $\mathbf{Diff}(S^1)$ is given by the vector space $\mathbf{Vec}(S^1)$ of smooth periodic vector fields $v(\theta)\partial/\partial\theta$ on the circle. The Weil-Petersson metric for $v \in \mathbf{Vec}(S^1)$ can be expressed as [45]:

$$\|v\|^2_{WP} = \sum_{n\in\widehat{\mathbb{Z}}} |n^3 - n||v_n|^2 \tag{10}$$

$$= \int_{S^1} Lv(\theta)v(\theta)d\theta. \tag{11}$$

Here $v(\theta) = \sum_{n=-\infty}^{\infty} v_n e^{in\theta}$ (where $\overline{v_n} = v_{-n}$ for the vector field to be real), and $\widehat{\mathbb{Z}} = \mathbb{Z}\backslash\{n = 0, \pm 1\}$. The Weil-Petersson operator L is an integro-differential operator with the form

$$L = -\mathcal{H}(\partial_\theta^3 + \partial_\theta). \tag{12}$$

Above, $\mathcal{H}$ is the periodic Hilbert transform, defined as a convolution with $\frac{1}{2\pi}\cot(\theta/2)$.

The null space of the L operator is given by the vector fields whose only Fourier coefficients are v_{-1}, v_0 and v_1, i.e. vector fields of the type $(a + b\cos\theta + c\sin\theta)\partial/\partial\theta$. These vector fields are exactly in the Lie algebra $\mathfrak{sl}_2(\mathbb{R})$ of the Lie group $\mathrm{PSL}_2(\mathbb{R})$.

It is a known fact [39] that the above WP norm on vector fields can be extended to the right-invariant *WP-Riemannian metric* on the coset space $\mathrm{PSL}_2(\mathbb{R})\backslash\mathbf{Diff}(S^1)$.

Consider any two diffeomorphisms $\psi_0, \psi_1 \in \mathrm{PSL}_2(\mathbb{R})\backslash\mathbf{Diff}(S^1)$. The Riemannian distance induced by the WP norm on vector fields is given by

$$L = \int_0^1 \|v(s)\|_{WP} ds, \text{ where } v(\theta,t) = \partial\psi(\psi^{-1}(\theta,t),t)/\partial t. \tag{13}$$

Vector fields v that minimize the distance (13) are geodesics on $\mathrm{PSL}_2(\mathbb{R})\backslash\mathbf{Diff}(S^1)$.

4.2. The geodesic equation

The Euler-Poincaré equation for diffeomorphisms (hereafter 'EPDiff') is a variant of Euler's equations for fluid flow. It describes geodesics on the Lie group of diffeomorphisms of $\mathbb{R}^n$ in any right invariant metric given on vector fields by $\|v\|^2 = \int_{\mathbb{R}^n} \langle Lv, v\rangle dx$ for some positive definite self-adjoint operator L (where $\langle\cdot,\cdot\rangle$ is the canonical L^2 pairing). The general EPDiff($\mathbb{R}^n$) is derived in [24] and has the form

$$\frac{\partial}{\partial t} Lv + (v\cdot\nabla)(Lv) + \operatorname{div} v\, Lv + Dv^t \cdot Lv = 0,$$

where v is a smooth vector field in $\mathbb{R}^n$, $\nabla = (\frac{\partial}{\partial x_1}, \ldots, \frac{\partial}{\partial x_n})^T$ is the divergence operator, L is a self-adjoint differential operator and Dv is a Jacobian matrix.

The above formula extends to the setting of the homogeneous space $\mathrm{PSL}_2(\mathbb{R})\backslash\mathbf{Diff}(S^1)$ [37, 39]. Thus given a path $\phi_t(\theta) = \phi(\theta,t)$ in $\mathbf{Diff}(S^1)$, let $v(\theta,t) = \frac{\partial\phi}{\partial t}(\phi^{-1}(\theta,t),t)$ be the scalar vector field it defines on a circle and let L be the Weil-Petersson differential operator $L = -\mathcal{H}(\partial_\theta^3 + \partial_\theta)$. Then EPDiff takes the form

$$(Lv)_t + v.(Lv)_\theta + 2v_\theta.Lv = 0. \tag{14}$$

Above, $v(\theta,t)$ is called the velocity of the path, $m(\theta,t) = Lv(\theta,t)$ is the momentum, and this equation is the same as introduced in (1). We note in particular that the momentum can be a distribution. The $v \to m$ map may be inverted by the relation $v(\theta,t) = G * m(\theta,t)$, where G is the Green's function $G(\theta)$ of the WP operator L. The Green's function $G(\theta)$ is obtained as a solution to $LG = \mathrm{Proj}(\delta_0)$, where δ_0 is the Dirac measure centered

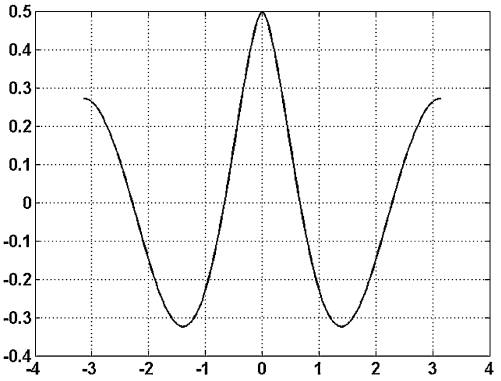

Figure 8: Green's function (15) of the Weil-Petersson operator $L = -\mathcal{H}(\partial_\theta^3 + \partial_\theta)$.

at $\theta = 0$, and $\mathrm{Proj}(\delta_0)$ is the projection of δ_0, such that 0th, $\pm$1st Fourier coefficients of $\mathrm{Proj}(\delta_0)$ are zero. Expression for $G(\theta)$ has been computed in [38], see Fig. 8:

$$G(\theta) = (1 - \cos\theta)\log\left[2(1 - \cos\theta)\right] + \frac{3}{2}\cos\theta - 1. \tag{15}$$

4.3. Teichons, singular solutions of EPDiff

The EPDiff equation (14) admits momenta solutions that, once initialized as a sum of N Dirac measures, remain a sum of N Dirac measures for all time [27, 33]. In reference to this self-similarity property, these singular solutions are named Teichons (or N-Teichons).

For a solution to EPDiff (14), we employ the N-Teichon ansatz

$$m(\theta, t) = \sum_{j=1}^{N} p_j(t)\delta(\theta - q_j(t)), \tag{16a}$$

$$v(\theta, t) = \sum_{j=1}^{N} p_j(t)G(\theta - q_j(t)), \tag{16b}$$

where $\delta \triangleq \delta_0$ is the origin-centered Dirac mass. Plugging these expressions into EPDiff (14), we obtain a system of ODEs describing the evolution of

the momentum coefficients p_k and the Teichon locations q_k:

$$\begin{cases} \dot{p}_k = -p_k \sum_{j=1}^{N} p_j G'(q_k - q_j), \\ \dot{q}_k = \sum_{j=1}^{N} p_j G(q_k - q_j). \end{cases} \tag{17}$$

Momenta m must lie in the horizontal space, i.e. $m(\theta, t)$ must have vanishing 0^{th} and $\pm 1^{\text{st}}$ Fourier coefficients. Using (16a), we obtain a set of three constraints for (q_k, p_k), linear in p_k:

$$\sum_{j=1}^{N} p_j = \sum_{j=1}^{N} p_j e^{iq_j} = \sum_{j=1}^{N} p_j e^{-iq_j} = 0. \tag{18}$$

If they are satisfied at time $t = 0$ they will be satisfied for all t. The Teichons never collide: the Teichon locations q_k retain their initial ordering on S^1 for all time.

The system (17) is a Hamiltonian system and can be efficiently solved using a symplectic integrator, for example Lobatto IIIA-B (for details see [38]).

4.4. Numerical simulations

We start with the initial shape as a circle, and evolve it with an N-Teichon. In other words we have the vector field of the form

$$v(\theta, t) = \sum_{k=1}^{N} p_k(t) G(\theta - q_k). \tag{19}$$

The evolution of the shape is the result of the corresponding evolution of the fingerprint

$$\frac{\partial \psi(\theta, t)}{\partial t} = v(\psi(\theta, t), t), \quad \psi(\theta, 0) = \theta. \tag{20}$$

Note, we are not solving a boundary value problem, but an initial value problem. In other words we specify the initial velocity (i.e. $p_k(0)$ and $q_k(0)$), solve forward equations (17) to obtain evolution of $\{p_k(t)\}_{k=1}^{N}, \{q_k(t)\}_{k=1}^{N}$. We plug this into (19), thus obtaining the time evolution of the velocity field. Integrating the flow equation (20) will yield the evolution of the fingerprint $\psi(\theta, t)$. Performing the welding procedure will give us the time evolution of the corresponding shape. We consider a few cases below.

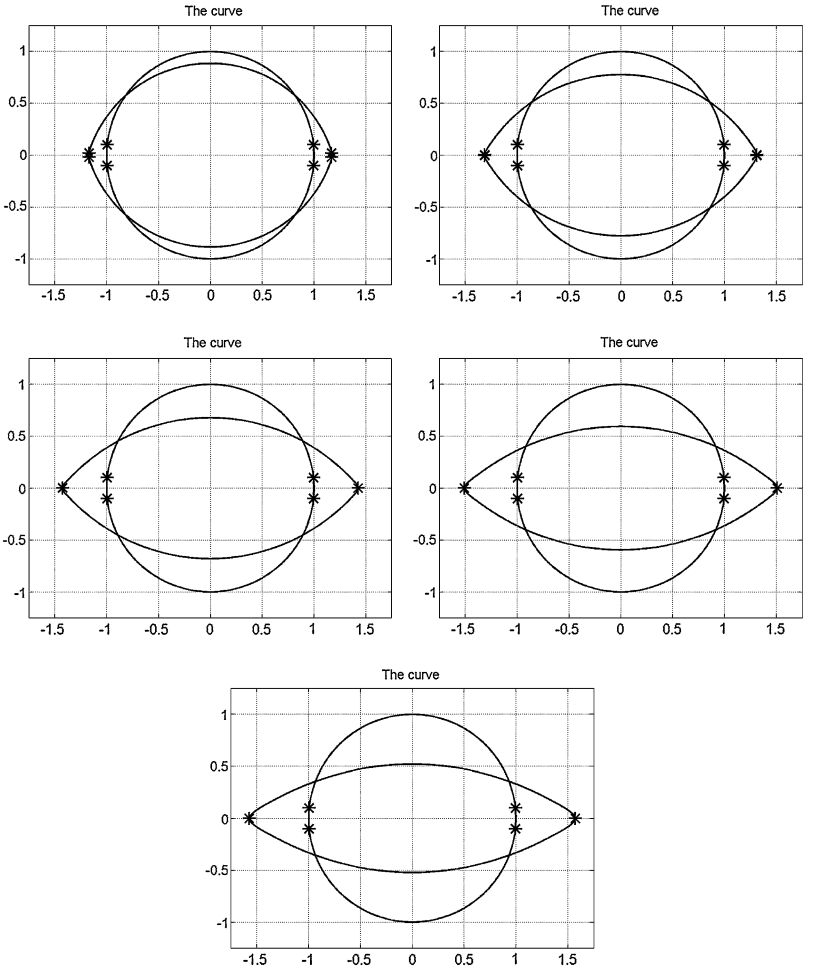

Figure 9: Final shapes of the 4-Teichon evolution for $d_0 = 0.2$ and $p_0 = 1, 2, 3, 4$ and 5. Asterisks indicate location of Teichons: on the circle at the initial time $t = 0$; on the final shape at time $t = 1$.

4.4.1. Case $N = 4, d_0 = 0.2$ The initial configuration that determines the velocity field is

$$p(t=0) = (p_0, -p_0, p_0, -p_0), \quad q(t=0) = \left(2\pi - \frac{d_0}{2}, \frac{d_0}{2}, \pi - \frac{d_0}{2}, \pi + \frac{d_0}{2}\right).$$

In Fig. 9 you can see an eye-like shape, that is the result of the evolution from the circle by the above velocity field, for $d_0 = 0.2$ and variable p_0. The estimates for the evolution of $p(t), d(t)$ and the shape were derived in [38].

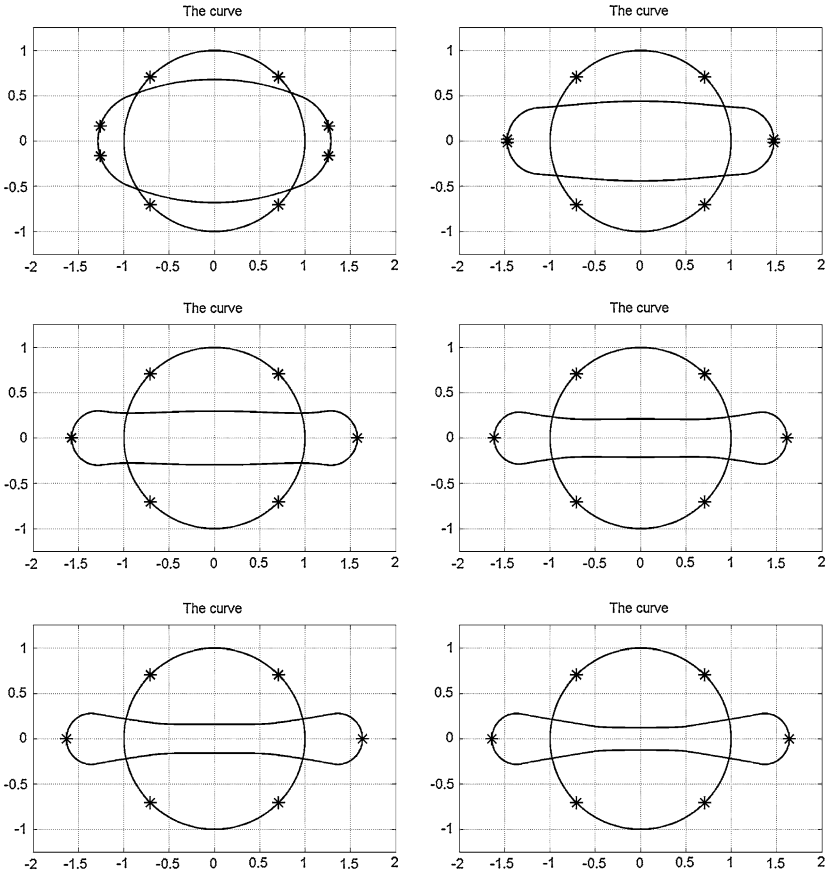

Figure 10: Final shapes at $T = 1$ of the 4-Teichon evolution for $d_0 = \pi/2$ and $p_0 = 0.5, 1, 1.5, 2, 2.5, 3$ (from left to right, from top to bottom). For $p_0 = 0$ the final shape is the initial circle. Asterisks indicate location of Teichons: on the circle at the initial time $t = 0$; on the final shape at time $t = 1$.

4.4.2. Case $N = 4, d = \pi/2$ In this case we consider the four 1-Teichons to be distributed equidistantly on the circle at $t = 0$. Initial configuration is

$$p(t = 0) = (p_0, -p_0, p_0, -p_0), q(t = 0) = (\pi/4, 3\pi/4, 5\pi/4, 7\pi/4).$$

In Fig. 10 you can see the range of shapes that one gets with the initial configuration of Teichons with $d_0 = \pi/2$ and p_0 ranging from 0.5 to 3. The integration scheme failed to compute evolution of $p(t), q(t)$ for $p_0 > 3$.

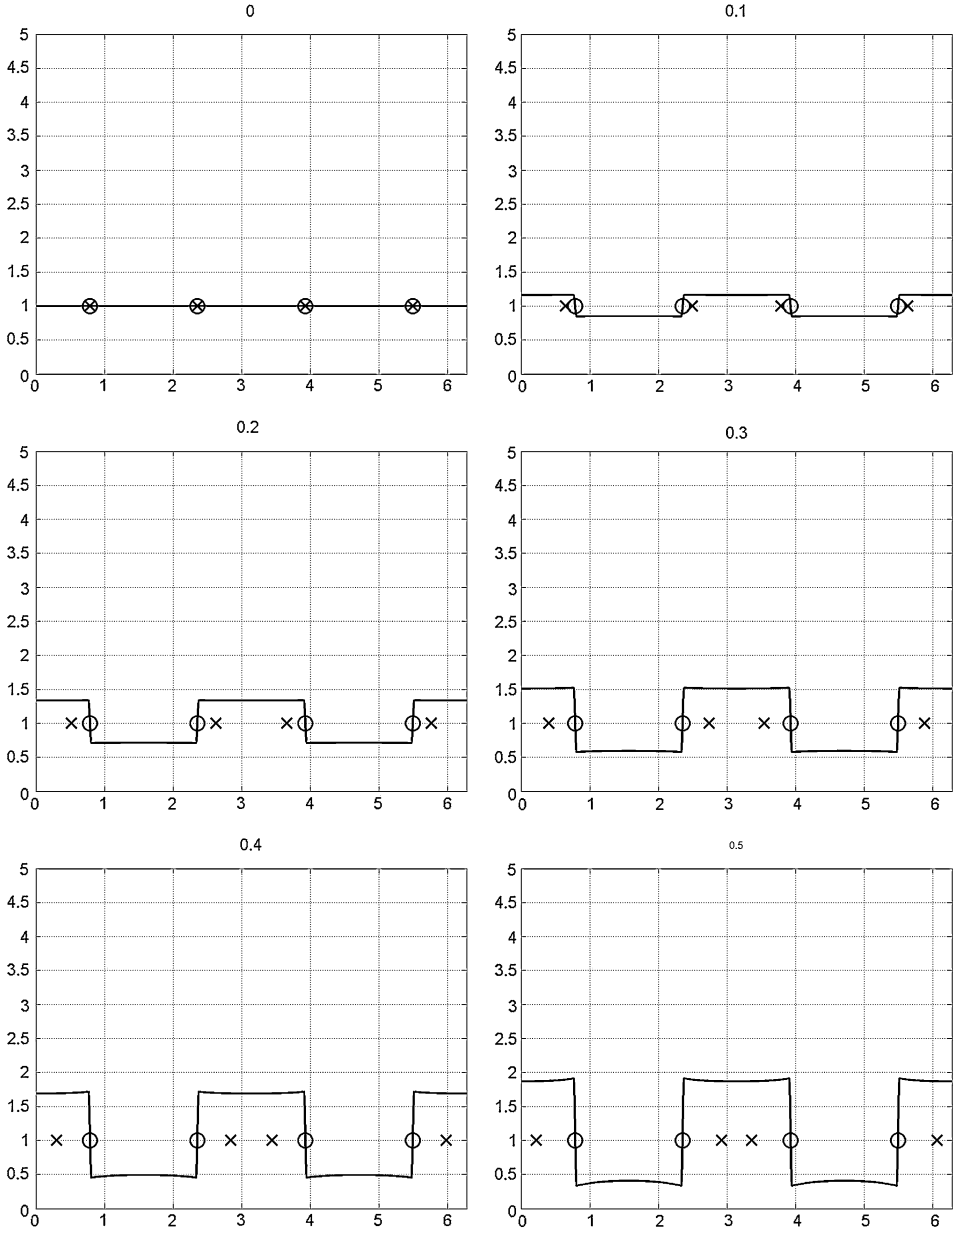

Figure 11: Snapshots of the evolution of the curvature of the shape, taken at times $t = k/10, k = 0, \ldots, 9$ with initial conditions $p_0 = 1, d_0 = \pi/2$. Crosses show positions of Teichons at a given time, circles show initial positions of Teichons.

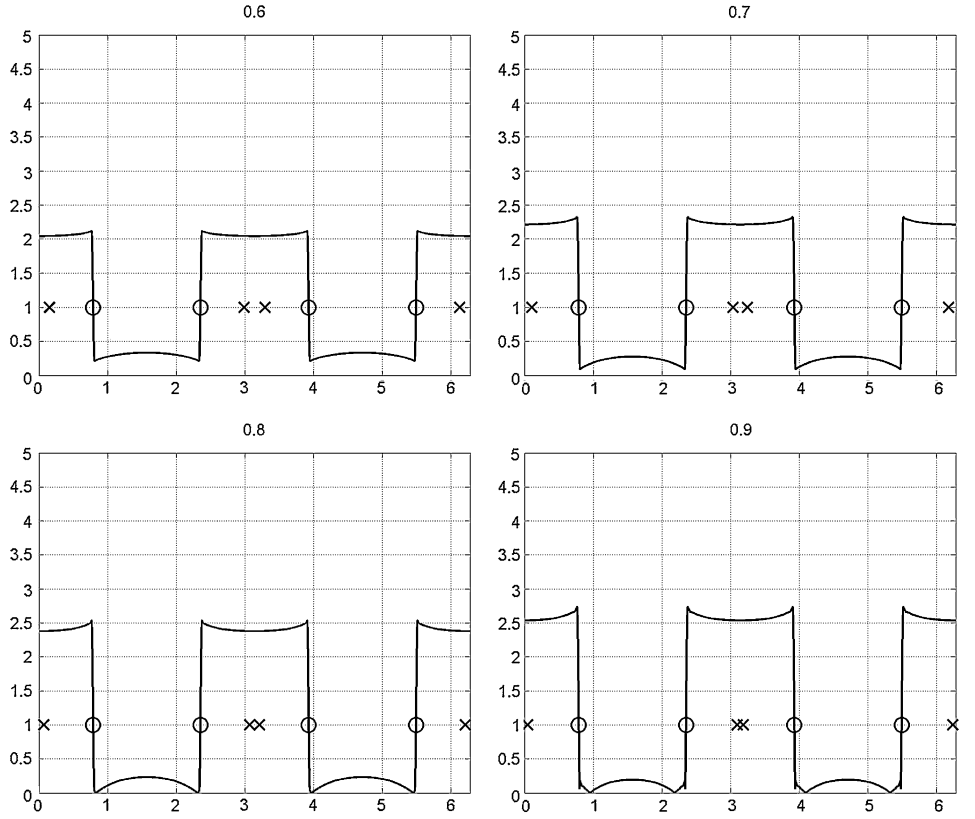

Figure 11: (Continued.)

In Fig. 11 we show the evolution of curvature from $t = 0$ to $t = 1$ for a configuration of Teichons where $d_0 = \pi/2, p_0 = 1$. Curvature has been estimated by fitting a second degree polynomial. Namely, assume that the curve is described by a finite sequence of points $\{m(k)\}_{k=1}^{N}$. We fix an approximation scale, an integer $D \geq 1$. For each k we are looking for three two-dimensional vectors $a(k), b(k), c(k)$ such that:

$$m(k+l) \approx a(k)\frac{l^2}{2} + b(k)l + c(k).$$

The curvature at a point $m(k)$ is then approximated via the formula:

$$\kappa(k) = \frac{\det[b(k), a(k)]}{|b(k)|^3}.$$

In the graphs in this section we have used approximation scale $D = 3$.

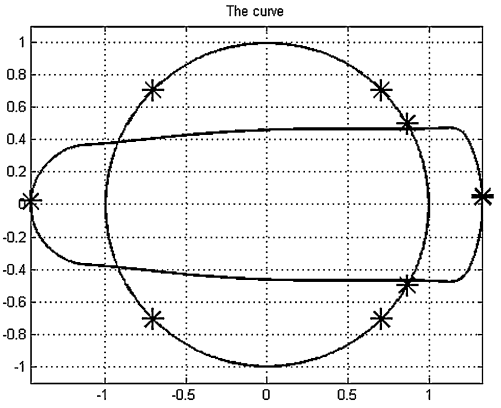

Figure 12: Final shape at $t = 1$ of a 6-Teichon. Asterisks indicate locations of Teichons: on the circle at the initial time $t = 0$; on the final shape at time $t = 1$.

It shows an interesting behavior of an N-Teichon solution: the resulting curve consists of arcs of the circles, at least for some initial time. If Teichons run into each other, then the curvature between them becomes larger, thus they produce a circle of the radius that gets smaller with time. On the intervals where Teichons are moving away from each other, the curvature decreases, thus producing circles of larger and larger radii.

4.4.3. Case $N = 6$ We consider an example of a 6-Teichon with

$$p(t = 0) = (2, -1.4142, 1.4142, -2, 1, -1),$$
$$q(t = 0) = \left(-\frac{\pi}{4}, -\frac{\pi}{6}, \frac{\pi}{6}, \frac{\pi}{4}, \frac{3\pi}{4}, \frac{5\pi}{4}\right).$$

The values of p_k were adjusted such that the constraints (18) are met. The resulting shape at $t = 1$ is presented in Fig. 12.

The curvature evolution for a 6-Teichon is shown in Figure 13. We observe a compound effect: Teichons that move toward each other produce increases in curvature, and the ones that are closer at the initial time produce higher curvature than the ones that are further apart. Notice how this is reflected in the shape: the right part of the shape has two sharp corners (higher curvature) connected by a less curved arc, while the opposite side has the same appearance as in the ellipse example. We derive the curvature evolution estimates in Section 4.5.

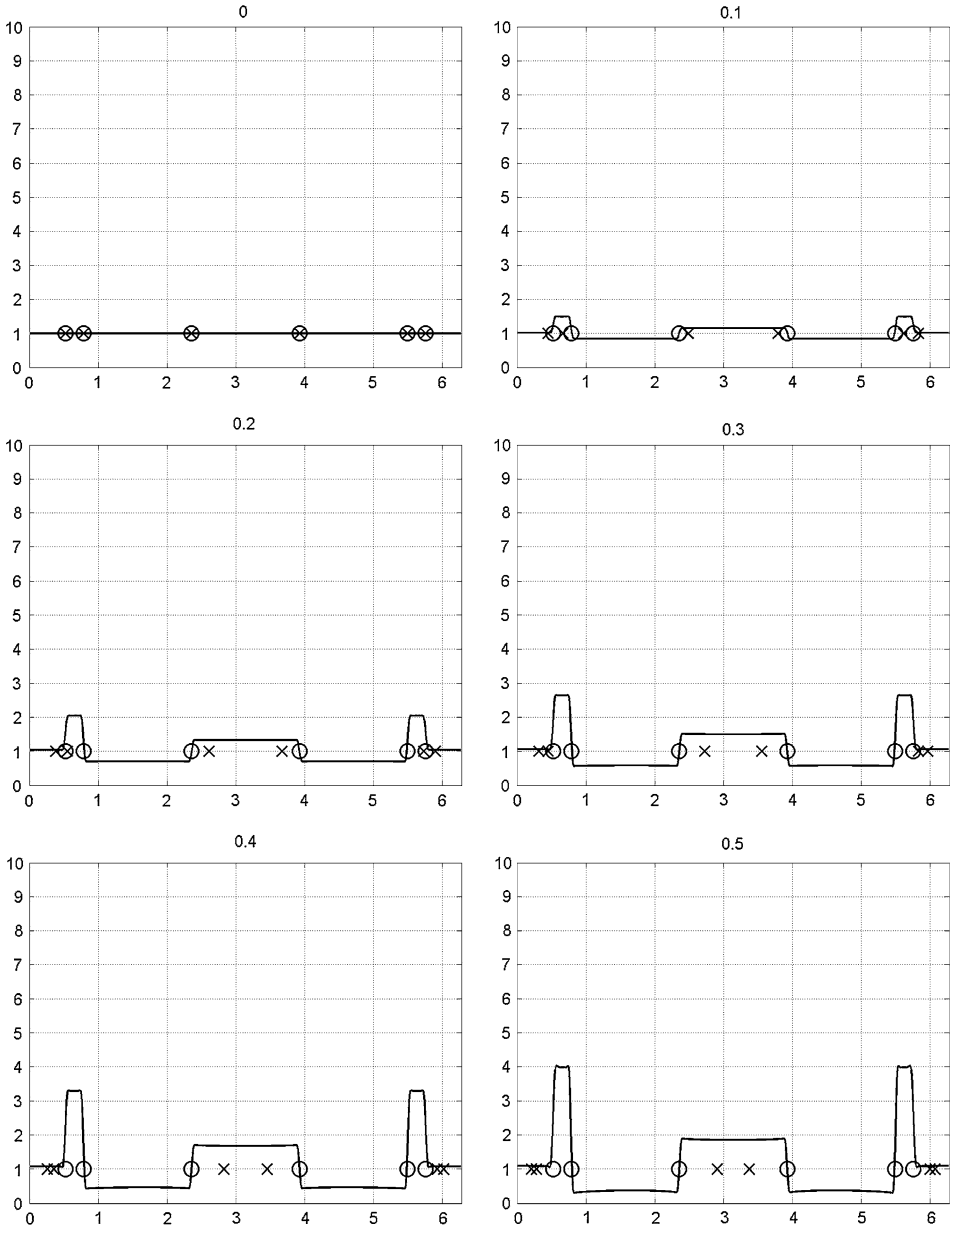

Figure 13: Snapshots of the evolution of the curvature of the shape from Fig. 12, taken at times $t = k/10, k = 0, 1, \ldots, 9$ for a 6-Teichon. Crosses show positions of Teichons at a given time, circles show initial positions of Teichons.

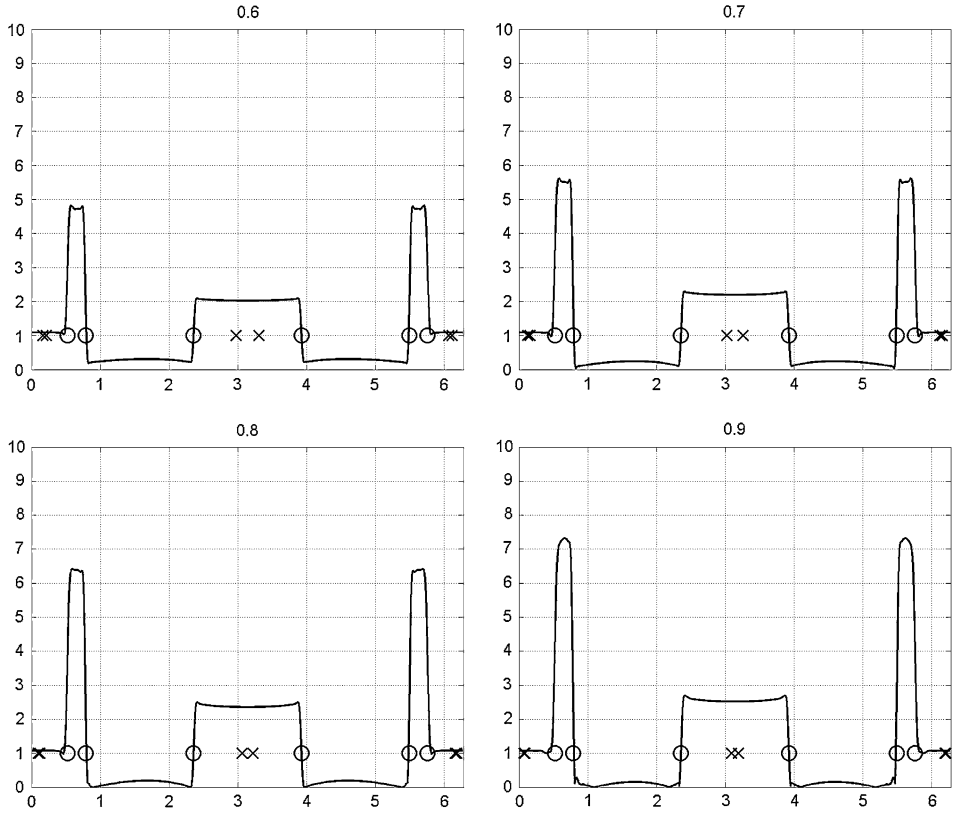

Figure 13: (Continued.)

4.5. Curvature evolution estimates

We will estimate the values of the curvature of the N-Teichon at the small initial time ε. Recall, that the N-Teichon velocity is given by $v(\theta, t) = \sum_{k=1}^{N} p_k(t)G(\theta - q_k(t))$. The flow equation for the fingerprint at time ε will be

$$\psi(\theta, \varepsilon) = \theta + \varepsilon v(\theta, \varepsilon) + o(\varepsilon)$$

It has been shown in [28] that the differential of the welding map at the origin (i.e. a unit circle) has the form

$$v(\theta) \mapsto \frac{1}{2}(iv(\theta) + \mathcal{H}v(\theta))e^{i\theta}\frac{\partial}{\partial z}.$$

In other words, if the fingerprint of the circle, $\psi(\theta) = \theta$, is perturbed by $\varepsilon v(\theta)$, then the resulting shape will be a perturbation of the circle $e^{i\theta}$:

$$r(\theta) = e^{i\theta} + \frac{\varepsilon}{2}(iv(\theta) + \mathcal{H}v(\theta))e^{i\theta} + o(\varepsilon).$$

Computing derivatives with respect to θ yields (writing only the linear terms):

$$r'(\theta) = ie^{i\theta} + \varepsilon\frac{e^{i\theta}}{2}\Big[\mathcal{H}R_1(\theta) + iR_1(\theta)\Big],$$
$$r''(\theta) = -e^{i\theta} + \varepsilon\frac{e^{i\theta}}{2}\Big[\mathcal{H}R_2(\theta) + iR_2(\theta)\Big].$$

Above, $R_1(\theta) = v'(\theta) + \mathcal{H}v(\theta)$ and $R_2(\theta) = v''(\theta) + 2\mathcal{H}v'(\theta) - v(\theta)$.

Treating r', r'' as vectors in $\mathbb{R}^2$ and using the fomula $\kappa = |r' \times r''|/\langle r', r'\rangle^{3/2}$ we get the curvature of the shape after time ε to be:

$$\kappa(\varepsilon) = 1 - \frac{\varepsilon}{2}\Big(\mathcal{H}v''(\theta) + \mathcal{H}v(\theta)\Big) + o(\varepsilon).$$

Recall, that the WP operator is $L = -\mathcal{H}(\partial_\theta^3 + \partial_\theta)$ and $LG(\theta) = 2\sum_{k=2}^{\infty}\cos(k\theta)$. Therefore

$$\begin{aligned} -\mathcal{H}(G''(\theta) + G(\theta)) &= 2\sum_{k=2}^{\infty}\frac{\sin(k\theta)}{k} \\ &= \operatorname{sign}(\theta)\pi - \theta - 2\sin\theta. \end{aligned}$$

The last equality is based on the sum of the series $\sum_k \sin(k\theta)/k$ [31]. Then, since $v = \sum p_k G(\theta - q_k)$, we get

$$\begin{aligned} -\mathcal{H}(v''(\theta) + v(\theta)) &= -\sum_{k=1}^{N} p_k\left[\mathcal{H}G''(\theta - q_k) + \mathcal{H}G(\theta - q_k)\right] \\ &= \sum_{k=1}^{N} p_k\left[\operatorname{sign}(\theta - q_k)\pi - \theta + q_k - 2\sin(\theta - q_k)\right] \\ &= \sum_{k=1}^{N} p_k\left[\operatorname{sign}(\theta - q_k)\pi + q_k\right]. \end{aligned}$$

The two terms in the second line in the above equations equal to zero because of the constraints (18). Finally, the curvature of the shape produced by a

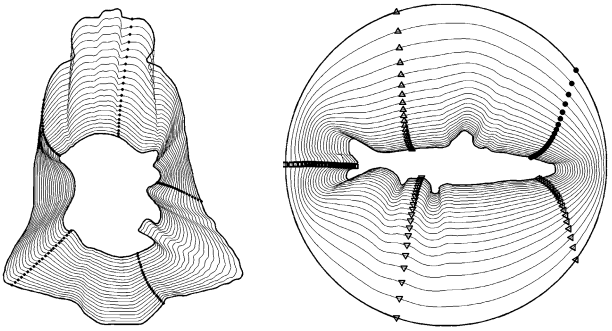

Figure 14: Examples of the geodesic computed by the shooting algorithm that uses N-Teichons to approximate WP geodesics [40]. Left: Shooting from a bell to a blob, two shapes in the MPEG-7 CE-Shape-1 database. Right: Shooting from a circle to a fish. An artificial scaling is employed to make the evolution clearer. Highlighted points show how the 4 specific points evolve on each shape.

Teichon for small enough initial time has the approximation:

$$\kappa(\theta) = 1 + \varepsilon/2 \sum_{k=1}^{N} p_k \left[\operatorname{sign}(\theta - q_k)\pi + q_k\right] + o(\varepsilon). \tag{21}$$

One can see that for small ε the curvature is a piecewise constant function of θ. Thus the Teichon evolution produces shapes that consist of piecewise circular arcs of various radii. From (21) one can see that the radii becomes smaller (i.e. curvature becomes larger) when two Teichons run into each other, for example when $q_1 < q_2$ and $p_1 = -p_2 > 0$. When Teichons' momenta point away from each other, the region between them experiences the decrease in curvature, thus producing circles of larger radii. This behavior is evident in Figs. 11 and 13.

In the future work we would like to investigate the behavior of the curvature for all time $t \in [0, 1]$. The current approximation of curvature can serve as a useful tool to make a better Teichons' initialization for the shooting method described in [40]. In Fig. 14 one can see examples of the geodesics computed by the shooting algorithm using N-Teichons.

4.6. Shooting with an N-Teichon

In this section we will briefly describe the shooting algorithm from [40]. The idea of the shooting method is the following: given the initial shape

(represented by fingerprint ψ_0) and target shape (represented by fingerprint ψ_1), start with an initial guess for the positions $q_k(0)$ and momenta (or 'strengths') $p_k(0)$ of Teichons, construct the initial momentum $m(\theta, 0) = \sum p_k(0)G(\theta - q_k(0))$, and then solve forward the equation (17) to obtain time evolution of $q_k(t), p_k(t)$. This in turn will produce a time-varying velocity field $v(\theta, t)$, which we integrate via the equation

$$v(\theta, t) = \frac{\partial \phi}{\partial t}(\phi^{-1}(\theta, t), t), \tag{22a}$$

$$\phi(\theta, t = 0) = \psi_0. \tag{22b}$$

to obtain $\phi(\theta, t = 1)$. The final computed fingerprint $\phi(\theta, t = 1)$, is compared with the target fingerprint, ψ_1, using the error term E. Based on this comparison, we modify the initial shot configuration $\{p_k(0), q_k(0)\}_{k=1}^N$ using the gradient of the error term E with respect to $\{p_k(0), q_k(0)\}_{k=1}^N$ and repeat the process. Because the fingerprints are members of a coset space $\mathrm{PSL}_2(\mathbb{R})\backslash\mathbf{Diff}(S^1)$ we cannot use an L^2 error function E. The suitable error function is based on complex cross-ratios of points on the fingerprint, since cross-ratios are invariant under $\mathrm{PSL}_2(\mathbb{R})$ transformation.

The computational cost is highly varied. On a desktop computer it takes 5-20mins to compute a geodesics between simple shapes (for example, two rotated ellipses). Shapes with that are less circular, with multiple concavities and convexities require coarse-to-fine type of algorithm and can take 6-8 hours. In the future work, we would like to incorporate the formula (21) into the initial guess for $\{p_k(0), q_k(0)\}_{k=1}^N$, thus cutting down the computational time to convergence.

5. Discussion

The first part of the paper discusses measures of massiveness (ε-entropy) of a C^1-type of metric on the space of planar curves and a related concept of coding and compression of shapes. The resulting optimal curve approximations are piecewise circular arcs with jump discontinuities. In the second part of the paper, variants of Sobolev type of metric are explored and they prove to be robust and rich shape descriptors for the purpose of finding similarity between planar curves. The paper concludes with a discussion of an $H^{3/2}$-like metric, the Weil-Petersson metric. In particular a special kind of singular solutions are suggested as a way to parametrize and describe a space of planar curves using a low-dimensional N-Teichon, thus providing a way for a finite-dimensional encoding of the 2D shape.

There are exciting prospects for these topics to be united. First, both the Teichon representation and the geometric Sobolev representation provide appealing finite-dimensional candidates for computing the ε-entropy of their respective shape spaces. Epsilon balls are easy to define in both cases, and in particular the geometric Sobolev machinery makes it relatively straightforward to describe the ε-balls in geometric terms.

Connections between approximation in these spaces are also of interest. Results in the compression section using tangent angle functions can be thought of as point-wise analogues of the β-angle approximations in the shape matching section. Bounds such as those in Lemma 2 of [2] give a natural entry point into discussions of approximation. In the C^1 case, a tangent angle function almost completely determines point location for the associated curve. It would be interesting to examine the extent to which this remains true in the H^α setting. With regard to the WP metric, connections are less straightforward, yet both approximation schemes result in curves consisting of piecewise circular arcs. It would be interesting to determine how the resulting approximations differ and why.

Another exciting prospect is to establish connections and bounds relating the length of the Weil-Petersson geodesic to the distance computed using angles $\beta(\theta, t)$ when working at the $H^{3/2}$ exponent. Indeed, β is known to be control the Schwarzian derivative of the conformal maps used to construct the Weil-Petersson welding map [10], and one can characterize the Weil-Petersson metric in terms of Schwarzian derivatives. β alone will not prevent the curve from self-intersecting, but it could be used to explain the local behaviour of the WP metric.

Further, denoising with the Sobolev-type metric could be an initial step before approximating the shape with Teichons: since Teichon approximation depends on the curvature, distinguishing between noise and shape features is a necessary step for any shape representation. In particular, $H^{3/2}$ is the critical exponent at which the boundary curve will admit a continuous tangent. Thus, a mostly smooth boundary arc which has even the slightest corner or cusp will lie at infinite distance from the origin (such a curve is not $H^{3/2}$). In practice, this means denoising the curves prior to computation is essential if one wishes to measure anything other than noise when computing geodesics.

References

[1] S. Belongie, J. Malik, and J. Puzicha. Shape matching and object recognition using shape contexts. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 24:509–522, 2001.

[2] M. Feiszli and P. W. Jones. Curve denoising by multiscale singularity detection and geometric shrinkage. *Appl. Comput. Harmon. Anal.*, 31:392–409, 2011. MR2836030

[3] H. Ling and D. Jacobs. Shape classification using the inner-distance. *IEEE Trans. Pattern Anal. Mach. Intell.*, 29(2):286–299, 2007.

[4] F. Mokhtarian, S. Abbasi, and J. Kittler. Efficient and robust retrieval by shape content through curvature scale space. In A. Smeulders and R. Jain, editors, *Image Databases and Multi-Media Search*, pages 51–58. World Scientific, 1997.

[5] E. M. Stein. *Singular integrals and differentiability properties of functions*. Princeton University Press, 1970. MR0290095

[6] E. M. Stein. The development of square functions in the work of a. zygmund. *Bull. Amer. Math. Soc.*, 7(2):359–376, 1982. MR0663787

[7] P. F. Felzenszwalb. Representation and detection of deformable shapes. *PAMI*, 27:208–220, 2004.

[8] P. W. Jones. Square functions, cauchy integrals, analytic capacity, and harmonic measure. *Springer Lecture Notes in Math.*, 1384:24–68, 1989. MR1013815

[9] P. W. Jones. Rectifiable sets and the traveling salesman problem. *Inventiones Mathematicae*, 102(1):1–15, 12 1990. MR1069238

[10] C. Bishop, P. Jones. Harmonic measure, L^2 estimates and the Schwarzian derivative. *Journal D'Analyse*, 62:77–113, 1994 MR1269200

[11] A. Efrat, S. Venkatasubramanian and Q. Fan, Curve matching, time warping, and light fields: New algorithms for computing similarity between curves, J. Mathematic Imaging and Vision, 2007 MR2325848

[12] V. V. Chepyzhov. *Attractors for equations of mathematical physics*, p. 163–197. American Mathematical Society, 2002. MR1868930

[13] D. L. Donoho. Counting bits with Kolmogorov and Shannon. Manuscript, 1998.

[14] D. L. Donoho, M. Vetterli, R. A. Devore, I. Daubechies. Data compression and harmonic analysis. Manuscript, 1998. MR1658775

[15] E. Klassen, A. Srivastava, W. Mio, and S. Joshi, Analysis of planar shapes using geodesic paths on shape spaces, 2003

[16] A. N. Kolmogorov and V. M. Tikhomirov. ε-entropy and ε-capacity. American Mathematical Society Translations, Series 2, **17**:277–364, 1959.

[17] G. Langdon, J. Rissanen. Compression of black and white images with arithmetic coding. IEEE Transactions in Communication, COMM-**29**:858–867, 1981.

[18] K. Leonard. Efficient shape modeling: ε-entropy, adaptive coding, and Blum's medial axis versus the boundary curve. Int. J. Comp. Vis., **74**, 2007, 183–199.

[19] W. Mio and A. Srivastava. Elastic-String Models for Representation and Analysis of Planar Shapes. Proceedings of the IEEE Computer Society International Conference on Computer Vision and Pattern Recognition (CVPR), Washington, DC, 2004.

[20] J. Rissanen. *Stochastic complexity in statistical inquiry*. World Scientific Press, 1989. MR1082556

[21] A. Pelczynski. On approximation of S-spaces by finite dimensional spaces. Bulletin Acad. Polon. Sci. Cl. III **5**:879–881, 1957. MR0092114

[22] E. C. Posner, E. R. Rodemich. ε-entropy and data compression. Annals of Mathematical Statistics, **42**:6: 2079–2125, 1971. MR0297458

[23] L. Younes. Computable Elastic Distance Between Shapes. SIAM Journal of Applied Mathematics, **58**, 565–586, 1998. MR1617630

[24] V. Arnold, Sur la géométrie différentielle des groupes de Lie de dimension infinie et ses applications à l'hydrodynamique des fluides parfaits, Annales de l'institut Fourier, 16 (1966), 319–361. MR0202082

[25] V. I. Arnold and B. A. Khesin, Topological Methods in Hydrodynamics, Springer, Apr. 1998. MR1612569

[26] M. J. Bowick and S. G. Rajeev, String theory as the Kähler geometry of loop space, Physical Review Letters, 58 (1987), 1158. MR0884852

[27] R. Camassa and D. D. Holm, An integrable shallow water equation with peaked solitons, Physical Review Letters, 71 (1993), 1661–1664. MR1234453

[28] M. Feiszli, Conformal shape representation, PhD thesis, Brown University, Providence, RI, May 2008. MR2717596

[29] O. Fringer and D. D. Holm, Integrable vs nonintegrable geodesic soliton behavior, Physica D, 150 (2001), 237–263. MR1820736

[30] F. Gay-Balmaz, J. E. Marsden, and T. S. Ratiu, The geometry of the Universal Teichmüller space and the Euler-Weil-Petersson equations, tech. report, 2009.

[31] I. Gradshteyn and R. I. M., *Table of integrals*, Academic, New York, 1980, 1980.

[32] U. Grenander and M. Miller., *Computational anatomy: an emerging discipline*, Quarterly of Applied Mathematics, LVI(4) (1998), 617–694. MR1668732

[33] D. D. Holm and J. E. Marsden, *Momentum maps and measure-valued solutions (peakons, filaments, and sheets) for the EPDiff equation*, in The Breadth of Symplectic and Poisson Geometry, J. E. Marsden and T. S. Ratiu, eds., vol. 232 of Progress in Mathematics, Birkhäuser Boston, 2005, 203–235. MR2103008

[34] D. D. Holm, J. Tilak R., A. Trouvé, and L. Younes, *Soliton dynamics in computational anatomy*, NeuroImage, 23, Supplement 1 (2004), S170–S178.

[35] J. H. Hubbard, *Teichmüller theory and applications to geometry, topology, and dynamics*, Matrix Editions, Ithaca, NY, 2006. MR2245223

[36] H. Karcher, *Riemannian center of mass and mollifier smoothing*, Comm. Pure Appl. Math, 30 (1977), 509–541. MR0442975

[37] B. Khesin and G. Misiołek, *Euler equations on homogeneous spaces and virasoro orbits*, Advances in Mathematics, 176 (2003), 116–144. MR1978343

[38] S. Kushnarev, *Teichons: Solitonlike geodesics on Universal Teichmüller space*, Experimental Mathematics, 18 (2009), 325–336. MR2555702

[39] S. Kushnarev, *The Geometry of the Space of 2D Shapes and the Weil-Petersson Metric*, PhD thesis, Brown University, Providence, RI, May 2010.

[40] S. Kushnarev and A. Narayan, *Approximating the Weil-Petersson geodesics on the universal Teichmüller space by singular solutions*, SIAM Journal on Imaging Sciences, accepted (2013).

[41] O. Lehto, *Univalent Functions and Teichmüller Spaces (Graduate Texts in Mathematics 109)*, Springer, 1 ed., Dec. 1986. MR0867407

[42] M. I. Miller, A. Trouvé, and L. Younes, *On metrics and Euler-Lagrange equations of computational anatomy*, Ann. Rev. Biomed. Engng, 4 (2002), 375–405.

[43] M. I. Miller, A. Trouvé, and L. Younes, *Geodesic shooting for computational anatomy*, Journal of Mathematical Imaging and Vision, 24 (2006), 209–228. MR2227097

[44] D. Mumford and A. Desolneux, *Pattern theory: the Stochastic Analysis of Real World Signals*, AK Peters Ltd, 2010. MR2723182

[45] S. Nag and A. Verjovsky, *Diff*(S^1) *and the Teichmüller spaces*, Communications in Mathematical Physics, 130 (1990), 123–138. MR1055689

[46] E. Sharon and D. Mumford, *2D-Shape analysis using conformal mapping*, International Journal of Computer Vision, 70 (2006), 55–75.

[47] L. A. Takhtajan and L.-P. Teo, *Weil-Petersson Metric on the Universal Teichmüller Space*, American Mathematical Society, Aug. 2006.

[48] A. Trouvé, *An infinite dimensional group approach for physics based model*, Technical report (electronically available at http://www.cis.jhu.edu), (1995).

[49] D. G. Kendall, *Shape Manifolds, Procrustean Metrics and Complex Projective Spaces*, Bull. of London Math. Soc., vol. 16, 81–121, (1984). MR0737237

[50] F. L. Bookstein, *Size and Shape Spaces for Landmark Data in Two Dimensions*, Statistical Science, vol. 1, 181–242, (1986).

[51] I. L. Dryden, K. V. Mardia, *Statistical Shape Analysis*, John Wiley and Son, (1998). MR1646114

[52] J. T. Kent, K. V. Mardia, *Shape, Procrustes Tangent Projections and Bilateral Symmetry*, Biometrika, vol. 88, 469–485, (2001). MR1844846

[53] T. B. Sebastian, P. N. Klein, and B. Kimia, *Recognition of Shapes by Editing Shock Graphs*, ICCV. Vol. 1. (2001).

[54] V. Caselles, R. Kimmel, and G. Sapiro, *Geodesic active contours*, International journal of computer vision 22.1, 61–79, (1997).

Matt Feiszli
13 Hine Place
New Haven, CT 06511
USA
E-mail address: mattfeiszli@gmail.com

Sergey Kushnarev
Engineering Systems & Design
Singapore University of Technology & Design
20 Dover Drive, 138682
Singapore
E-mail address: sergey.kushnarev@gmail.com

Kathryn Leonard
Mathematics Department
California State University Channel Islands
One University Drive
Camarillo, CA 93012
USA
E-mail address: kathryn.leonard@csuci.edu

Received August 4, 2013