Precise Matching of PL Curves in $\mathbb{R}^N$ in the Square Root Velocity Framework Part II

S. Lahiri, D. Robinson, E. Klassen, A. Srivastava

July 17, 2018

Goals of this talk

- Briefly review the SRVF (Square Root Velocity Function) method of putting a metric on the space of absolutely continuous curves in $\mathbb{R}^N$, in a way that is invariant under reparametrization.

Goals of this talk

- Briefly review the SRVF (Square Root Velocity Function) method of putting a metric on the space of absolutely continuous curves in $\mathbb{R}^N$, in a way that is invariant under reparametrization.
- Discuss the dense subspace of this metric space consisting of PL (Piecewise Linear) curves.

Goals of this talk

- Briefly review the SRVF (Square Root Velocity Function) method of putting a metric on the space of absolutely continuous curves in $\mathbb{R}^N$, in a way that is invariant under reparametrization.
- Discuss the dense subspace of this metric space consisting of PL (Piecewise Linear) curves.
- Demonstrate a method for computing precise geodesics between PL curves in the shape space.

Definition

A function $f : I \to \mathbb{R}^N$ is absolutely continuous (AC) if these two conditions hold:

- The derivative f' exists almost everywhere.

Definition

A function $f : I \to \mathbb{R}^N$ is absolutely continuous (AC) if these two conditions hold:

- The derivative f' exists almost everywhere.
- For all $t \in I$, $f(t) = f(0) + \int_0^t f'(u)du$.

Definition

A function $f : I \to \mathbb{R}^N$ is absolutely continuous (AC) if these two conditions hold:

- The derivative f' exists almost everywhere.
- For all $t \in I$, $f(t) = f(0) + \int_0^t f'(u)du$.

Note: The condition of absolute continuity is weaker than C^1, and much weaker than smoothness! Example: piecewise smooth curves are AC. Also, AC curves can be constant on subintervals of I.

Definition

A function $f : I \to \mathbb{R}^N$ is absolutely continuous (AC) if these two conditions hold:

- The derivative f' exists almost everywhere.
- For all $t \in I$, $f(t) = f(0) + \int_0^t f'(u)du$.

Note: The condition of absolute continuity is weaker than C^1, and much weaker than smoothness! Example: piecewise smooth curves are AC. Also, AC curves can be constant on subintervals of I.

Notation:

$$AC(I, \mathbb{R}^N) := \{\text{absolutely continuous functions } f : I \to \mathbb{R}^N\}$$

$$AC_0(I, \mathbb{R}^N) := \{\text{absolutely continuous } f : I \to \mathbb{R}^N : f(0) = 0\}$$

Bijection Between $AC_0(I, \mathbb{R}^N)$ and $L^2(I, \mathbb{R}^N)$:

Given $f \in AC_0(I, \mathbb{R}^N)$, define $q_f \in L^2(I, \mathbb{R}^N)$ as follows:

$$q_f(t) = \begin{cases} \frac{f'(t)}{\sqrt{|f'(t)|}} & \text{if } f'(t) \neq 0 \\ 0 & \text{if } f'(t) = 0. \end{cases}$$

Bijection Between $AC_0(I, \mathbb{R}^N)$ and $L^2(I, \mathbb{R}^N)$:

Given $f \in AC_0(I, \mathbb{R}^N)$, define $q_f \in L^2(I, \mathbb{R}^N)$ as follows:

$$q_f(t) = \begin{cases} \frac{f'(t)}{\sqrt{|f'(t)|}} & \text{if } f'(t) \neq 0 \\ 0 & \text{if } f'(t) = 0. \end{cases}$$

The mapping

$$AC_0(I, \mathbb{R}^N) \to L^2(I, \mathbb{R}^N)$$

given by

$$f \mapsto q_f$$

is easily verified to be a bijection. We often refer to q_f as the Square Root Velocity Function (SRVF) of f, or the q-function of f.

Since $L^2(I, \mathbb{R}^N)$ is a Hilbert space, it is a complete Riemannian manifold. Thus our bijection gives $AC_0(I, \mathbb{R}^N)$ the structure of a complete Riemannian manifold. The geodesics in $AC_0(I, \mathbb{R}^N)$ correspond to straight lines in $L^2(I, \mathbb{R}^N)$.

Since $L^2(I, \mathbb{R}^N)$ is a Hilbert space, it is a complete Riemannian manifold. Thus our bijection gives $AC_0(I, \mathbb{R}^N)$ the structure of a complete Riemannian manifold. The geodesics in $AC_0(I, \mathbb{R}^N)$ correspond to straight lines in $L^2(I, \mathbb{R}^N)$.

Remark 1:

Note that we are using this bijection to pull back all three structures from L^2 to AC_0 : topological, differentiable and Riemannian.

Remark 2:

This differential structure differs from the usual one on AC_0 in any neighborhood of a function $f \in AC_0(I, \mathbb{R}^N)$ with the property that $f'(t) = 0$ on a set of measure > 0.

Advantages of including absolutely continuous curves in our space of curves (instead of just immersions):

Advantages of including absolutely continuous curves in our space of curves (instead of just immersions):

1. We can now model a much larger variety of curves (for example, piecewise smooth).

Advantages of including absolutely continuous curves in our space of curves (instead of just immersions):

1. We can now model a much larger variety of curves (for example, piecewise smooth).
2. Our set of parametrized curves is now a complete metric space.

Advantages of including absolutely continuous curves in our space of curves (instead of just immersions):

1. We can now model a much larger variety of curves (for example, piecewise smooth).
2. Our set of parametrized curves is now a complete metric space.
3. Every pair of parametrized curves can now be joined by a geodesic.

Unit Length Curves

Define

$$S^\infty = \{f \in L^2(I, \mathbb{R}^N : \langle f, f \rangle = 1\}$$

S^∞ is the unit sphere in $L^2(I, \mathbb{R}^N)$. Under our bijection,

$$\{\text{AC curves of length } 1\} \leftrightarrow \text{unit sphere } S^\infty$$

S^∞ is a complete Riemannian manifold; geodesics in S^∞ are great circles. Thus, if we wish to mod out by the rescaling group, this is easily accomplished by normalizing all curves to have length 1, which means restricting our attention to the unit sphere S^∞.

We wish to consider two parametrized curves as "equivalent" if one is a reparametrization of the other; this is best understood by the action of the following reparametrization group.

We wish to consider two parametrized curves as "equivalent" if one is a reparametrization of the other; this is best understood by the action of the following reparametrization group.

Definition

The group Γ of reparametrizations is defined to be the set of functions $\gamma : I \to I$ satisfying

- $\gamma(0) = 0$ and $\gamma(1) = 1$.
- γ is absolutely continuous.
- $\gamma'(t) > 0$ almost everywhere.

These conditions imply that each $\gamma \in \Gamma$ is bijective and that its inverse is also in Γ.

Note that Γ acts on $AC_0(I, \mathbb{R}^N)$ from the right by composition.

Definition

The monoid $\tilde{\Gamma}$ of *singular reparametrizations* is defined to be the set of functions $\gamma : I \to I$ satisfying

- $\gamma(0) = 0$ and $\gamma(1) = 1$.
- γ is absolutely continuous.
- $\gamma'(t) \geq 0$ almost everywhere.

Because a function $\gamma \in \tilde{\Gamma}$ can be constant on subintervals of I, it will not in general have an inverse. Thus, $\tilde{\Gamma}$ is only a monoid.

The monoid $\tilde{\Gamma}$ also acts on $AC_0(I, \mathbb{R}^N)$ from the right by composition. Clearly,

$$\Gamma \subset \tilde{\Gamma}.$$

Corresponding Actions of Γ and $\tilde{\Gamma}$ on $L^2(I, \mathbb{R}^N)$

Since Γ and $\tilde{\Gamma}$ act on $AC_0(I, \mathbb{R}^N)$, and we have a bijection between $AC_0(I, \mathbb{R}^N)$ and $L^2(I, \mathbb{R}^N)$, it follows that there is an induced action on $L^2(I, \mathbb{R}^N)$. We denote that action by $q * \gamma$, where $q \in L^2(I, \mathbb{R}^N)$ and $\gamma \in \tilde{\Gamma}$ (or Γ); its formula is given as follows:

$$(q * \gamma)(t) = \sqrt{\gamma'(t)}q(\gamma(t)).$$

The actions of $\tilde{\Gamma}$ on $AC_0(I, \mathbb{R}^N)$ and on $L^2(I, \mathbb{R}^N)$ are related as follows:

$$q_{(f \circ \gamma)} = (q_f) * \gamma$$

for all $f \in AC_0(I, \mathbb{R}^N)$ and all $\gamma \in \tilde{\Gamma}$.

Theorem

The elements of $\tilde{\Gamma}$ (and of Γ) act on the Hilbert space $L^2(I, \mathbb{R}^N)$ as linear isometries, i.e., they preserve the Hilbert space structure. In symbols: $\langle q, w\rangle = \langle q * \gamma, w * \gamma\rangle$ for all $q, w \in L^2(I, \mathbb{R}^N)$ and for all $\gamma \in \tilde{\Gamma}$. Note: elements of $\tilde{\Gamma} - \Gamma$ act injectively, not bijectively.

This theorem is proved using integration by substitution. This is one important reason for our choice of reparametrization functions: absolutely continuous reparametrizations are precisely the functions for which integration by substitution is valid.

Theorem

The elements of $\tilde{\Gamma}$ (and of Γ) act on the Hilbert space $L^2(I, \mathbb{R}^N)$ as linear isometries, i.e., they preserve the Hilbert space structure. In symbols: $\langle q, w\rangle = \langle q * \gamma, w * \gamma\rangle$ for all $q, w \in L^2(I, \mathbb{R}^N)$ and for all $\gamma \in \tilde{\Gamma}$. Note: elements of $\tilde{\Gamma} - \Gamma$ act injectively, not bijectively.

This theorem is proved using integration by substitution. This is one important reason for our choice of reparametrization functions: absolutely continuous reparametrizations are precisely the functions for which integration by substitution is valid.

We have replaced the action of the reparametrization group (and monoid) on the space of curves, by an action by isometries on a complete metric space (in fact a Hilbert space).

Modding Out by Reparametrizations

Our primary goal is to make sense of the quotient of $L^2(I, \mathbb{R}^N)$ by the action of the group Γ, and to view this quotient as a metric space in its own right. Let $q\Gamma$ denote the orbit of q under Γ.

Natural attempt to put a metric on $L^2(I, \mathbb{R}^N)/\Gamma$

If $q, w \in L^2(I, \mathbb{R}^N)$, define $d(q\Gamma, w\Gamma) = \inf_{\tilde{q} \in q\Gamma, \tilde{w} \in w\Gamma} d(\tilde{q}, \tilde{w})$.

Modding Out by Reparametrizations

Our primary goal is to make sense of the quotient of $L^2(I, \mathbb{R}^N)$ by the action of the group Γ, and to view this quotient as a metric space in its own right. Let $q\Gamma$ denote the orbit of q under Γ.

Natural attempt to put a metric on $L^2(I, \mathbb{R}^N)/\Gamma$

If $q, w \in L^2(I, \mathbb{R}^N)$, define $d(q\Gamma, w\Gamma) = \inf_{\tilde{q} \in q\Gamma, \tilde{w} \in w\Gamma} d(\tilde{q}, \tilde{w})$.

Problem

The orbits $q\Gamma$ are not closed; hence there exist distinct orbits between which the infimum in the above formula is zero. Therefore, $L^2(I, \mathbb{R}^N)/\Gamma$ is not a metric space with respect to d.

Formal Solution

Define an equivalence relation $\sim$ on $L^2(I, \mathbb{R}^N)$ by

$$q \sim w \Leftrightarrow w \in Cl(q\Gamma)$$

where $CL(q\Gamma)$ denotes the L^2-closure of $q\Gamma$. The quotient space $(L^2(I, \mathbb{R}^N)/\sim)$ will then inherit a metric from $L^2(I, \mathbb{R}^N)$. Let $[q] = Cl(q\Gamma)$ denote the equivalence class of q under $\sim$.

Structure of the closed-up orbit $[q]$

Theorem

For every $q \in L^2(I, \mathbb{R}^N)$, $[q] = w\tilde{\Gamma}$, where w is the SRVF of the constant speed parametrization of the curve corresponding to q.

The theorem can be restated as: Two parametrized curves are equivalent if and only if both of them have the same constant speed parametrization.

Structure of the closed-up orbit $[q]$

Theorem

For every $q \in L^2(I, \mathbb{R}^N)$, $[q] = w\tilde{\Gamma}$, where w is the SRVF of the constant speed parametrization of the curve corresponding to q.

The theorem can be restated as: Two parametrized curves are equivalent if and only if both of them have the same constant speed parametrization.

Definition

Define the shape space by $\mathcal{S}(I, \mathbb{R}^N) = L^2(I, \mathbb{R}^N)/\sim$.

Note: $\mathcal{S}(I, \mathbb{R}^N)$ is a complete metric space, but not a manifold.

Definition

An *optimal matching* of a pair $q_1, q_2 \in L^2(I, \mathbb{R}^N)$ is a pair $w_1 \in [q_1]$ and $w_2 \in [q_2]$ such that $d(w_1, w_2) = d([q_1], [q_2])$.

Definition

An *optimal matching* of a pair $q_1, q_2 \in L^2(I, \mathbb{R}^N)$ is a pair $w_1 \in [q_1]$ and $w_2 \in [q_2]$ such that $d(w_1, w_2) = d([q_1], [q_2])$.

Fundamental Question about $\mathcal{S}(I, \mathbb{R}^N)$:

Given $[q_1]$ and $[q_2]$ in $\mathcal{S}(I, \mathbb{R}^N)$, under what circumstances does there exist an optimal matching between q_1 and q_2? Note: Whenever such an optimal matching exists, there also exists a geodesic (in the metric space sense) joining $[q_1]$ and $[q_2]$ in the shape space $\mathcal{S}(I, \mathbb{R}^N)$.

M. Bruveris[1] has some nice recent results on this question. He has shown that if the two curves are both C^1, then an optimal matching exists. He has also produced a pair of very "badly behaved" AC curves for which no optimal matching exists.

S. Lahiri et al.[2] have taken a different approach, in which one or both curves is piecewise linear. If at least one of them is PL, this paper proves that an optimal matching exists. If both are PL, they prove that the optimal matching exists and is PL and provide an algorithm for the precise computation of this matching.

[1] M. Bruveris, *Optimal reparametrizations in the square root velocity framework*, SIAM J. Math. Anal., 2016.

[2] S. Lahiri, D. Robinson, E. Klassen, *Precise matching of PL curves in $\mathbb{R}^N$ in the square root velocity framework*, Geometry, Imaging and Computing, 2015.

We will now describe the algorithm of Lahiri et al. During the development of this algorithm, Lahiri implemented it in matlab and used her code to produce examples for the paper, but never made the code itself public. We have been informed that M. Bruveris and A. Salili have written code for the same algorithm, and made it available on github.

Definition

A function $q \in L^2(I, \mathbb{R}^N)$ is a step function if there is a finite partition $0 = s_0 < t_1 < t_2 < \cdots < s_n = 1$ of I such that q is constant on each open interval (q_{i-1}, q_i).

A Computationally Useful Subspace of $\mathcal{S}(I, \mathbb{R}^N)$

Observation: $f \in AC_0(I, \mathbb{R}^N)$ is PL $\Leftrightarrow$ $q_f \in L^2(I, \mathbb{R}^N)$ is a step function.

Definition

$\mathcal{S}_{st}(I, \mathbb{R}^N) := \{[q] \in \mathcal{S}(I, \mathbb{R}^N) : [q]$ contains a step function$\}$

Thus, $\mathcal{S}_{st}(I, \mathbb{R}^N)$ is the subset of $\mathcal{S}(I, \mathbb{R}^N)$ corresponding to curves that admit PL parametrizations.

A Computationally Useful Subspace of $\mathcal{S}(I, \mathbb{R}^N)$

Observation: $f \in AC_0(I, \mathbb{R}^N)$ is PL $\Leftrightarrow$ $q_f \in L^2(I, \mathbb{R}^N)$ is a step function.

Definition

$\mathcal{S}_{st}(I, \mathbb{R}^N) := \{[q] \in \mathcal{S}(I, \mathbb{R}^N) : [q] \text{ contains a step function}\}$

Thus, $\mathcal{S}_{st}(I, \mathbb{R}^N)$ is the subset of $\mathcal{S}(I, \mathbb{R}^N)$ corresponding to curves that admit PL parametrizations.

Theorem

$\mathcal{S}_{st}(I, \mathbb{R}^N)$ is dense in $\mathcal{S}(I, \mathbb{R}^N)$.

Proof: Elementary measure theory.

Theorem

Given $[q], [w] \in \mathcal{S}(I, \mathbb{R}^N)$:

1. *if at least one of them is in $\mathcal{S}_{st}(I, \mathbb{R}^N)$, then an optimal matching exists;*
2. *if both of them are in $\mathcal{S}_{st}(I, \mathbb{R}^N)$, then this optimal matching can be taken to consist of step functions and the geodesic between them lies entirely in $\mathcal{S}_{st}(I, \mathbb{R}^N)$;*
3. *there is a finite combinatorial algorithm that computes this optimal matching and the corresponding geodesic.*

For $N = 1$, a simpler version of this algorithm had been discovered earlier by D. Robinson and implemented in his 2012 PhD dissertation at FSU. A few years later it was published[3] in 2017. For $N \geq 1$, it was implemented and published by S. Lahiri, in a paper already cited. The $N \geq 1$ version is more intricate and slower computationally.

More recently, M. Bruveris and A. Salili have implemented this algorithm for $N \geq 1$ and made it publicly available on github.

[3] Robinson, Duncan, Srivastava, Kassen, *Exact Function Alignment Under Elastic Riemannian Metric*, in Graphs in Biomedical Image Analysis, Computational Anatomy and Imaging Genetics, Lecture Notes in Computer Science, vol 10551, Springer, 2017

Idea of proof of (1) for general case $N \geq 1$

Case 1: Suppose $w(t) \equiv w_0$ is a constant map; this is equivalent to assuming that one of the two curves in a straight line. Then we can write an explicit formula for a reparametrization γ that maximizes $\langle q, w * \gamma \rangle$. The proof of maximality is elementary, using only the Cauchy-Schwarz inequality.

Case 2: Suppose w is a step function. Then once we decide which parameter value of q to match to each change-point of w, Case 1 determines the optimal reparametrization of each linear piece of w. But the space of all of these choices is the compact finite dimensional simplex $0 = s_0 \leq s_1 \leq \cdots \leq s_n = 1$. Since the distance function is continuous, a minimum distance must be achieved.

Setup for Precise Matching of PL Curves in $\mathbb{R}^N$

Given:

Let f_1 and f_2 be continuous PL curves in $\mathbb{R}^N$; let q_1 and q_2 be their SRVFs (which will be step functions).

Goal:

Find reparametrizations $\gamma_1, \gamma_2 \in \tilde{\Gamma}$ which minimize $d(q_1 * \gamma_1, q_2 * \gamma_2)$. Because $\tilde{\Gamma}$ acts by isometries, this is the same as maximizing $\langle q_1 * \gamma_1, q_2 * \gamma_2 \rangle$.

Then $(q_1 * \gamma_1, q_2 * \gamma_2)$ will be an optimal matching of (q_1, q_2), and the straight line between $q_1 * \gamma_1$ and $q_2 * \gamma_2$ in $L^2(I, \mathbb{R}^n)$ will yield a geodesic in $\mathcal{S}(I, \mathbb{R}^N)$ between $[q_1]$ and $[q_2]$.

Description of Algorithm

Choose partitions:

$$0 = s_0 < s_1 < \cdots < s_m = 1 \text{ and } 0 = t_0 < t_1 < \cdots < t_n = 1$$

so that

- f_1 is linear on each $[s_{i-1}, s_i]$; (hence q_1 is constant on each (s_{i-1}, s_i))
- f_2 is linear on each $[t_{j-1}, t_j]$; (hence q_2 is constant on each (t_{j-1}, t_j))

For each i and j, define $u_i = q_1(s_{i-1}, s_i)$ and $v_j = q_2(t_{j-1}, t_j)$

For each i, j, let $W_{ij} = u_i \cdot v_j$. We call $\{W_{ij}\}$ the weight matrix of q_1 and q_2. We subdivide the square $I \times I$ into rectangular blocks $G_{ij} = [s_{i-1}, s_i] \times [t_{j-1}, t_j]$. We assign to each block G_{ij} the real weight W_{ij}.

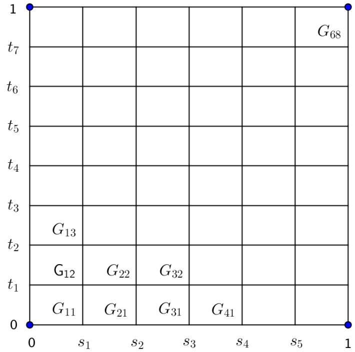

Our optimal matching $\gamma = (\gamma_1, \gamma_2)$ is a parametrized path in $I \times I$, from $(0, 0)$ to $(1, 1)$. Because γ_1 and γ_2 are weakly increasing, the direction of this path must always be towards the upper right – i.e., its slope must always be an element of $[0, \infty]$.

Our matching algorithm for a pair of PL curves is based on the following theorem that we have proved, concerning certain laws that an optimal γ must obey as it passes through the various blocks G_{ij}.

Theorem

An optimal matching γ between two PL curves can be parametrized so that:

1. *it is PL;*
2. *it consists of a sequence of P-segments and N-segments, with no two consecutive N-segments;*
3. *it satisfies certain inequalities relating the final slope of a P-segment to the initial slope of the next P-segment.*

Definition

A P-segment

- starts at a vertex, ends at a vertex, and passes through no other vertices;
- is linear as it passes through any single block.
- The initial and final blocks that it passes through must have positive weights, and the slope of γ in the initial and final blocks must lie in $(0, \infty)$, i.e., cannot be vertical or horizontal.
- It is either vertical or horizontal whenever it passes through a block with weight ≤ 0.
- When γ passes through an edge from one block to another, the change in slope is determined by the weights of the two blocks involved.

As a result of the last clause of this definition, a P-segment is determined by its slope as it passes through its initial block.

Example of a P-segment:

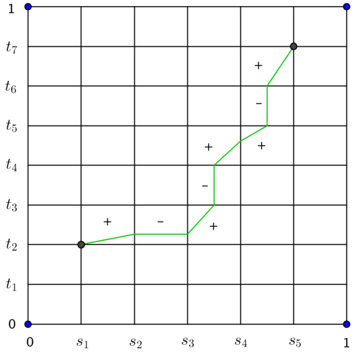

Slope Change Rule for passing though a Vertical Edge

Assume $W_{i,j}$ and $W_{i+1,j}$ are both positive. Suppose a P-segment passes through a vertical edge from $G_{i,j}$ to $G_{i+1,j}$. Let $H_{i,j}$ and $H_{i+1,j}$ denote the slopes of the P-segment in these blocks. Then the required relation is:

$$H_{i+1,j} = \left(\frac{W_{i+1,j}}{W_{i,j}} \right)^2 H_{i,j}$$

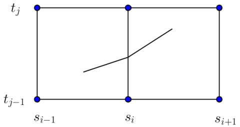

Slope Change Rule for passing though a Horizontal Edge

Assume $W_{i,j}$ and $W_{i,j+1}$ are both positive. Suppose a P-segment passes through a horizontal edge from $G_{i,j}$ to $G_{i,j+1}$. Let $H_{i,j}$ and $H_{i,j+1}$ denote the slopes of the P-segment in these blocks. Then the required relation is:

$$H_{i,j+1} = \left(\frac{W_{i,j}}{W_{i,j+1}}\right)^2 H_{i,j}$$

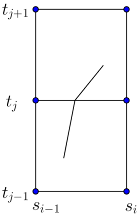

Definition

An N-segment

- starts at a vertex and ends at a vertex;
- has the property that the rectangle spanned by the beginning and ending vertices contains only blocks with weights ≤ 0; also some of the adjoining blocks must have this property;
- consists entirely of horizontal and vertical segments.

Examples of three N-segments from (s_1, t_2) to (s_4, t_6):

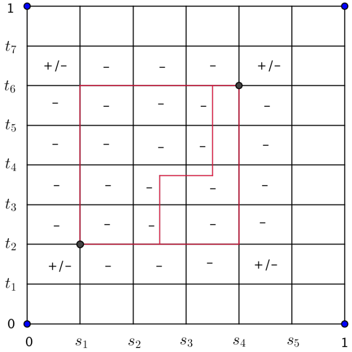

Relation between final slope of a P-segment, and initial slope of the next P-segment: Let $H_{i,j} > 0$ be the final slope of a P-segment, and $H_{i+1,j+1} > 0$ be the initial slope of the next P-segment. Then we must have the following relationship:

$$H_{i+1,j+1} = \mu H_{i,j}$$

where

$$\mu \in \left[\frac{D^2}{AB}, \frac{AB}{C^2}\right].$$

A, B, C, and D denote the weights of the relevant blocks as shown in the following diagram.

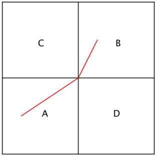

The formula given is for the case in which the weights A, B, C, and D are all positive. If C and/or D is not positive, then the bounds in the μ-interval will be replaced by 0 and/or ∞. If this interval is empty, then no optimal matching can pass through this vertex! This μ-interval is important computationally, because it reduces the number of P-segments we must search over for each vertex. In fact, the greater the number the sample points, the narrower this interval tends to become.

Algorithm for Producing Optimal Matching Between PL Curves:

Our algorithm is formally similar to Dynamic Programming. We examine each vertex (s_i, t_j) in the grid, starting with $(0, 0)$ and proceeding left to right along each row, covering all the rows from bottom to top. By the time we arrive at a vertex, we have determined the best allowable path from $(0, 0)$ to that vertex, so we then examine all allowable P-segments or N-segments beginning at the new vertex, keeping track of where each one ends and of the contribution of the new segment. In this manner, by the time we reach the final vertex $(1, 1)$, we have determined the optimal path from $(0, 0)$ to $(1, 1)$. What makes this algorithm possible is that minimizing the distance $d(q_1 * \gamma_1, q_2 * \gamma_2)$ is equivalent to maximizing the inner product $\langle q_1 * \gamma_1, q_2 * \gamma_2 \rangle$, which is simply an integral along the parameter space and, hence, is additive along the segments of the path.

Interesting Fact: If the optimal matching between a pair of curves includes at least one N-segment, then there are infinitely many non-equivalent optimal matchings, obtained by substituting any other N-segment (composed entirely of vertical and horizontal segments) between the same two vertices. As a result, there are infinitely many geodesics between these two curves in the shape space! This is analogous to the fact that two antipodal points on a sphere can be joined by an infinite number of geodesics.

Examples of Optimal Matchings between functions $I \to \mathbb{R}^1$: In each 1-dimensional example, we show the graphs of the unaligned functions on the left, the optimally aligned functions on the right, and we show the optimal matching $\gamma = (\gamma_1, \gamma_2)$ below.

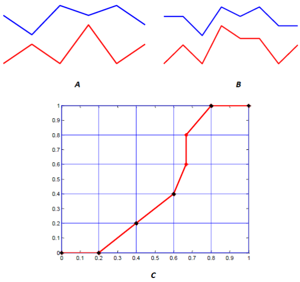

Figure: Example $1(1D)$. Distance before alignment is 1.4815. Distance after alignment is 0.5071.

Another 1-Dimensional Example:

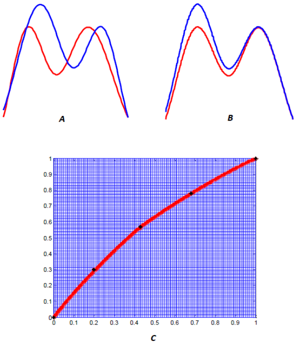

Figure: Example 2($1D$). Distance before alignment is 1.4312. Distance after alignment is 0.1195.

2-Dimensional Examples: On the left, we show the optimal matching between the curves; on the right, the shortest geodesic between the curves, and below, we show the matching function $\gamma = (\gamma_1, \gamma_2)$.

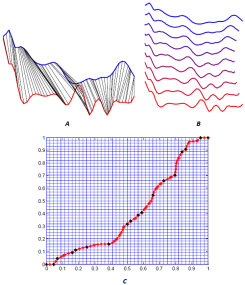

Figure: Example 3(2D). Distance before alignment is 7.0108. Distance after alignment is 4.0721.

Another 2D Example:

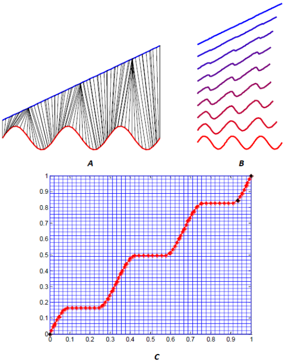

Figure: Example 4($2D$). Distance before alignment is 3.9107. Distance after alignment is 2.8418.

Another 2D Example: (semi-circles traversed in opposite directions; includes an N-segment)

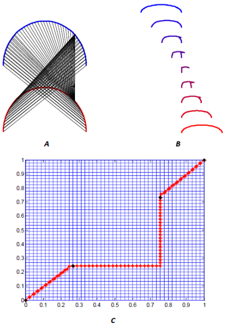

Figure: Example 5(2D). Distance before alignment is 2.5064. Distance after alignment is 2.0683.

Another 2D Example:

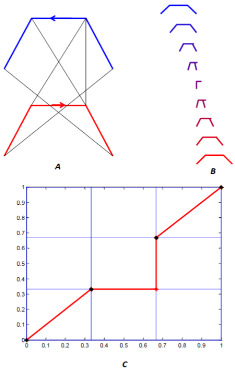

Figure: Example 6(2D). Distance before alignment is 2.4495. Distance after alignment is 2.

Dan Robinson's results on $\mathcal{S}_{st}(I, R)$

Given $f \in AC_0(I, R)$ with unit arclength,

$$[q_f] \in \mathcal{S}_{st}(I, R)$$
$$\Updownarrow$$

There is a finite partition $I = [0, t_1] \cup [t_1, t_2] \cup \ldots [t_{n-1}, 1]$ such that on each $[t_{i-1}, t_i]$, f is weakly monotonic.

Robinson has constructed and implemented a combinatorial algorithm that performs the following: Given $f, g \in AC_0(I, R)$ with unit arclength, and assuming both $[q_f], [q_g] \in \mathcal{S}_{st}(I, R)$, his algorithm precisely determines a pair $\tilde{q}_f \in [q_f]$ and $\tilde{q}_g \in [q_g]$ such that $d(\tilde{q}_f, \tilde{q}_g) = d([q_f], [q_g])$.

Stated informally, Dan's algorithm determines the precise matching between the domain of f and the domain of g that minimizes the distance between their $SRVF$'s. It should be noted that "usually" both $\tilde{q}_f$ and $\tilde{q}_g$ involve reparametrizations by elements of $\tilde{\Gamma}$, not just Γ.

The next frame contains the simplest non-trivial example of functional alignment using the SRVF method. We give two functions $f(t)$ and $g(t)$, and then their optimally aligned versions. (In this case we only need to alter f, not g.) Note that we reparametrize f to give it a stationary point during the parameter interval in which g is going the "opposite direction."

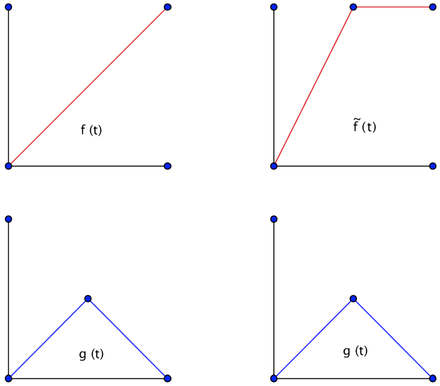

Original functions f(t) and g(t)
distance = $\pi/2$

Optimally reparametrized functions $\tilde{f}$(t) and g(t)
distance = $\pi/4$

The following more complicated example gives a better demonstration of Robinson's method. In it, we give

- two random functions f_1 and f_2;
- PL functions $\tilde{f}_1 \in [f_1]$ and $\tilde{f}_2 \in [f_2]$ that minimize $d(q_{\tilde{f}_1}, q_{\tilde{f}_2})$;
- the geodesic between $\tilde{f}_1$ and $\tilde{f}_2$;
- a reparametrization of this geodesic that approximates the original f_1. (It is impossible to make it match exactly, since the relevant reparametrizations are in $\tilde{\Gamma}$, hence are not invertible.)

It is a characteristic of these optimal matchings that a decreasing portion of one function is never matched against an increasing portion of the other.

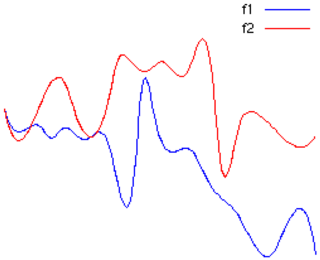

Original functions

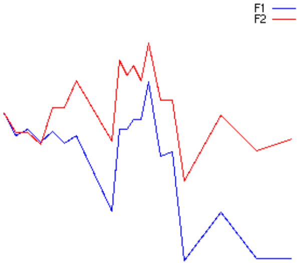

Matched PLFs

$\mathbb{L}^2$ product of SRVFs: 0.92387

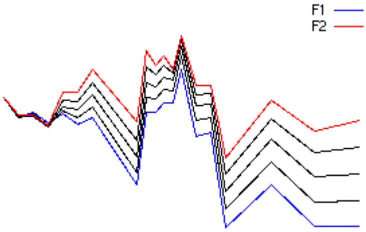

Geodesic between PLFs

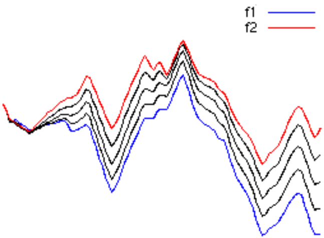

Geodesic from original f1

Multiple alignment of functions: Berkeley Growth Rate Curves

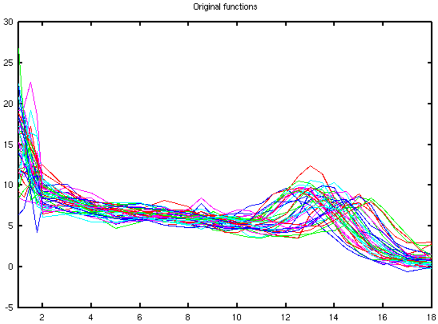

Figure: Original growth rate curves.

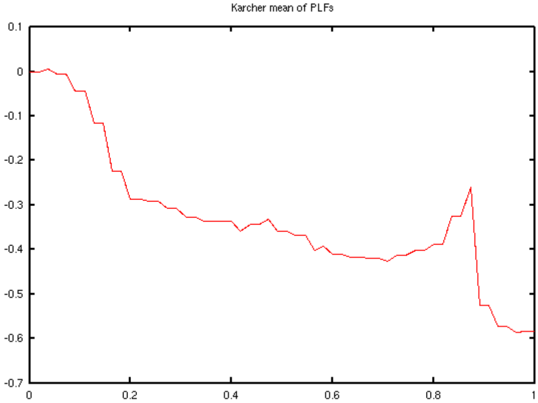

Figure: PL parametrization of Karcher mean.

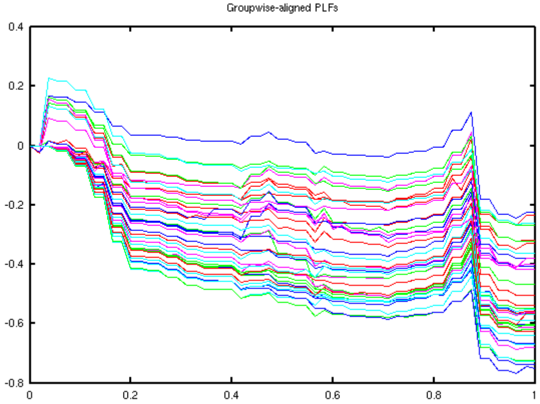

Figure: PL reparametrizations of original functions optimally aligned to match mean.

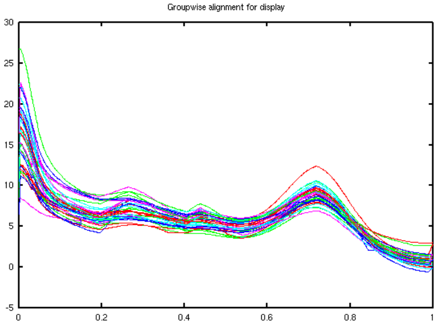

Figure: Aligned functions reparametrized by inverse of average of the reparametrizations.

Comments on the geometry of $\mathcal{S}_{st}(I, \mathbb{R})$

- $\mathcal{S}_{st}(I, \mathbb{R})$ is dense in $\mathcal{S}(I, \mathbb{R})$.
- $\mathcal{S}_{st}(I, \mathbb{R})$ has the structure of a 'CW complex', with two cells in each dimension (but not weak topology).
- For each n, the n-skeleton of $\mathcal{S}_{st}(I, \mathbb{R})$ is homotopy equivalent to S^n.
- Given each $[q], [w] \in \mathcal{S}_{st}(I, \mathbb{R})$, there exist a finite number of geodesics between them, and Robinson's algorithm determines all of them combinatorially.
- $\mathcal{S}_{st}(I, \mathbb{R})$ is a very interesting metric space!

A pair of points in $\mathcal{S}_{st}(I, \mathbb{R})$ that can be joined by two different shortest geodesics

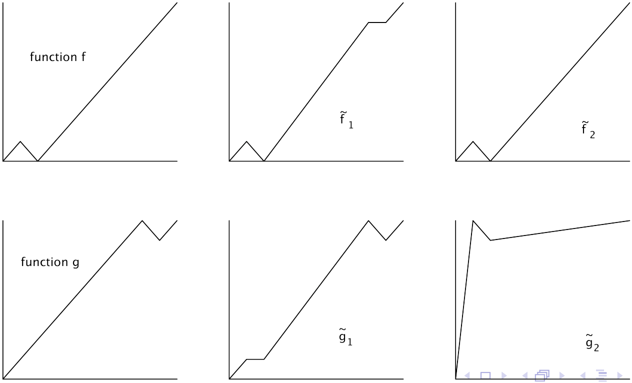

References

M. Bauer, M. Bruveris, and P. Michor. *R-Transforms for Sobolev H^2 metrics on spaces of plane curves*, Geometry, Imaging and Computing, 1(1):1-56, 2014.

S. Lahiri, D. Robinson, and E. Klassen, *Precise Matching of PL Curves in $\mathbb{R}^N$ in the Square Root Velocity Framework*, arXiv:1501.00577, 2015.

W. Mio, A Srivastava, S Joshi, *On Shape of Plane Elastic Curves*, IJCV, 2007.

Daniel Robinson, *Functional Data Analysis and Partial Shape Matching in the Square Root Velocity Framework*, PhD Dissertation, FSU, 2012.

References

A. Srivastava, E. Klassen, S. Joshi, and I. Jermyn, *Shape Analysis of Elastic Curves in Euclidean Spaces*, IEEE Transactions on Pattern Analysis and Machine Intelligence , vol. 33, issue 7, pages 1415-1428, July 2011.

L. Younes, P. Michor, J. Shah, and D. Mumford, *A Metric on Shape Space with Explicit Geodesics*, Matematica E Applicazioni, 19(1): 25-57, 2008.

Thank You!!