Signatures in Shape Analysis

Nikolas Tapia (WIAS/TU Berlin) joint with E. Celledoni & P. E. Lystad (NTNU)

FG6

Shape Analysis: Practical Setup

The problem is to find a similarity measure (metric) between shapes that is:

1. accurate enough, in that it distinguishes different types of motion (clustering), and
2. easy (and fast) to compute.

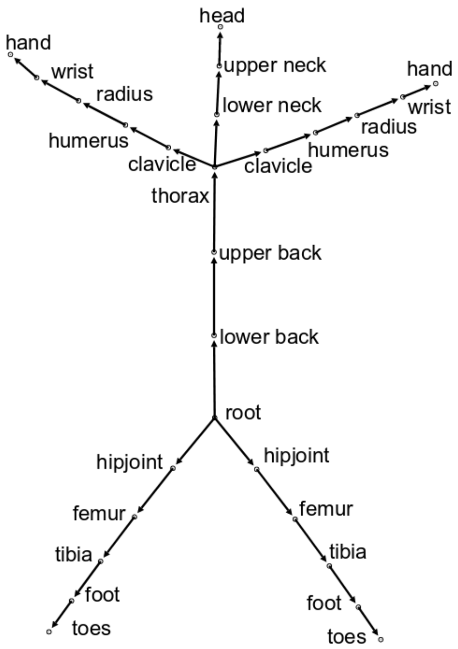

Our main application is to *computer motion capture*.
For each motion, we get a set of curves in SO(3), representing the rotation of joints relative to a fixed origin (root).
Given two motions, we want to compute some kind of distance between them.
We use data from the Carnegie Melon University MoCap Database `http://mocap.cs.cmu.edu`.

Shape Analysis: Technical Setup

Shapes are viewed as *unparametrized* curves taking values, in our case, on a finite-dimensional Lie group G whose Lie algebra is denoted by $\mathfrak{g}$.

We identify curves modulo reparametrization.

For technical reasons, we restrict to the space of *immersions*

$$\mathrm{Imm} := \{c\colon [0,1] \to G \mid c' \neq 0\}$$

The group D^+ of orientation-preserving diffeomorphisms of $[0,1]$ acts on Imm by composition

$$c.\varphi := c \circ \varphi.$$

We denote $\mathcal{S} := \mathrm{Imm}/D^+$.

Shape Analysis: Technical Setup

Similarity between *shapes* is then measured by

$$d_{\mathcal{S}}([c],[c']) := \inf_{\varphi} d_{\mathcal{P}}(c, c'.\varphi).$$

The (pseudo-)distance $d_{\mathcal{P}}$ on *parametrized* curves must be reparametrization invariant.

The standard choice is obtained through a Riemannian metric on Imm. In the end one gets

$$d_{\mathcal{P}}(c,c') = \sqrt{\int_0^1 \|q(t) - q'(t)\|^2 \,\mathrm{d}t}$$

where

$$q(t) := \frac{(R_{c(t)}^{-1})_*(\dot{c}(t))}{\sqrt{|\dot{c}(t)|}}$$

is called the *Square root velocity transform* (SRVT) of the curve c.

Shape Analysis: Technical Setup

Some observations about $d_{\mathcal{P}}$:

1. it is only a pseudometric.
2. it corresponds to the geodesic distance of a weak Riemannian metric on Imm, obtained as the pullback of the usual L^2 metric on curves in $\mathfrak{g}$ under the SRVT.
3. it is reparametrization invariant.

Hence, the similarity measure for shapes is

$$d_{\mathcal{S}}([c],[c']) = \inf_{\varphi \in D^+} \left(\int_0^1 \left\| q - (q'.\varphi)\sqrt{\dot{\varphi}} \right\|^2 \right)^{1/2}.$$

This optimization problem is often solved using *dynamic programming*.

Signatures

Signatures, introduced by K.T. Chen (1957) and generalized by Lyons (1998) provide a "universal" representation of paths in $\mathbb{R}^d$.

Definition

Let $x : [0,1] \longrightarrow \mathbb{R}^d$, given a multi-index $\mathcal{I} = (i_1, \ldots, i_n) \in \{1, \ldots, d\}^n$ define

$$S(x)^{\mathcal{I}} := \int\cdots\int_{0<s_1<\cdots<s_n<1} \dot{x}^{i_1}(s_1)\cdots\dot{x}^{i_n}(s_n)\,\mathrm{d}s_1\cdots\mathrm{d}s_n.$$

*The **full** collection $S(x) := (S(x)^{\mathcal{I}} : \mathcal{I})$ is called the signature of x.*

Some examples:

$$S(x)^i = \int_0^1 \dot{x}_s^i\,\mathrm{d}s = x^i(1) - x^i(0), \qquad S(x)^{ij} = \int_0^1 (x^i(s) - x^i(0))\dot{x}^j(s)\,\mathrm{d}s.$$

Signatures

Some of its properties are:

1. It is *reparametrization invariant*, i.e. if $\varphi \in D^+$ then $S(x \circ \varphi) = S(x)$.
2. Characterizes the path up to the removal of excursions.
3. It is invertible under certain assumptions.

Some observations:

1. This representation of the path is infinite dimensional, but it can be made finite-dimensional at the cost of losing some information about the path.
2. There is some redundancy. However *log-signature* provides a fully compressed representation of x.
3. The first few levels have a clear geometrical interpretation.

Signatures

In practice, we are only given *discrete* data, i.e. a sequence of samples $(X_1, \ldots, X_N)$.

We form a path x by linearly interpolating these values.[a]

In this case, the iterated integrals are actually explicit, so no need for numerical integration. In fact, if $x(t) = a + bt$ then

$$S(x) = \exp_{\otimes}(b)$$

in the *tensor algebra* $T(\mathbb{R}^d)$.

Since there is a fixed number of multi-indices having lenght less than some value, the size of the representation does not depend on the number of samples N. The same holds for the log-signature.

[a] However, there exists a fully discrete approach recently introduced by Diehl, Ebrahimi-Fard and T.

Signatures on Lie groups

Let G be a d-dimensional Lie group with Lie algebra $\mathfrak{g}$.

Definition

The Maurer–Cartan form of G is the $\mathfrak{g}$-valued 1-form

$$\omega_g(v) := (R_g^{-1})_*(v), \quad g \in G, v \in T_g G.$$

In particular, if $X_1, \ldots, X_d$ is a basis for $\mathfrak{g}$ then $\omega_g(v) = \omega_g^1(v) X_1 + \cdots + \omega_g^d(v) X_d$.

Definition (Chen (1954))

Consider a curve $\alpha : [0,1] \to G$. The signature $S(\alpha)$ is defined as the signature of the $\mathbb{R}^d$-valued path x such that

$$\dot{x}(t) = (\omega^1_{\alpha(t)}(\dot{\alpha}(t)), \ldots, \omega^d_{\alpha(t)}(\dot{\alpha}(t))).$$

It is known that $S(x) = Y(1)$ where Y solves

$$\dot{Y} = Y \otimes \dot{x}$$

in $T(\mathbb{R}^d)$.

This allows for an interpretation of $S(x)$ as a (infinite dimensional) Lie group exponential.

$$\begin{tikzcd}
\alpha \in C([0,1],G) \arrow[r, "\omega=\delta^r"] \arrow[d, "Y"] & C([0,1],\mathfrak{g}) \arrow[d] \\
C([0,1],\mathcal{G}) & C([0,1],\mathcal{L}(\mathbb{R}^d)) \arrow[l, "\mathrm{Evol}"']
\end{tikzcd}
\qquad
\begin{tikzcd}
\alpha \in C([0,1],G) \arrow[r, "\omega=\delta^r"] \arrow[d, "S"] & C([0,1],\mathfrak{g}) \arrow[d] \\
\mathcal{G} & C([0,1],\mathcal{L}(\mathbb{R}^d)) \arrow[l, "\mathrm{evol}"']
\end{tikzcd}$$

Signatures on Lie groups: The case of SO(3)

In our setting, we have just a sample of the curve, that is, a sequence of rotation matrices $A_1, \ldots, A_N \in \mathrm{SO}(3)$.

We do "geodesic interpolation", which corresponds to linear interpolation in this context.

Given $A, B \in \mathrm{SO}(3)$, define $\alpha \colon [0,1] \longrightarrow \mathrm{SO}(3)$ by

$$\alpha(t) := \exp(t \log(BA^\top))A,$$

so that $\alpha(0) = A$ and $\alpha(1) = B$.

For this choice,

$$\omega_{\alpha(t)}(\dot{\alpha}(t)) = \log(BA^\top).$$

Therefore, we have also an easy expression for the signature of α.

Clustering: Comparing signatures

We need a way of comparing signatures. There are several choices:

1. using the metric inherited from $T(\mathbb{R}^d)$, i.e.
$$d(g,h) := \|h - g\|_{T(\mathbb{R}^d)},$$
2. using an ad-hoc metric for signatures:
$$\rho^n(S(x), S(y)) := \|S(x) \otimes S(\tilde{y})\|,$$
3. compare log-signatures.

We also compare with currently used methods based on the SRVT, i.e. dynamic programming.

Clustering

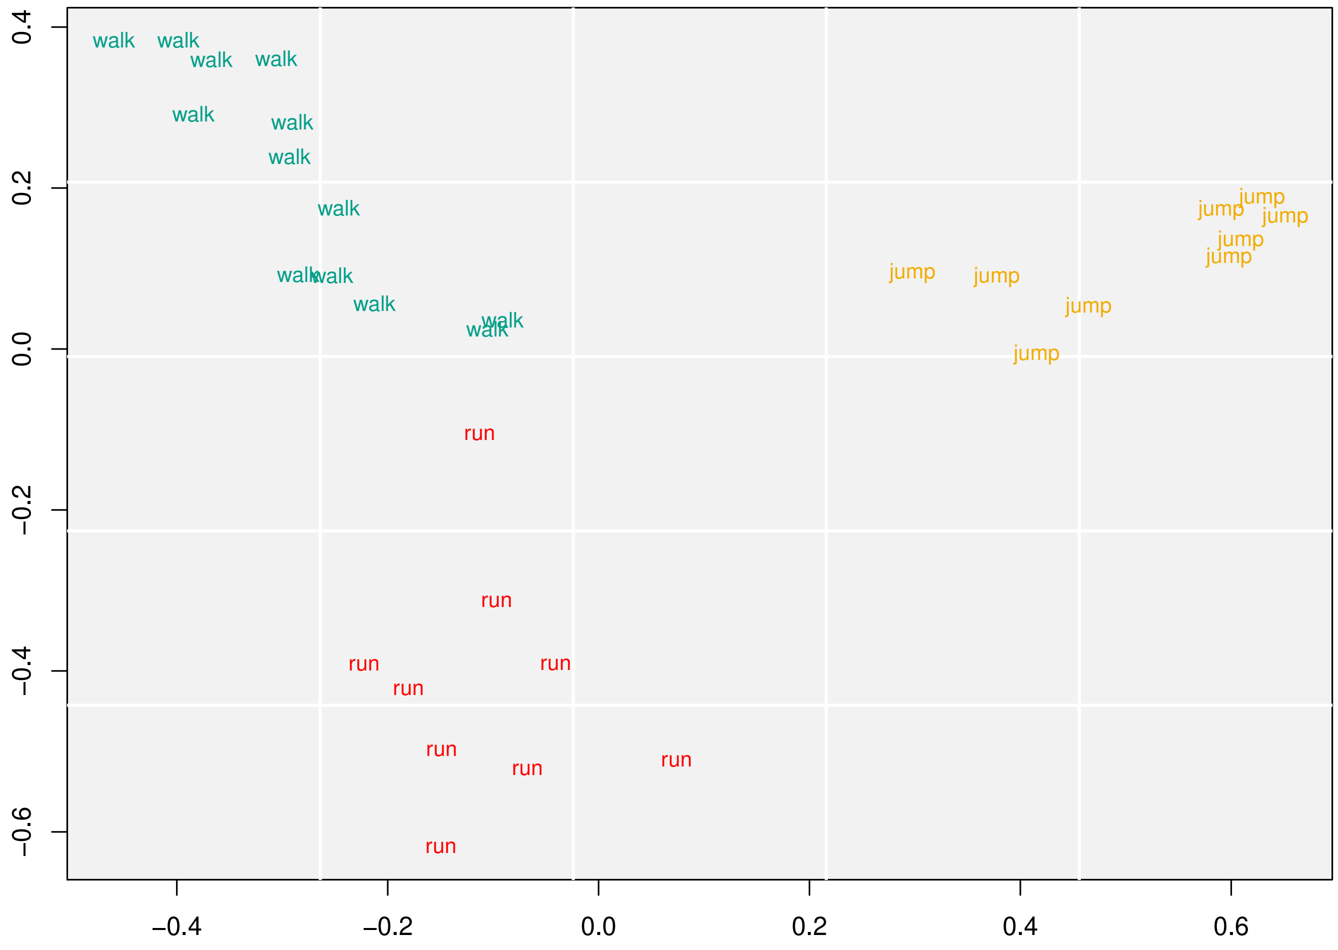

Clustering

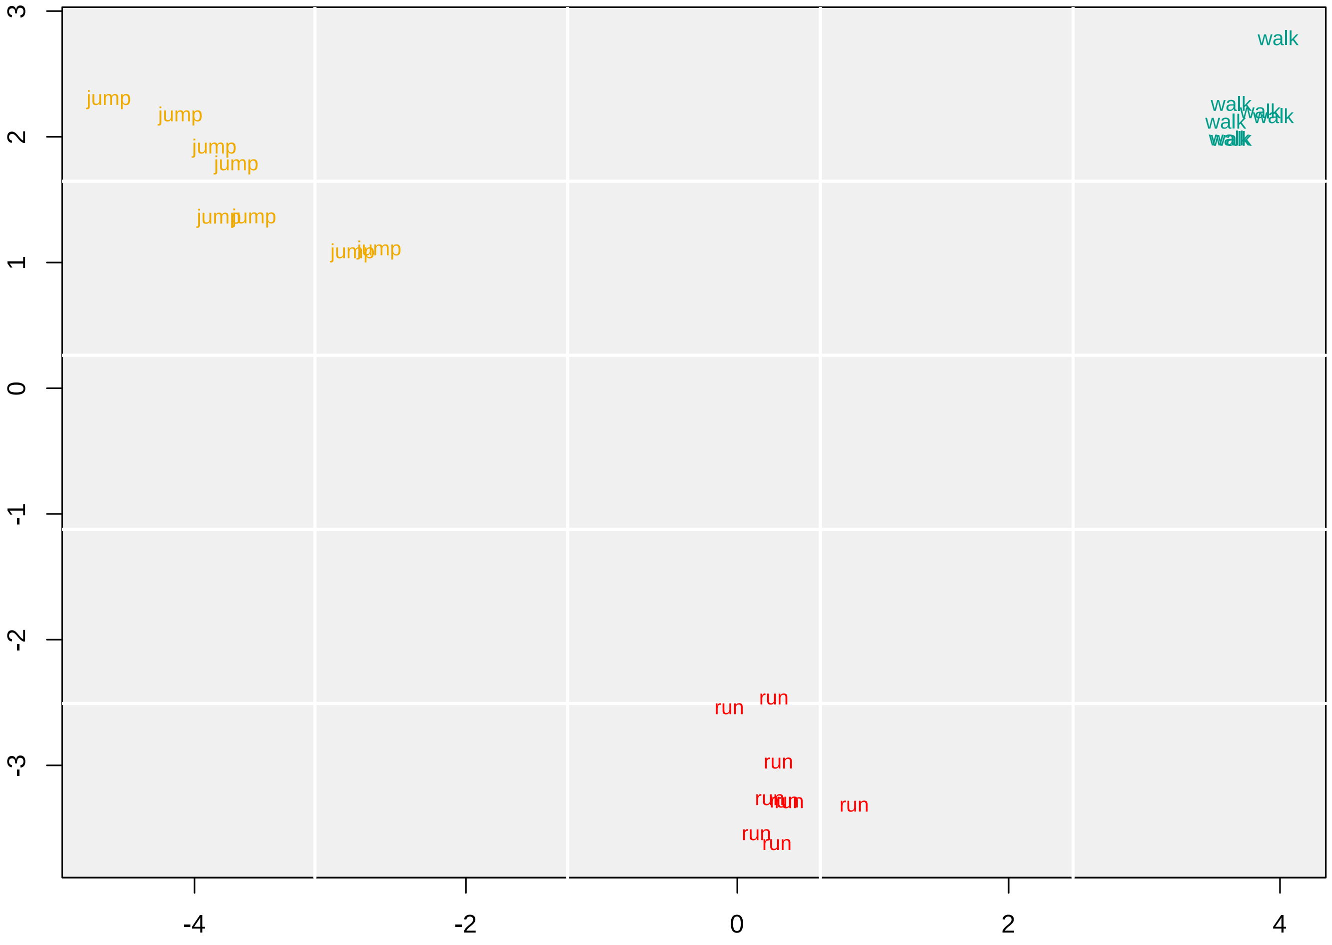

walk
walk
walk
walk
walk
walk
jump
jump
jump
jump
jump
jump
jump
jump
run
run
run
run
run
run
run
run
3
2
1
0
-1
-2
-3
-4
-2
0
2
4

Clustering: a concrete example

Now, let c_a and c_b correspond to "walking" and "jogging" animations.

We generate a geodesic interpolation $\bar{c}$ between the curves, i.e. $c(0,\cdot) = c_a$, $c(1,\cdot) = c_b$ and for $s \in (0,1)$ the animation $c(s,\cdot)$ is a mixture of both.

In practice, this is generated using the SRVT so in fact we are doing linear interpolation at the level of the Lie algebra.

Signatures were computed using the `iisignature` Python package by J. Reizenstein and B. Graham.

We can then look at the behaviour of the different similarity measures when s varies.

Remark

Since the distance $d_{\mathcal{S}}$ coincides with the geodesic distance, we will see a straight line for this metric.

Clustering: a concrete example (cont.)

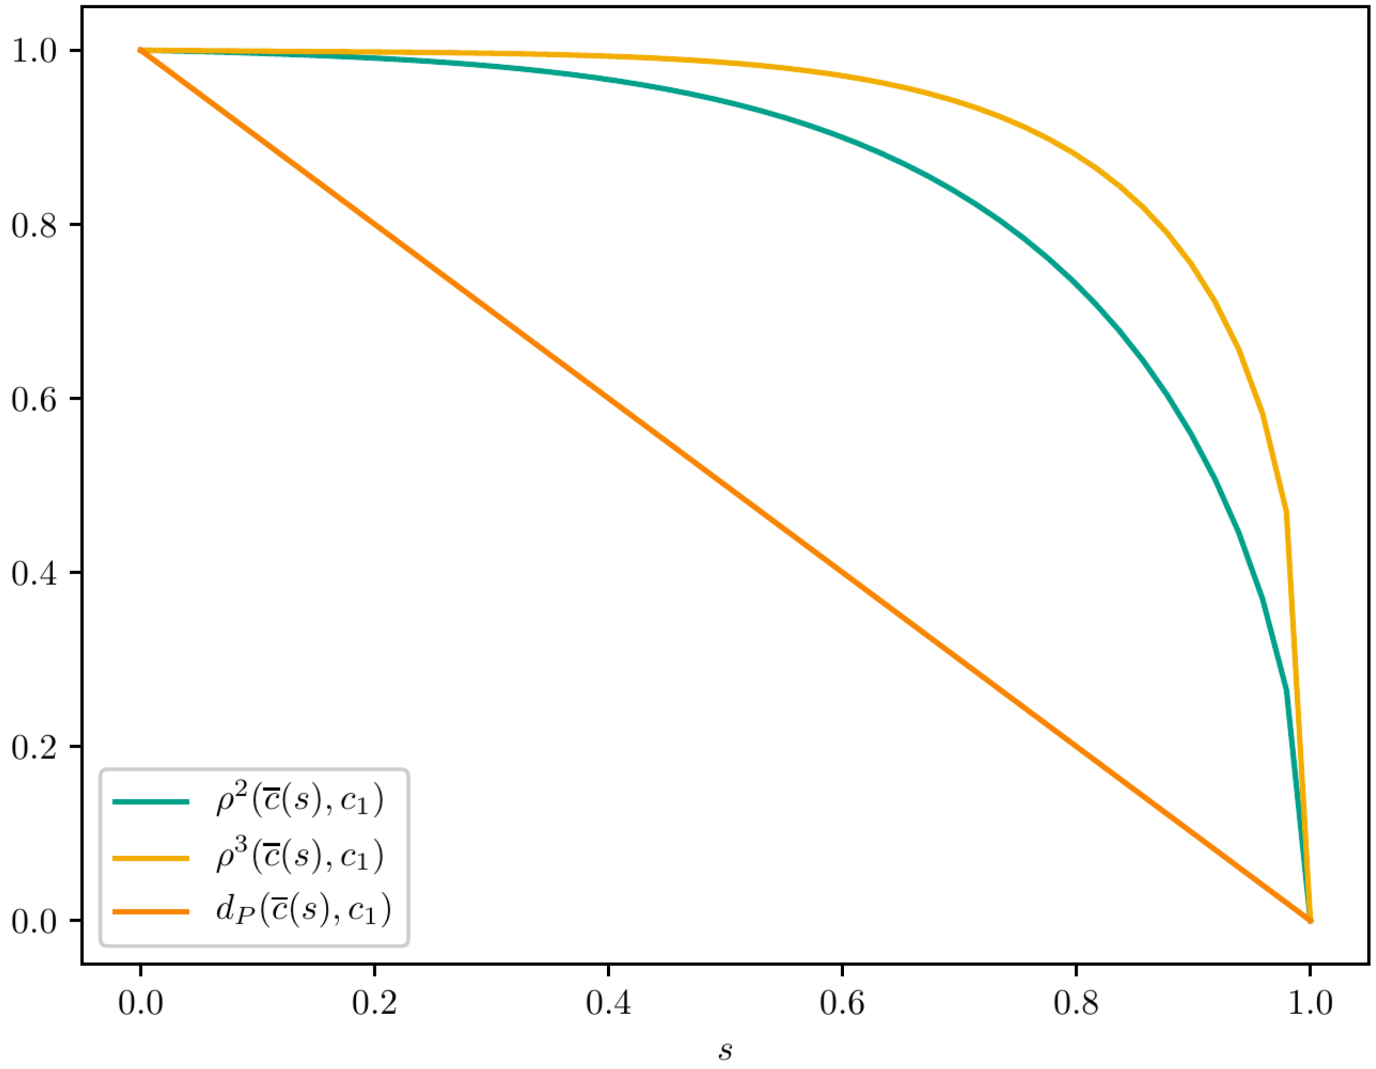

Clustering: a concrete example (cont.)

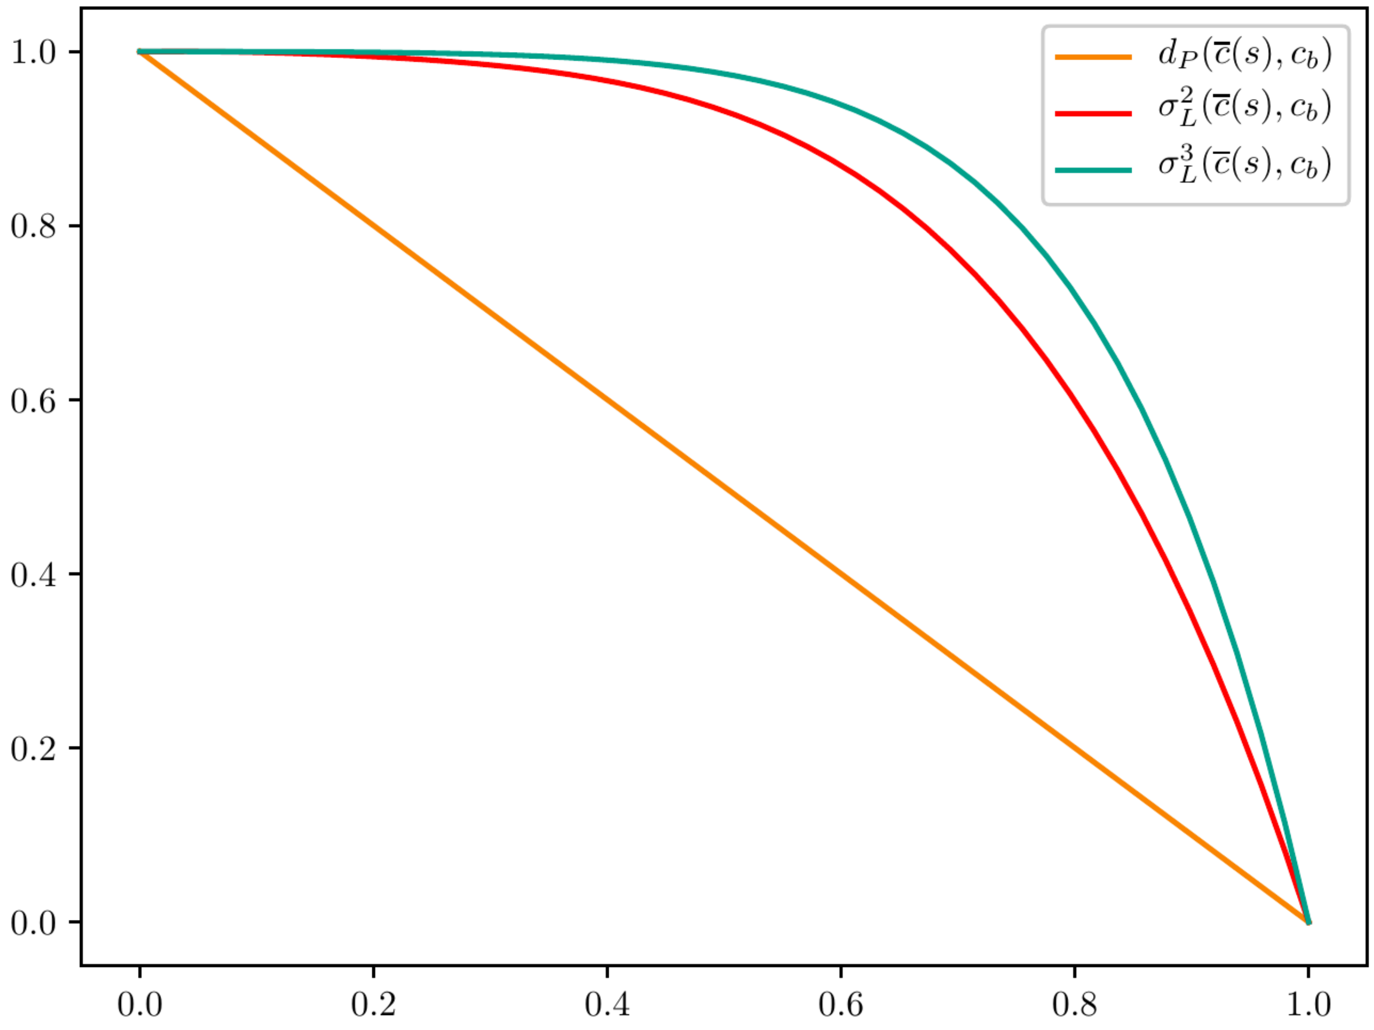

Clustering: a concrete example (cont.)

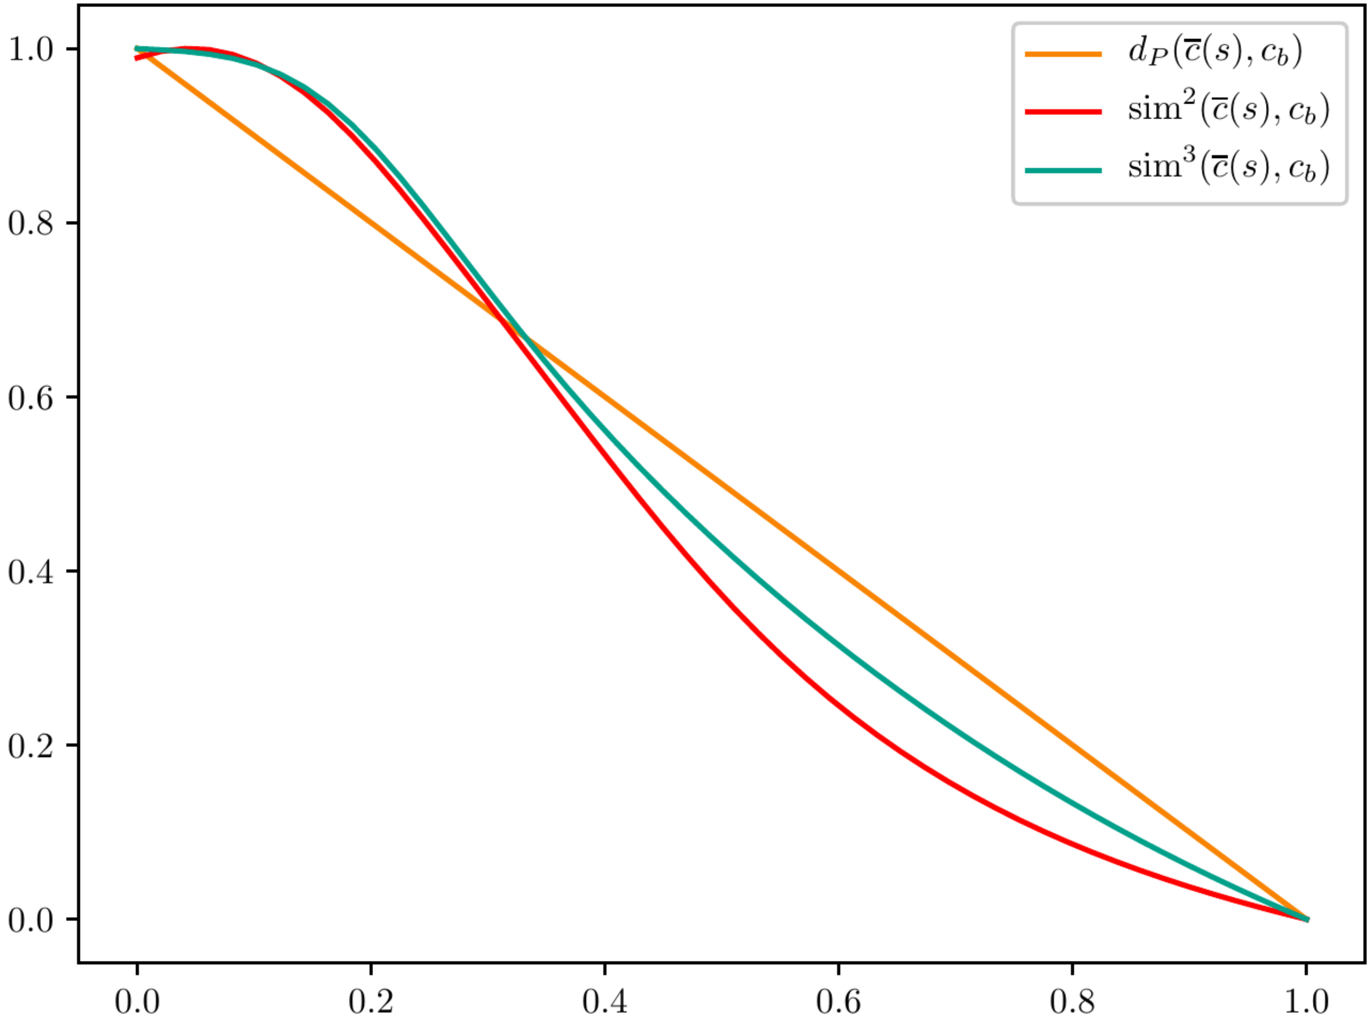

Perspectives

Questions:

1. Pullback metric from signatures to curves.
2. Better understanding of the various metrics. Dependence on the underlying norm. Weighted norm with learned wieghts.
3. Clarification of the geometrical interpretation of the signature.

Thanks!