

Global divergences between measures: from Hausdorff distance to Optimal Transport

Jean Feydy Alain Trouvé

ShapeMI workshop, Miccai, Granada – 20th September, 2018


Écoles Normales Supérieures de Paris et Paris-Saclay

Collaboration with B. Charlier, J. Glaunès (KeOps library);

S.-i. Amari, G. Peyré, T. Séjourné, F.-X. Vialard (OT theory)


Today: focus on shape registration

Source **A**, target **B**,




Today: focus on shape registration

Source A , target B , mapping φ




Today: focus on shape registration

Source A , target B , mapping φ



Today: focus on shape registration


Source A , target B , mapping φ



Today: focus on shape registration

Source A , target B , mapping φ

$$A \xrightarrow[\text{Model}]{\varphi} \varphi(A) = A' \xleftrightarrow[\text{Loss}]{\quad} B$$



A good Loss function is a guarantee of robustness

Iterative Matching Algorithm

- 1: $A' \leftarrow A$
 - 2: **repeat**
 - 3: $L, v \leftarrow \text{Loss}(A', B), -\partial_{A'} \text{Loss}(A', B)$
 - 4: $A' \leftarrow A' + \text{Model}(v)$
 - 5: **until** $L < \text{tol}$
 Output: deformed shape $A' = \varphi(A)$.
-

A good Loss function is a guarantee of robustness

Iterative Matching Algorithm

- 1: $A' \leftarrow A$
 - 2: **repeat**
 - 3: $L, v \leftarrow \text{Loss}(A', B), -\partial_{A'} \text{Loss}(A', B)$
 - 4: $A' \leftarrow A' + \text{Model}(v)$
 - 5: **until** $L < \text{tol}$
- Output:** deformed shape $A' = \varphi(A)$.
-

“Model” encodes the **prior knowledge** on admissible deformations:

- *smoothing* convolution
- LDDMM/SVF *backprop* + regularization + *shooting*
- *trained* neural network

A good Loss function is a guarantee of robustness

Iterative Matching Algorithm


- 1: $A' \leftarrow A$
 - 2: **repeat**
 - 3: $L, \nu \leftarrow \text{Loss}(A', B), -\partial_{A'} \text{Loss}(A', B)$
 - 4: $A' \leftarrow A' + \text{Model}(\nu)$
 - 5: **until** $L < \text{tol}$
- Output:** deformed shape $A' = \varphi(A)$.
-

“Model” encodes the **prior knowledge** on admissible deformations:

- *smoothing* convolution
- LDDMM/SVF *backprop* + regularization + *shooting*
- *trained* neural network


\Rightarrow The *raw* Loss gradient ν is what **drives** the registration

On labeled shapes, use a spring energy



Anatomical landmarks from *A morphometric approach for the analysis of body shape in bluefin tuna*, Addis et al., 2009.

On labeled shapes, use a spring energy



Anatomical landmarks from *A morphometric approach for the analysis of body shape in bluefin tuna*, Addis et al., 2009.

Encoding unlabeled shapes as measures

Let's enforce sampling invariance:


$$A \longrightarrow \alpha = \sum_{i=1}^N \alpha_i \delta_{x_i}, \quad B \longrightarrow \beta = \sum_{j=1}^M \beta_j \delta_{y_j}.$$

Encoding unlabeled shapes as measures

Let's enforce sampling invariance:

$$A \longrightarrow \alpha = \sum_{i=1}^N \alpha_i \delta_{x_i},$$


$$B \longrightarrow \beta = \sum_{j=1}^M \beta_j \delta_{y_j}.$$




Encoding unlabeled shapes as measures

Let's enforce sampling invariance:


$$A \longrightarrow \alpha = \sum_{i=1}^N \alpha_i \delta_{x_i}, \quad B \longrightarrow \beta = \sum_{j=1}^M \beta_j \delta_{y_j}.$$



A baseline setting: density registration




A baseline setting: density registration



$$\alpha = \sum_{i=1}^N \alpha_i \delta_{x_i}, \quad \beta = \sum_{j=1}^M \beta_j \delta_{y_j}.$$


A baseline setting: density registration



$$\alpha = \sum_{i=1}^N \alpha_i \delta_{x_i}, \quad \beta = \sum_{j=1}^M \beta_j \delta_{y_j}.$$

$$\sum_{i=1}^N \alpha_i = 1 = \sum_{j=1}^M \beta_j$$

A baseline setting: density registration




$$\alpha = \sum_{i=1}^N \alpha_i \delta_{x_i}, \quad \beta = \sum_{j=1}^M \beta_j \delta_{y_j}.$$

$$\sum_{i=1}^N \alpha_i = 1 = \sum_{j=1}^M \beta_j$$

Display $v = -\nabla_{x_i} d(\alpha, \beta)$.

A baseline setting: density registration



$$\alpha = \sum_{i=1}^N \alpha_i \delta_{x_i}, \quad \beta = \sum_{j=1}^M \beta_j \delta_{y_j}.$$

$$\sum_{i=1}^N \alpha_i = 1 = \sum_{j=1}^M \beta_j$$

Display $v = -\nabla_{x_i} d(\alpha, \beta)$.

Seamless extensions to:

- $\sum_i \alpha_i \neq \sum_j \beta_j$ [Chizat et al., 2018],
- curves and surfaces [Kaltenmark et al., 2017],
- variable weights α_i .


Computing fidelities between **measures**:

1. **Computer graphics**: Hausdorff distance
2. **Statistics**: kernel distances
3. **Optimal Transport**: Wasserstein distance


Computing fidelities between **measures**:

1. **Computer graphics**: Hausdorff distance
2. **Statistics**: kernel distances
3. **Optimal Transport**: Wasserstein distance
4. Efficient GPU routines: **KeOps**

The weighted Hausdorff distance




The weighted Hausdorff distance



p -Hausdorff distance:

$$\text{Loss}(\alpha, \beta) = \frac{1}{2} \sum_i \alpha_i \cdot \min_j \|x_i - y_j\|^p$$


The weighted Hausdorff distance



p -Hausdorff distance:

$$\text{Loss}(\alpha, \beta) = \frac{1}{2} \sum_i \alpha_i \cdot \min_j \|x_i - y_j\|^p + \frac{1}{2} \sum_j \beta_j \cdot \min_i \|x_i - y_j\|^p$$

The weighted Hausdorff distance




p -Hausdorff distance:

$$\begin{aligned} \text{Loss}(\alpha, \beta) &= \frac{1}{2} \sum_i \alpha_i \cdot \min_j \|x_i - y_j\|^p + \frac{1}{2} \sum_j \beta_j \cdot \min_i \|x_i - y_j\|^p \\ &= \frac{1}{2} \langle \alpha, b \rangle + \frac{1}{2} \langle \beta, a \rangle \end{aligned}$$

with $a(x) = d(x, \text{supp}(\alpha))^p$

$b(x) = d(x, \text{supp}(\beta))^p$

The weighted Hausdorff distance




p -Hausdorff distance:

$$\begin{aligned} \text{Loss}(\alpha, \beta) &= \frac{1}{2} \sum_i \alpha_i \cdot \min_j \|x_i - y_j\|^p + \frac{1}{2} \sum_j \beta_j \cdot \min_i \|x_i - y_j\|^p \\ &= \frac{1}{2} \langle \alpha, b \rangle + \frac{1}{2} \langle \beta, a \rangle \end{aligned}$$

with $a(x) = d(x, \text{supp}(\alpha))^p$

$b(x) = d(x, \text{supp}(\beta))^p$

The weighted Hausdorff distance




p -Hausdorff distance:

$$\begin{aligned}\text{Loss}(\alpha, \beta) &= \frac{1}{2} \sum_i \alpha_i \cdot \min_j \|x_i - y_j\|^p + \frac{1}{2} \sum_j \beta_j \cdot \min_i \|x_i - y_j\|^p \\ &= \frac{1}{2} \langle \alpha, b \rangle + \frac{1}{2} \langle \beta, a \rangle \\ &= \frac{1}{2} \langle \alpha, b - a \rangle + \frac{1}{2} \langle \beta, a - b \rangle\end{aligned}$$

$$\begin{aligned}\text{with } a(x) &= d(x, \text{supp}(\alpha))^p \\ b(x) &= d(x, \text{supp}(\beta))^p\end{aligned}$$

The weighted Hausdorff distance



p -Hausdorff distance:


$$\begin{aligned}\text{Loss}(\alpha, \beta) &= \frac{1}{2} \sum_i \alpha_i \cdot \min_j \|x_i - y_j\|^p + \frac{1}{2} \sum_j \beta_j \cdot \min_i \|x_i - y_j\|^p \\ &= \frac{1}{2} \langle \alpha, b \rangle + \frac{1}{2} \langle \beta, a \rangle \\ &= \frac{1}{2} \langle \alpha, b - a \rangle + \frac{1}{2} \langle \beta, a - b \rangle \\ &= \frac{1}{2} \langle \alpha - \beta, b - a \rangle\end{aligned}$$

$$\text{with } a(x) = d(x, \text{supp}(\alpha))^p$$

$$b(x) = d(x, \text{supp}(\beta))^p$$


Naive projections in Hausdorff cause imbalance

$$\text{Loss}(\alpha, \beta) = \frac{1}{2} \langle \alpha, b - a \rangle + \frac{1}{2} \langle \beta, a - b \rangle$$




Naive projections in Hausdorff cause imbalance

$$\text{Loss}(\alpha, \beta) = \frac{1}{2} \langle \alpha, b - a \rangle + \frac{1}{2} \langle \beta, a - b \rangle$$




Naive projections in Hausdorff cause imbalance

$$\text{Loss}(\alpha, \beta) = \frac{1}{2} \langle \alpha, b - a \rangle + \frac{1}{2} \langle \beta, a - b \rangle$$



Naive projections in Hausdorff cause imbalance

$$\text{Loss}(\alpha, \beta) = \frac{1}{2} \langle \alpha, b - a \rangle + \frac{1}{2} \langle \beta, a - b \rangle$$



Kernel distances: distance fields computed through convolutions

Kernel distances, aka. **blurred SSDs**:

$$\text{choose } \mathbf{a}(x) = -(k \star \alpha)(x) = -\sum_i \alpha_i k(x, x_i)$$

$$\text{and use } \frac{1}{2} \langle \alpha - \beta, \mathbf{b} - \mathbf{a} \rangle = \frac{1}{2} \langle \alpha - \beta, k \star (\alpha - \beta) \rangle.$$

Kernel distances: distance fields computed through convolutions

Kernel distances, aka. **blurred SSDs**:

$$\text{choose } a(x) = -(k \star \alpha)(x) = -\sum_i \alpha_i k(x, x_i)$$

$$\text{and use } \frac{1}{2} \langle \alpha - \beta, b - a \rangle = \frac{1}{2} \langle \alpha - \beta, k \star (\alpha - \beta) \rangle.$$

The **Energy Distance**: an underrated kernel, $k(x, y) = -\|x - y\|$.

$$a(x) = \sum_i \alpha_i \|x - x_i\| \quad \text{instead of} \quad a(x) = \min_i \|x - x_i\|$$

$$b(x) = \sum_j \beta_j \|x - y_j\| \quad \text{instead of} \quad b(x) = \min_j \|x - y_j\|.$$

Kernel distances: distance fields computed through convolutions

Kernel distances, aka. **blurred SSDs**:

$$\text{choose } \mathbf{a}(x) = -(k \star \alpha)(x) = -\sum_i \alpha_i k(x, \mathbf{x}_i)$$

$$\text{and use } \frac{1}{2} \langle \alpha - \beta, \mathbf{b} - \mathbf{a} \rangle = \frac{1}{2} \langle \alpha - \beta, k \star (\alpha - \beta) \rangle.$$


The **Energy Distance**: an underrated kernel, $k(x, y) = -\|x - y\|$.

$$\mathbf{a}(x) = \sum_i \alpha_i \|x - \mathbf{x}_i\| \quad \text{instead of} \quad \mathbf{a}(x) = \min_i \|x - \mathbf{x}_i\|$$


$$\mathbf{b}(x) = \sum_j \beta_j \|x - \mathbf{y}_j\| \quad \text{instead of} \quad \mathbf{b}(x) = \min_j \|x - \mathbf{y}_j\|.$$

$$\begin{aligned} \text{Loss}(\alpha, \beta) &= \sum_i \sum_j \alpha_i \beta_j \|\mathbf{x}_i - \mathbf{y}_j\| \\ &\quad - \frac{1}{2} \sum_i \sum_j \alpha_i \alpha_j \|\mathbf{x}_i - \mathbf{x}_j\| - \frac{1}{2} \sum_i \sum_j \beta_i \beta_j \|\mathbf{y}_i - \mathbf{y}_j\| \end{aligned}$$

The Hausdorff distance is local, the Energy Distance is global



Hausdorff, min



Kernel, Σ

An idea from Optimal Transport theory: Sinkhorn divergences

Optimal Transport = Hausdorff + mass spreading constraint

Computational OT [Cuturi, 2013, Peyré and Cuturi, 2018]:

Start from an ε -smoothed **Hausdorff** distance, but let the influence fields a and b **interact** with each other.

Enforce a **mass spreading** constraint on the spring system:

all of α should be linked to all of β .

Optimal Transport = Hausdorff + mass spreading constraint

Computational OT [Cuturi, 2013, Peyré and Cuturi, 2018]:

Start from an ε -smoothed **Hausdorff** distance, but let the influence fields a and b **interact** with each other.

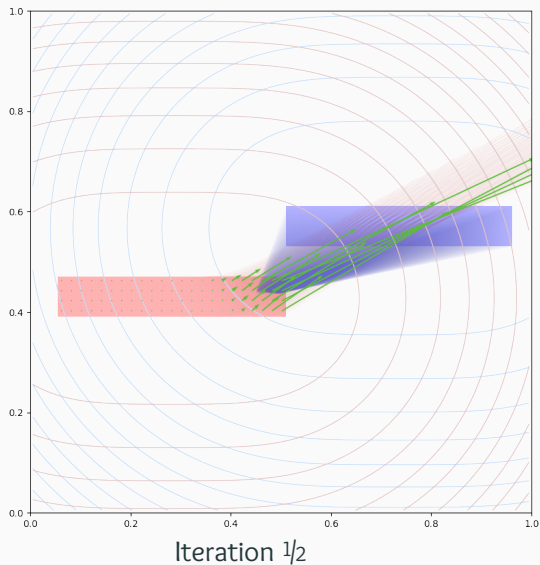
Enforce a **mass spreading** constraint on the spring system:
all of α should be linked to all of β .

In practice: use the 5-line **Sinkhorn** algorithm.


Updates a and b alternatively.

Converges in about 10-20 steps – x2 convolutions.


The Sinkhorn algorithm, in practice; with $\|x - y\|^2$ and $\sqrt{\varepsilon} = .1$




The Sinkhorn algorithm, in practice; with $\|x - y\|^2$ and $\sqrt{\varepsilon} = .1$




The Sinkhorn algorithm, in practice; with $\|x - y\|^2$ and $\sqrt{\varepsilon} = .1$




The Sinkhorn algorithm, in practice; with $\|x - y\|^2$ and $\sqrt{\varepsilon} = .1$




The Sinkhorn algorithm, in practice; with $\|x - y\|^2$ and $\sqrt{\varepsilon} = .1$




The Sinkhorn algorithm, in practice; with $\|x - y\|^2$ and $\sqrt{\varepsilon} = .1$




The Sinkhorn algorithm, in practice; with $\|x - y\|^2$ and $\sqrt{\varepsilon} = .1$




The Sinkhorn algorithm, in practice; with $\|x - y\|^2$ and $\sqrt{\varepsilon} = .1$




The Sinkhorn algorithm, in practice; with $\|x - y\|^2$ and $\sqrt{\varepsilon} = .1$




The Sinkhorn algorithm, in practice; with $\|x - y\|^2$ and $\sqrt{\varepsilon} = .1$



The Sinkhorn algorithm, in practice; with $\|x - y\|^2$ and $\sqrt{\varepsilon} = .1$




The Sinkhorn algorithm, in practice; with $\|x - y\|^2$ and $\sqrt{\varepsilon} = .1$




Iteration 10

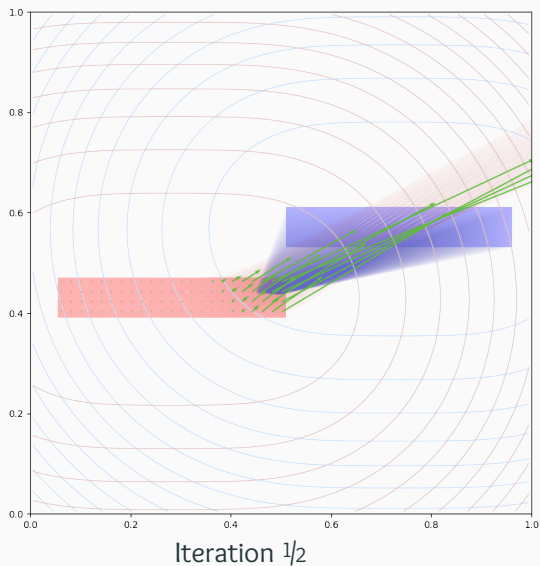
The Sinkhorn algorithm, in practice; with $\|x - y\|^2$ and $\sqrt{\varepsilon} = .1$



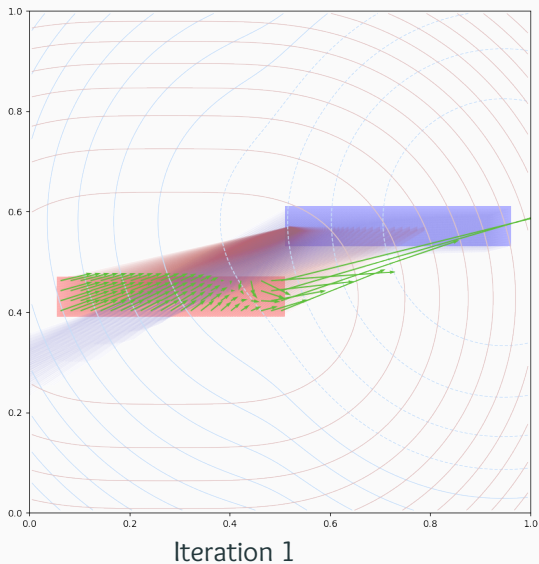
The Sinkhorn algorithm, in practice; with $\|x - y\|^2$ and $\sqrt{\varepsilon} = .1$




The Sinkhorn algorithm, in practice; with $\|x - y\|^2$ and $\sqrt{\varepsilon} = .1$




The Sinkhorn algorithm, in practice; with $\|x - y\|^2$ and $\sqrt{\varepsilon} = .1$



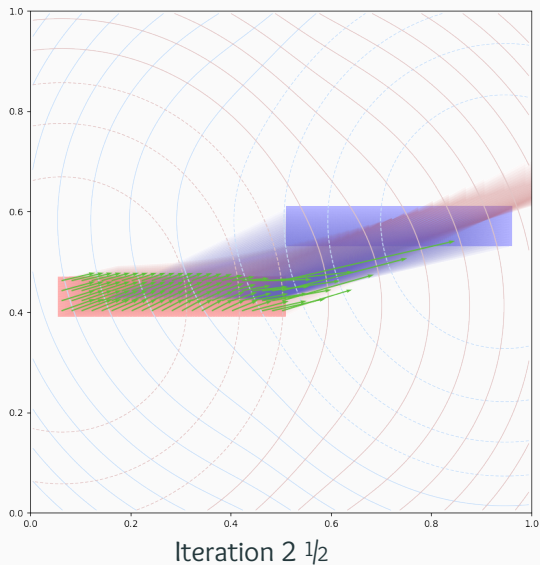
The Sinkhorn algorithm, in practice; with $\|x - y\|^2$ and $\sqrt{\varepsilon} = .1$




The Sinkhorn algorithm, in practice; with $\|x - y\|^2$ and $\sqrt{\varepsilon} = .1$



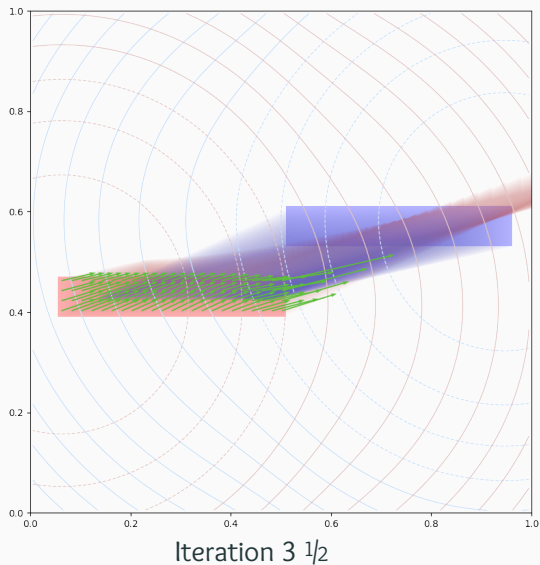
The Sinkhorn algorithm, in practice; with $\|x - y\|^2$ and $\sqrt{\varepsilon} = .1$




The Sinkhorn algorithm, in practice; with $\|x - y\|^2$ and $\sqrt{\varepsilon} = .1$




The Sinkhorn algorithm, in practice; with $\|x - y\|^2$ and $\sqrt{\varepsilon} = .1$



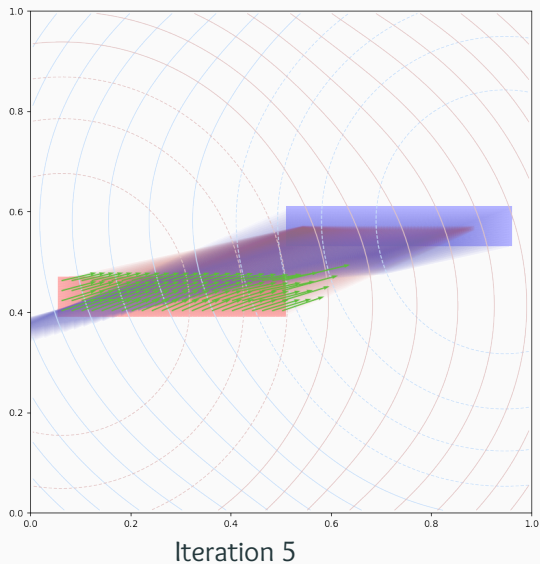
The Sinkhorn algorithm, in practice; with $\|x - y\|^2$ and $\sqrt{\varepsilon} = .1$




The Sinkhorn algorithm, in practice; with $\|x - y\|^2$ and $\sqrt{\varepsilon} = .1$




The Sinkhorn algorithm, in practice; with $\|x - y\|^2$ and $\sqrt{\varepsilon} = .1$




The Sinkhorn algorithm, in practice; with $\|x - y\|^2$ and $\sqrt{\varepsilon} = .1$




The Sinkhorn algorithm, in practice; with $\|x - y\|^2$ and $\sqrt{\varepsilon} = .1$




The Sinkhorn algorithm, in practice; with $\|x - y\|^2$ and $\sqrt{\varepsilon} = .1$




The Sinkhorn algorithm, in practice; with $\|x - y\|^2$ and $\sqrt{\varepsilon} = .1$




The Sinkhorn algorithm, in practice; with $\|x - y\|^2$ and $\sqrt{\varepsilon} = .1$




The Sinkhorn algorithm, in practice; with $\|x - y\|^2$ and $\sqrt{\varepsilon} = .1$




The Sinkhorn algorithm, in practice; with $\|x - y\|^2$ and $\sqrt{\varepsilon} = .1$




The Sinkhorn algorithm, in practice; with $\|x - y\|^2$ and $\sqrt{\varepsilon} = .1$




The Sinkhorn algorithm, in practice; with $\|x - y\|^2$ and $\sqrt{\varepsilon} = .1$




The Sinkhorn algorithm, in practice; with $\|x - y\|^2$ and $\sqrt{\varepsilon} = .1$




The Sinkhorn algorithm, in practice; with $\|x - y\|^2$ and $\sqrt{\varepsilon} = .1$




The Sinkhorn algorithm, in practice; with $\|x - y\|^2$ and $\sqrt{\varepsilon} = .1$




The Sinkhorn algorithm, in practice; with $\|x - y\|^2$ and $\sqrt{\varepsilon} = .1$




The Sinkhorn algorithm, in practice; with $\|x - y\|^2$ and $\sqrt{\varepsilon} = .1$




The Sinkhorn algorithm, in practice; with $\|x - y\|^2$ and $\sqrt{\varepsilon} = .1$




The Sinkhorn algorithm, in practice; with $\|x - y\|^2$ and $\sqrt{\varepsilon} = .1$



The Sinkhorn algorithm, in practice; with $\|x - y\|^2$ and $\sqrt{\varepsilon} = .1$



The Sinkhorn algorithm, in practice; with $\|x - y\|^2$ and $\sqrt{\varepsilon} = .1$



Theorem (F., Séjourné, Vialard, Trounev, Amari, Peyré; 2018)

We define a symmetric Sinkhorn divergence:

$$d_{\varepsilon\text{-Sinkhorn}}(\alpha, \beta).$$

Then, for all probability measures α, β and regularization $\varepsilon > 0$:

Theorem (F., Séjourné, Vialard, Trouvé, Amari, Peyré; 2018)

We define a symmetric Sinkhorn divergence:

$$d_{\varepsilon\text{-Sinkhorn}}(\alpha, \beta).$$

Then, for all probability measures α, β and regularization $\varepsilon > 0$:

$$0 \leq d_{\varepsilon\text{-Sinkhorn}}(\alpha, \beta) \quad \text{with equality iff. } \alpha = \beta$$

Theorem (F., Séjourné, Vialard, Trouvé, Amari, Peyré; 2018)

We define a symmetric Sinkhorn divergence:

$$d_{\varepsilon\text{-Sinkhorn}}(\alpha, \beta).$$

Then, for all probability measures α, β and regularization $\varepsilon > 0$:

$$0 \leq d_{\varepsilon\text{-Sinkhorn}}(\alpha, \beta) \quad \text{with equality iff. } \alpha = \beta$$

$$\alpha \mapsto d_{\varepsilon\text{-Sinkhorn}}(\alpha, \beta) \text{ is convex and differentiable}$$

Theorem (F., Séjourné, Vialard, Trouvé, Amari, Peyré; 2018)

We define a symmetric Sinkhorn divergence:

$$d_{\varepsilon\text{-Sinkhorn}}(\alpha, \beta).$$

Then, for all probability measures α, β and regularization $\varepsilon > 0$:

$$0 \leq d_{\varepsilon\text{-Sinkhorn}}(\alpha, \beta) \quad \text{with equality iff. } \alpha = \beta$$

$\alpha \mapsto d_{\varepsilon\text{-Sinkhorn}}(\alpha, \beta)$ is convex and differentiable

$$\text{Wasserstein}(\alpha, \beta) \xleftarrow{\varepsilon \rightarrow 0} d_{\varepsilon\text{-Sinkhorn}}(\alpha, \beta) \xrightarrow{\varepsilon \rightarrow +\infty} \text{Energy}(\alpha, \beta)$$

Theorem (F., Séjourné, Vialard, Trouvé, Amari, Peyré; 2018)

We define a symmetric Sinkhorn divergence:

$$d_{\varepsilon\text{-Sinkhorn}}(\alpha, \beta).$$

Then, for all probability measures α, β and regularization $\varepsilon > 0$:

$$0 \leq d_{\varepsilon\text{-Sinkhorn}}(\alpha, \beta) \quad \text{with equality iff. } \alpha = \beta$$


$\alpha \mapsto d_{\varepsilon\text{-Sinkhorn}}(\alpha, \beta)$ is convex and differentiable

$$\text{Wasserstein}(\alpha, \beta) \xleftarrow{\varepsilon \rightarrow 0} d_{\varepsilon\text{-Sinkhorn}}(\alpha, \beta) \xrightarrow{\varepsilon \rightarrow +\infty} \text{Energy}(\alpha, \beta)$$

These results can be generalized to arbitrary **feature** spaces


– e.g. (position, orientation, curvature).

The ε -Sinkhorn divergence; with $\|x - y\|^2$ and $\sqrt{\varepsilon} = .1$




A high-quality gradient.

The ε -Sinkhorn divergence; with $\|x - y\|^2$ and $\sqrt{\varepsilon} = .1$



A high-quality gradient.

The ε -Sinkhorn divergence; with $\|x - y\|^2$ and $\sqrt{\varepsilon} = .1$




A high-quality gradient.

(Data from the Spectral Log-Demons paper.)

In practice

Kernel OPERATIONS, with autodiff, without memory overflows


Kernel norm + gradient with N vertices on a cheap laptop's GPU (GTX960M)



Kernel OPERATIONS, with autodiff, without memory overflows

⇒ pip install pykeops ⇐
(Thanks Benjamin and Joan!)


Kernel norm + gradient with N vertices on a cheap laptop's GPU (GTX960M)



Kernel OPERATIONS, with autodiff, without memory overflows

⇒ pip install pykeops ⇐
(Thanks Benjamin and Joan!)


Fidelity + gradient with N vertices on a **cheap laptop's GPU** (GTX960M)



Kernel OPERATIONS, with autodiff, without memory overflows


⇒ pip install pykeops ⇐
(Thanks Benjamin and Joan!)

Fidelity + gradient with N vertices on a **high-end GPU (Tesla P100)**



We provide a reference PyTorch implementation

github.com/jeanfeydy/global-divergences.



Gradient of the Energy Distance, computed in 0.5s on my laptop.

Data from the OsteoArthritis Initiative:

52,319 and 34,966 voxels out of a 192-192-160 volume.

Robust, **geometry-aware** loss functions are easy to compute.

Robust, **geometry-aware** loss functions are easy to compute.

- Try using $k(x,y) = -\|x - y\|$!

Robust, **geometry-aware** loss functions are easy to compute.

- Try using $k(x,y) = -\|x - y\|$!
- Sinkhorn = Hausdorff + mass **spreading** constraint
 - \simeq best you can do without topology or landmarks
 - \simeq 20-50 convolutions through the data
 - Is it worth it?

Our work:

- Miccai2017: proof of concept

Our work:

- Miccai2017: proof of concept
- ShapeMi2018:
 - link with statistics and **computer graphics**
 - reference **implementation** on sparse data
 - theoretical **guarantees**

Our work:


- Miccai2017: proof of concept
- ShapeMi2018:
 - link with statistics and **computer graphics**
 - reference **implementation** on sparse data
 - theoretical **guarantees**
- 2019:
 - **evaluation** in varied settings
 - separable **volumetric** implementation

Thank you for your attention.

Any questions ?


**An idea from statistics:
Kernel distances**

Kernel fidelities: the simplest formula for $d(\alpha, \beta)$



Raw signal ($\alpha - \beta$).

Kernel fidelities: the simplest formula for $d(\alpha, \beta)$




Blurred signal $g \star (\alpha - \beta)$.

Choose a symmetric blurring function g , a **kernel** $k = g \star g$:

$$d_k(\alpha, \beta) = \frac{1}{2} \|g \star \alpha - g \star \beta\|_{L^2}^2$$

Kernel fidelities: the simplest formula for $d(\alpha, \beta)$




Blurred signal $g \star (\alpha - \beta)$.

Choose a symmetric blurring function g , a **kernel** $k = g \star g$:

$$\begin{aligned}d_k(\alpha, \beta) &= \frac{1}{2} \|g \star \alpha - g \star \beta\|_{L^2}^2 \\ &= \frac{1}{2} \langle \alpha - \beta | k \star (\alpha - \beta) \rangle\end{aligned}$$

Kernel fidelities: the simplest formula for $d(\alpha, \beta)$




Blurred signal $g \star (\alpha - \beta)$.

Choose a symmetric blurring function g , a **kernel** $k = g \star g$:

$$\begin{aligned}d_k(\alpha, \beta) &= \frac{1}{2} \|g \star \alpha - g \star \beta\|_{L^2}^2 \\&= \frac{1}{2} \langle \alpha - \beta | k \star (\alpha - \beta) \rangle \\&= - \sum_{i,j} k(x_i, y_j) \alpha_i \beta_j + \dots\end{aligned}$$

Kernel fidelities: the simplest formula for $d(\alpha, \beta)$



Blurred signal $g \star (\alpha - \beta)$.


Choose a symmetric blurring function g , a **kernel** $k = g \star g$:

$$\begin{aligned}d_k(\alpha, \beta) &= \frac{1}{2} \|g \star \alpha - g \star \beta\|_{L^2}^2 \\&= \frac{1}{2} \langle \alpha - \beta | k \star (\alpha - \beta) \rangle \\&= - \sum_{i,j} k(x_i, y_j) \alpha_i \beta_j + \dots \\&= \frac{1}{2} \langle \alpha - \beta | b^k - a^k \rangle\end{aligned}$$

with $a^k = -k \star \alpha$, $b^k = -k \star \beta$.


The registration flows along the gradient of $b^k - a^k$

$$k(x - y) = \exp(-\|x - y\|^2 / .2^2)$$




The registration flows along the gradient of $b^k - a^k$

$$k(x-y) = \exp(-\|x-y\|^2 / .2^2)$$




The registration flows along the gradient of $b^k - a^k$

$$k(x-y) = \exp(-\|x-y\|^2 / .2^2)$$




The registration flows along the gradient of $b^k - a^k$

$$k(x-y) = \exp(-\|x-y\|^2 / .2^2)$$



The registration flows along the gradient of $b^k - a^k$

$$k(x-y) = \exp(-\|x-y\|^2 / .2^2)$$




$$d_k(\alpha, \beta) = \frac{1}{2} \langle \alpha - \beta | k \star (\alpha - \beta) \rangle$$

$$\nabla_{x_i} d_k(\alpha, \beta) = \nabla [k \star (\alpha - \beta)](x_i) = \nabla b^k(x_i) - \nabla a^k(x_i)$$

The registration flows along the gradient of $b^k - a^k$

$$k(x - y) = \exp(-\|x - y\| / .2)$$




$$d_k(\alpha, \beta) = \frac{1}{2} \langle \alpha - \beta \mid k \star (\alpha - \beta) \rangle$$

$$\nabla_{x_i} d_k(\alpha, \beta) = \nabla [k \star (\alpha - \beta)](x_i) = \nabla b^k(x_i) - \nabla a^k(x_i)$$

The registration flows along the gradient of $b^k - a^k$

$$k(x-y) = -\|x-y\|$$




$$d_k(\alpha, \beta) = \frac{1}{2} \langle \alpha - \beta \mid k \star (\alpha - \beta) \rangle$$

$$\nabla_{x_i} d_k(\alpha, \beta) = \nabla [k \star (\alpha - \beta)](x_i) = \nabla b^k(x_i) - \nabla a^k(x_i)$$


The Energy Distance is scale-invariant, robust

$$k(x - y) = \exp(-\|x - y\|^2 / .1^2)$$



The Energy Distance is scale-invariant, robust

$$k(x - y) = -\|x - y\|$$



The SoftMin interpolates between a sum and a minimum


$$\log(e^c + e^d) = \max(c, d) + \log(\underbrace{e^{c-\max(c,d)} + e^{d-\max(c,d)}}_{\in [1,2]})$$

The SoftMin interpolates between a sum and a minimum

$$\log(e^c + e^d) = \max(c, d) + \log(\underbrace{e^{c-\max(c,d)} + e^{d-\max(c,d)}}_{\in [1,2]})$$

Building on this, for a regularization parameter $\varepsilon > 0$, we define

$$b^\varepsilon(\mathbf{x}) = \min_{y \sim \beta}^\varepsilon \|\mathbf{x} - \mathbf{y}\| = -\varepsilon \log \sum_{j=1}^M \beta_j \exp\left(-\frac{1}{\varepsilon} \|\mathbf{x} - \mathbf{y}_j\|\right)$$



$b^\varepsilon(\mathbf{x})$, with $\beta = \frac{1}{2}\delta_{.25} + \frac{1}{2}\delta_{.75}$

An idea from computer graphics:
Hausdorff distances

Can we go further?

$$\begin{matrix} & \beta_1 & \beta_2 & \cdots & \beta_M \\ \alpha_1 & \left(\begin{array}{cccc} \|\mathbf{x}_1 - \mathbf{y}_1\| & \|\mathbf{x}_1 - \mathbf{y}_2\| & \cdots & \|\mathbf{x}_1 - \mathbf{y}_M\| \\ \|\mathbf{x}_2 - \mathbf{y}_1\| & \|\mathbf{x}_2 - \mathbf{y}_2\| & \cdots & \|\mathbf{x}_2 - \mathbf{y}_M\| \\ \vdots & \vdots & \ddots & \vdots \\ \|\mathbf{x}_N - \mathbf{y}_1\| & \|\mathbf{x}_N - \mathbf{y}_2\| & \cdots & \|\mathbf{x}_N - \mathbf{y}_M\| \end{array} \right) \end{matrix}$$

Can we go further?

$$\begin{matrix} & \beta_1 & \beta_2 & \cdots & \beta_M \\ \alpha_1 & \left(\begin{array}{cccc} \|x_1 - y_1\| & \|x_1 - y_2\| & \cdots & \|x_1 - y_M\| \\ \|x_2 - y_1\| & \|x_2 - y_2\| & \cdots & \|x_2 - y_M\| \\ \vdots & \vdots & \ddots & \vdots \\ \|x_N - y_1\| & \|x_N - y_2\| & \cdots & \|x_N - y_M\| \end{array} \right) \\ \alpha_2 & & & & \\ \vdots & & & & \\ \alpha_N & & & & \end{matrix}$$

Can we go further?

$$\begin{array}{l} \alpha_1 \\ \alpha_2 \\ \vdots \\ \alpha_N \end{array} \begin{pmatrix} \beta_1 & \beta_2 & \cdots & \beta_M \\ \|\mathbf{x}_1 - \mathbf{y}_1\| & \|\mathbf{x}_1 - \mathbf{y}_2\| & \cdots & \|\mathbf{x}_1 - \mathbf{y}_M\| \\ \|\mathbf{x}_2 - \mathbf{y}_1\| & \|\mathbf{x}_2 - \mathbf{y}_2\| & \cdots & \|\mathbf{x}_2 - \mathbf{y}_M\| \\ \vdots & \vdots & \ddots & \vdots \\ \|\mathbf{x}_N - \mathbf{y}_1\| & \|\mathbf{x}_N - \mathbf{y}_2\| & \cdots & \|\mathbf{x}_N - \mathbf{y}_M\| \end{pmatrix} \rightarrow \begin{array}{l} \sum_j \beta_j \|\mathbf{x}_1 - \mathbf{y}_j\| \\ \sum_j \beta_j \|\mathbf{x}_2 - \mathbf{y}_j\| \\ \vdots \\ \sum_j \beta_j \|\mathbf{x}_N - \mathbf{y}_j\| \end{array}$$

$$\text{Energy Distance} \quad : \quad \sum_j \beta_j \|\mathbf{x}_i - \mathbf{y}_j\| = b^k(\mathbf{x}_i)$$

Can we go further?

$$\begin{array}{l} \alpha_1 \\ \alpha_2 \\ \vdots \\ \alpha_N \end{array} \begin{pmatrix} \beta_1 & \beta_2 & \cdots & \beta_M \\ \|\mathbf{x}_1 - \mathbf{y}_1\| & \|\mathbf{x}_1 - \mathbf{y}_2\| & \cdots & \|\mathbf{x}_1 - \mathbf{y}_M\| \\ \|\mathbf{x}_2 - \mathbf{y}_1\| & \|\mathbf{x}_2 - \mathbf{y}_2\| & \cdots & \|\mathbf{x}_2 - \mathbf{y}_M\| \\ \vdots & \vdots & \ddots & \vdots \\ \|\mathbf{x}_N - \mathbf{y}_1\| & \|\mathbf{x}_N - \mathbf{y}_2\| & \cdots & \|\mathbf{x}_N - \mathbf{y}_M\| \end{pmatrix} \rightarrow \begin{array}{l} \min_j \|\mathbf{x}_1 - \mathbf{y}_j\| \\ \min_j \|\mathbf{x}_2 - \mathbf{y}_j\| \\ \vdots \\ \min_j \|\mathbf{x}_N - \mathbf{y}_j\| \end{array}$$

$$\text{Energy Distance} \quad : \quad \sum_j \beta_j \|\mathbf{x}_i - \mathbf{y}_j\| = b^k(\mathbf{x}_i)$$

$$\text{Hausdorff Distance} \quad : \quad \min_j \|\mathbf{x}_i - \mathbf{y}_j\| = d(\mathbf{x}_i, \text{supp}(\beta))$$

Can we go further?


$$\begin{array}{l} \alpha_1 \\ \alpha_2 \\ \vdots \\ \alpha_N \end{array} \begin{pmatrix} \beta_1 & \beta_2 & \cdots & \beta_M \\ \|\mathbf{x}_1 - \mathbf{y}_1\| & \|\mathbf{x}_1 - \mathbf{y}_2\| & \cdots & \|\mathbf{x}_1 - \mathbf{y}_M\| \\ \|\mathbf{x}_2 - \mathbf{y}_1\| & \|\mathbf{x}_2 - \mathbf{y}_2\| & \cdots & \|\mathbf{x}_2 - \mathbf{y}_M\| \\ \vdots & \vdots & \ddots & \vdots \\ \|\mathbf{x}_N - \mathbf{y}_1\| & \|\mathbf{x}_N - \mathbf{y}_2\| & \cdots & \|\mathbf{x}_N - \mathbf{y}_M\| \end{pmatrix} \rightarrow \begin{array}{l} \min_{y \sim \beta}^\epsilon \|\mathbf{x}_1 - \mathbf{y}\| \\ \min_{y \sim \beta}^\epsilon \|\mathbf{x}_2 - \mathbf{y}\| \\ \vdots \\ \min_{y \sim \beta}^\epsilon \|\mathbf{x}_N - \mathbf{y}\| \end{array}$$

$$\text{Energy Distance} \quad : \quad \sum_j \beta_j \|\mathbf{x}_i - \mathbf{y}_j\| = b^k(\mathbf{x}_i)$$


$$\epsilon\text{-SoftMin} \quad : \quad \min_{y \sim \beta}^\epsilon \|\mathbf{x}_i - \mathbf{y}\| = b^\epsilon(\mathbf{x}_i)$$

$$\text{Hausdorff Distance} \quad : \quad \min_j \|\mathbf{x}_i - \mathbf{y}_j\| = d(\mathbf{x}_i, \text{supp}(\beta))$$

The SoftMin fidelity interpolates between Hausdorff and ED




Hausdorff, min




Kernel, \sum


The SoftMin fidelity interpolates between Hausdorff and ED



Hausdorff, \min
 $\varepsilon = 0$




SoftMin, \min^ε




Kernel, \sum
 $\varepsilon = +\infty$

$$\max^\varepsilon(c, d) = \varepsilon \log \left(\exp\left(\frac{c}{\varepsilon}\right) + \exp\left(\frac{d}{\varepsilon}\right) \right)$$


The SoftMin fidelity interpolates between Hausdorff and ED



Hausdorff, min
 $\varepsilon = 0$



SoftMin, \min^ε



Kernel, \sum
 $\varepsilon = +\infty$

$$\max^\varepsilon(c, d) = \varepsilon \log \left(\exp\left(\frac{c}{\varepsilon}\right) + \exp\left(\frac{d}{\varepsilon}\right) \right)$$

You can also use it with e.g. $\|x - y\|^2$ instead of $\|x - y\|$.

Our papers:



- *Global divergences between measures: from Hausdorff distance to Optimal Transport*, F., Trouvé, 2018



Our papers:

- *Global divergences between measures: from Hausdorff distance to Optimal Transport*, F., Trouvé, 2018
- *Sinkhorn entropies and divergences*, F., Séjourné, Vialard, Amari, Trouvé, Peyré, 2018

Our papers:

- *Global divergences between measures: from Hausdorff distance to Optimal Transport*, F., Trouvé, 2018
- *Sinkhorn entropies and divergences*, F., Séjourné, Vialard, Amari, Trouvé, Peyré, 2018
- *Optimal Transport for diffeomorphic registration*, F., Charlier, Vialard, Peyré, 2017

-  Chizat, L., Peyré, G., Schmitzer, B., and Vialard, F.-X. (2018). **Unbalanced optimal transport: Dynamic and kantorovich formulations.**
Journal of Functional Analysis, 274(11):3090–3123.
-  Cuturi, M. (2013). **Sinkhorn distances: Lightspeed computation of optimal transport.**
In *Advances in neural information processing systems*, pages 2292–2300.

-  Kaltenmark, I., Charlier, B., and Charon, N. (2017).
A general framework for curve and surface comparison and registration with oriented varifolds.
In Computer Vision and Pattern Recognition (CVPR).
-  Peyré, G. and Cuturi, M. (2018).
Computational optimal transport.
arXiv preprint arXiv:1803.00567.