

Robust shape matching with Optimal Transport

Jean Feydy

Télécom Paristech – 15th November, 2018


Écoles Normales Supérieures de Paris et Paris-Saclay

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

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


What is shape matching?

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




What is shape matching?

Source **A**, target **B**, mapping φ




What is shape matching?

Source **A**, target **B**, mapping φ



What is shape matching?


Source A , target B , mapping φ



What is shape matching?

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


First setting: processing of point clouds



- φ is **rigid** or affine
- Occlusions
- Outliers

From the documentation of the
Point Cloud Library.


Second setting: medical imaging



- φ is a spline or a **diffeomorphism**
- Ill-posed problem
- Some occlusions

From Marc Niethammer's
Quicksilver slides.


Third setting: training a generative model



- φ is a **neural network**
- Very weak regularization
- High-dimensional space

Wasserstein Auto-Encoders,
Tolstikhin et al., 2018.

Third setting: training a generative model




- φ is a **neural network**
- Very weak regularization
- High-dimensional space

Wasserstein Auto-Encoders,
Tolstikhin et al., 2018.


Which **Loss** function
should we use?

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:

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

Computing fidelities between **measures**:


1. **Computer graphics**: weighted Hausdorff distance
2. **Statistics**: kernel distances
3. **Optimal Transport**: Wasserstein distance
 \simeq Robust Point Matching

Computing fidelities between **measures**:


1. **Computer graphics**: weighted Hausdorff distance
2. **Statistics**: kernel distances
3. **Optimal Transport**: Wasserstein distance
 \simeq Robust Point Matching
4. What's **new**, in 2018?
5. Efficient GPU routines: **KeOps**

The weighted Hausdorff distance: Iterative Closest Point algorithm

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:

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

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

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

The weighted Hausdorff distance



p -Hausdorff distance:


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

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

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

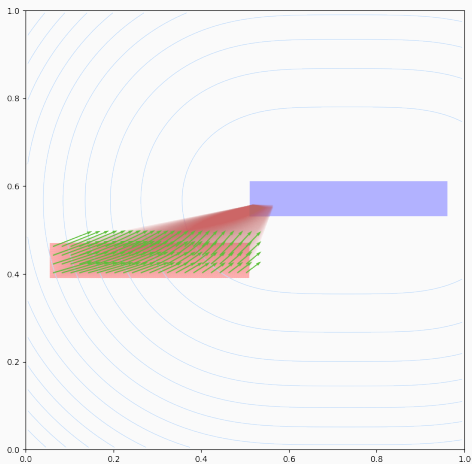
Naive projections in Hausdorff cause imbalance

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




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

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



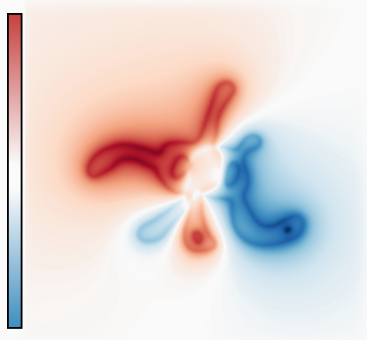
**An idea from statistics:
Kernel distances**

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** $k = g \star g$:

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

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




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

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

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

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




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

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

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

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



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

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

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

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

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

The Hausdorff distance is local, the Energy Distance is global



Hausdorff, min





Kernel, Σ

⇒ Can we get the best of both worlds?

An idea from Optimal Transport theory:
The SoftAssign algorithm

Introducing the Optimal Transport problem



Minimize over N -by- M matrices
(transport plans) π :

$$\text{OT}(\alpha, \beta) = \min_{\pi} \underbrace{\sum_{i,j} \pi_{i,j} \cdot |x_i - y_j|^2}_{\text{transport cost}}$$

subject to $\pi_{i,j} \geq 0$,

$$\sum_j \pi_{i,j} = \alpha_i, \quad \sum_i \pi_{i,j} = \beta_j.$$

Kantorovitch's dual formulation

With $C(x_i, y_j) = \|x_i - y_j\|^p$,

$$\text{OT}(\alpha, \beta) = \min_{\pi} \langle \pi, C \rangle \quad \longrightarrow \text{Assignment}$$

s.t. $\pi \geq 0, \quad \pi \mathbf{1} = \alpha, \quad \pi^T \mathbf{1} = \beta$

Kantorovitch's dual formulation

With $C(x_i, y_j) = \|x_i - y_j\|^p$,

$$\text{OT}(\alpha, \beta) = \min_{\pi} \langle \pi, C \rangle \quad \longrightarrow \text{Assignment}$$

$$\text{s.t. } \pi \geq 0, \quad \pi \mathbf{1} = \alpha, \quad \pi^T \mathbf{1} = \beta$$

$$= \max_{f, g} \langle \alpha, f \rangle + \langle \beta, g \rangle \quad \longrightarrow \text{FedEx}$$

$$\text{s.t. } f(x_i) + g(y_j) \leq C(x_i, y_j),$$

Kantorovitch's dual formulation

With $C(x_i, y_j) = \|x_i - y_j\|^p$,

$$\text{OT}(\alpha, \beta) = \min_{\pi} \langle \pi, C \rangle \quad \longrightarrow \text{Assignment}$$

$$\text{s.t. } \pi \geq 0, \quad \pi \mathbf{1} = \alpha, \quad \pi^T \mathbf{1} = \beta$$

$$= \max_{f, g} \langle \alpha, f \rangle + \langle \beta, g \rangle \quad \longrightarrow \text{FedEx}$$

$$\text{s.t. } f(x_i) + g(y_j) \leq C(x_i, y_j),$$

\implies **Combinatorial** problem on the simplex

Kantorovitch's dual formulation

With $C(x_i, y_j) = \|x_i - y_j\|^p$,

$OT(\alpha, \beta) = \min_{\pi} \langle \pi, C \rangle$ \longrightarrow Assignment

$$\text{s.t. } \pi \geq 0, \quad \pi \mathbf{1} = \alpha, \quad \pi^T \mathbf{1} = \beta$$



$= \max_{f, g} \langle \alpha, f \rangle + \langle \beta, g \rangle$ \longrightarrow FedEx

$$\text{s.t. } f(x_i) + g(y_j) \leq C(x_i, y_j),$$

\implies **Combinatorial** problem on the simplex

\implies Hungarian method in $O(N^3)$.

Entropic regularization: introducing Schrödinger's problem



For $\varepsilon > 0$:

$$\text{OT}_\varepsilon(\alpha, \beta) = \min_{\pi} \underbrace{\sum_{i,j} \pi_{i,j} \cdot |x_i - y_j|^2}_{\text{transport cost}} + \varepsilon \underbrace{\sum_{i,j} \pi_{i,j} \cdot \log \frac{\pi_{i,j}}{\alpha_i \beta_j}}_{\text{entropic barrier}}$$

subject to

$$\sum_j \pi_{i,j} = \alpha_i, \quad \sum_i \pi_{i,j} = \beta_j.$$

$$\text{OT}_\varepsilon(\alpha, \beta) = \min_{\pi} \langle \pi, C \rangle + \varepsilon \text{KL}(\pi, \alpha \otimes \beta) \longrightarrow \text{Fuzzy assignment}$$

s.t. $\pi \mathbf{1} = \alpha, \quad \pi^T \mathbf{1} = \beta$

Fenchel-Rockafellar to the rescue

$$\text{OT}_\varepsilon(\alpha, \beta) = \min_{\pi} \langle \pi, C \rangle + \varepsilon \text{KL}(\pi, \alpha \otimes \beta) \longrightarrow \text{Fuzzy assignment}$$

$$\text{s.t.} \quad \pi \mathbf{1} = \alpha, \quad \pi^T \mathbf{1} = \beta$$

$$= \max_{f, g} \langle \alpha, f \rangle + \langle \beta, g \rangle \longrightarrow \text{Cheeky FedEx}$$

$$- \underbrace{\varepsilon \langle \alpha \otimes \beta, e^{(f \oplus g - C)/\varepsilon} - 1 \rangle}_{\text{soft constraint } f \oplus g \leq C}$$

Fenchel-Rockafellar to the rescue

$$\text{OT}_\varepsilon(\alpha, \beta) = \min_{\pi} \langle \pi, C \rangle + \varepsilon \text{KL}(\pi, \alpha \otimes \beta) \longrightarrow \text{Fuzzy assignment}$$

$$\text{s.t.} \quad \pi \mathbf{1} = \alpha, \quad \pi^T \mathbf{1} = \beta$$

$$= \max_{f, g} \langle \alpha, f \rangle + \langle \beta, g \rangle \longrightarrow \text{Cheeky FedEx}$$

$$- \underbrace{\varepsilon \langle \alpha \otimes \beta, e^{(f \oplus g - C)/\varepsilon} - 1 \rangle}_{\text{soft constraint } f \oplus g \leq C}$$

\implies **Strictly convex** problem on the simplex

Fenchel-Rockafellar to the rescue

$$\text{OT}_\varepsilon(\alpha, \beta) = \min_{\pi} \langle \pi, C \rangle + \varepsilon \text{KL}(\pi, \alpha \otimes \beta) \longrightarrow \text{Fuzzy assignment}$$

$$\text{s.t.} \quad \pi \mathbf{1} = \alpha, \quad \pi^T \mathbf{1} = \beta$$

$$= \max_{f, g} \langle \alpha, f \rangle + \langle \beta, g \rangle \longrightarrow \text{Cheeky FedEx}$$


$$- \underbrace{\varepsilon \langle \alpha \otimes \beta, e^{(f \oplus g - C)/\varepsilon} - 1 \rangle}_{\text{soft constraint } f \oplus g \leq C}$$

\implies **Strictly convex** problem on the simplex


$$\text{At the optimum, } \pi = e^{(f \oplus g - C)/\varepsilon} \cdot \alpha \otimes \beta$$

$$\text{i.e.} \quad \pi_{i,j} = \alpha_i e^{f_i/\varepsilon} e^{-C(x_i, y_j)/\varepsilon} e^{g_j/\varepsilon} \beta_j.$$

Textbook interpretation: balancing of a kernel matrix



=



$$\pi_{ij} = \Delta(U\alpha) \cdot K_{x,y} \cdot \Delta(V\beta)$$

with

- a kernel function k


$$k(x_i - y_j) = e^{-C(x_i, y_j)/\epsilon}$$

- $U = e^{f/\epsilon}$ and $V = e^{g/\epsilon}$,
positive weights on
 $\{x_i\}$ and $\{y_j\}$.

→ Enforce the **constraints**

$$\pi \mathbf{1} = \alpha, \quad \pi^T \mathbf{1} = \beta$$

Enforcing $\pi \mathbf{1} = \alpha$ and $\pi^T \mathbf{1} = \beta$ alternatively



Source and target.

Sinkhorn Iterative Algorithm

Input : source $\alpha = \sum_i \alpha_i \delta_{x_i}$

target $\beta = \sum_j \beta_j \delta_{y_j}$

Parameter : $k : x \mapsto e^{-|x|^2/\epsilon}$

1: $U \leftarrow \text{ones}(\text{size}(\alpha))$

2: $V \leftarrow \text{ones}(\text{size}(\beta))$

3: **while** updates $>$ tol **do**


4: $U \leftarrow \mathbf{1} ./ K \cdot (V\beta)$

5: $V \leftarrow \mathbf{1} ./ K^T \cdot (U\alpha)$

6: **return** $\epsilon (\langle \alpha, \log(U) \rangle + \langle \beta, \log(V) \rangle)$

Output : fidelity $\text{OT}_\epsilon(\alpha, \beta)$

Enforcing $\pi\mathbf{1} = \alpha$ and $\pi^T\mathbf{1} = \beta$ alternatively



Seen by the kernel k .

Sinkhorn Iterative Algorithm

Input : source $\alpha = \sum_i \alpha_i \delta_{x_i}$

target $\beta = \sum_j \beta_j \delta_{y_j}$

Parameter : $k : x \mapsto e^{-|x|^2/\epsilon}$

1: $U \leftarrow \text{ones}(\text{size}(\alpha))$

2: $V \leftarrow \text{ones}(\text{size}(\beta))$

3: **while** updates $>$ tol **do**


4: $U \leftarrow \mathbf{1} ./ K \cdot (V\beta)$

5: $V \leftarrow \mathbf{1} ./ K^T \cdot (U\alpha)$

6: **return** $\epsilon (\langle \alpha, \log(U) \rangle + \langle \beta, \log(V) \rangle)$

Output : fidelity $\text{OT}_\epsilon(\alpha, \beta)$

Enforcing $\pi \mathbf{1} = \alpha$ and $\pi^T \mathbf{1} = \beta$ alternatively



Sinkhorn Iteration 000

Starting estimate.

Sinkhorn Iterative Algorithm

Input : source $\alpha = \sum_i \alpha_i \delta_{x_i}$


target $\beta = \sum_j \beta_j \delta_{y_j}$

Parameter : $k : x \mapsto e^{-|x|^2/\epsilon}$

- 1: $U \leftarrow \text{ones}(\text{size}(\alpha))$
- 2: $V \leftarrow \text{ones}(\text{size}(\beta))$
- 3: while updates $>$ tol do
- 4: $U \leftarrow \mathbf{1} ./ K \cdot (V\beta)$
- 5: $V \leftarrow \mathbf{1} ./ K^T \cdot (U\alpha)$
- 6: return $\epsilon (\langle \alpha, \log(U) \rangle + \langle \beta, \log(V) \rangle)$

Output : fidelity $\text{OT}_\epsilon(\alpha, \beta)$

Enforcing $\pi \mathbf{1} = \alpha$ and $\pi^T \mathbf{1} = \beta$ alternatively



Sinkhorn Iteration 250

Computing the OT plan.

Sinkhorn Iterative Algorithm

Input : source $\alpha = \sum_i \alpha_i \delta_{x_i}$

target $\beta = \sum_j \beta_j \delta_{y_j}$

Parameter : $k : x \mapsto e^{-|x|^2/\epsilon}$

1: $U \leftarrow \text{ones}(\text{size}(\alpha))$

2: $V \leftarrow \text{ones}(\text{size}(\beta))$

3: **while** updates $>$ tol **do**


4: $U \leftarrow \mathbf{1} ./ K \cdot (V\beta)$

5: $V \leftarrow \mathbf{1} ./ K^T \cdot (U\alpha)$

6: **return** $\epsilon (\langle \alpha, \log(U) \rangle + \langle \beta, \log(V) \rangle)$

Output : fidelity $\text{OT}_\epsilon(\alpha, \beta)$

Enforcing $\pi \mathbf{1} = \alpha$ and $\pi^T \mathbf{1} = \beta$ alternatively



Sinkhorn Iteration 250

Computing the OT plan.

Sinkhorn Iterative Algorithm

Input : source $\alpha = \sum_i \alpha_i \delta_{x_i}$

target $\beta = \sum_j \beta_j \delta_{y_j}$

Parameter : $k : x \mapsto e^{-|x|^2/\epsilon}$

1: $U \leftarrow \text{ones}(\text{size}(\alpha))$

2: $V \leftarrow \text{ones}(\text{size}(\beta))$


3: **while** updates $>$ tol **do**

4: $U \leftarrow \mathbf{1} ./ K \cdot (V\beta)$

5: $V \leftarrow \mathbf{1} ./ K^T \cdot (U\alpha)$


6: **return** $\epsilon (\langle \alpha, \log(U) \rangle + \langle \beta, \log(V) \rangle)$

Output : fidelity $\text{OT}_\epsilon(\alpha, \beta)$




TPS-RPM algorithm,
Chui and Rangarajan, CVPR 2000

21




Optimal Transport for diffeomorphic registration,
Feydy et al., MICCAI 2017



TPS-RPM algorithm,
Chui and Rangarajan, CVPR 2000

21



Optimal Transport for diffeomorphic registration,
Feydy et al., MICCAI 2017

⇒ We've added weights, orientations, convergence analysis...
But shouldn't we go a bit **further**?

**It's 2018 now:
What's new?**

Fact 1 : Sinkhorn is best implemented in the log-domain

Unfortunately,

$$k(x_i, y_j) \simeq 0 \quad \text{if } \varepsilon \text{ is too small.}$$

Fact 1 : Sinkhorn is best implemented in the log-domain

Unfortunately,

$$k(x_i, y_j) \simeq 0 \quad \text{if } \varepsilon \text{ is too small.}$$

$$\begin{aligned} \text{OT}_\varepsilon(\alpha, \beta) = \max_{f, g} & \langle \alpha, f \rangle + \langle \beta, g \rangle && \longrightarrow \text{Cheeky FedEx} \\ & - \underbrace{\varepsilon \langle \alpha \otimes \beta, e^{(f \oplus g - C)/\varepsilon} - 1 \rangle}_{\text{soft constraint } f \oplus g \leq C} \end{aligned}$$

Fact 1 : Sinkhorn is best implemented in the log-domain

Unfortunately,

$$k(\mathbf{x}_i, \mathbf{y}_j) \simeq 0 \quad \text{if } \varepsilon \text{ is too small.}$$

$$\begin{aligned} \text{OT}_\varepsilon(\boldsymbol{\alpha}, \boldsymbol{\beta}) &= \max_{f, g} \langle \boldsymbol{\alpha}, f \rangle + \langle \boldsymbol{\beta}, g \rangle && \longrightarrow \text{Cheeky FedEx} \\ &\quad - \underbrace{\varepsilon \langle \boldsymbol{\alpha} \otimes \boldsymbol{\beta}, e^{(f \oplus g - C)/\varepsilon} - 1 \rangle}_{\text{soft constraint } f \oplus g \leq C} \end{aligned}$$

Equivalent to the constraints on π , the optimality conditions read:

$$\begin{aligned} f(\mathbf{x}_i) &= -\varepsilon \log \sum_j \beta_j \exp \frac{1}{\varepsilon} (g(\mathbf{y}_j) - C(\mathbf{x}_i, \mathbf{y}_j)), \\ g(\mathbf{y}_j) &= -\varepsilon \log \sum_i \alpha_i \exp \frac{1}{\varepsilon} (f(\mathbf{x}_i) - C(\mathbf{x}_i, \mathbf{y}_j)). \end{aligned}$$

The SoftMin interpolates between a minimum and a sum


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

The SoftMin interpolates between a minimum and a sum

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

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

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



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

A new reduction on the distance matrix

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

A new reduction on the distance matrix

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

A new reduction on the distance matrix

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

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

A new reduction on the distance matrix

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

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

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

A new reduction on the distance matrix

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

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

$$\varepsilon\text{-SoftMin} \quad : \quad \min_{y \sim \beta} \|\mathbf{x}_i - \mathbf{y}\| = b_\varepsilon(\mathbf{x}_i) \simeq f(\mathbf{x}_i)$$

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

Optimal Transport = Hausdorff + mass spreading constraint

The optimality conditions read:

$$f(x_i) = b(x) = -\varepsilon \log \sum_j \beta_j \exp \frac{1}{\varepsilon} [g(y_j) - C(x_i, y_j)],$$

$$g(y_j) = a(y) = -\varepsilon \log \sum_i \alpha_i \exp \frac{1}{\varepsilon} [f(x_i) - C(x_i, y_j)].$$

Optimal Transport = Hausdorff + mass spreading constraint

The optimality conditions read:

$$f(x_i) = b(x) = \min_{y \sim \beta}^{\epsilon} [C(x,y) - a(y)] \quad ,$$

$$g(y_j) = a(y) = \min_{x \sim \alpha}^{\epsilon} [C(x,y) - b(x)] \quad .$$

Optimal Transport = Hausdorff + mass spreading constraint

The optimality conditions read:

$$f(x_i) = b(x) = \min_{y \sim \beta}^{\varepsilon} [C(x,y) - a(y)] \quad ,$$

$$g(y_j) = a(y) = \min_{x \sim \alpha}^{\varepsilon} [C(x,y) - b(x)] \quad .$$

Final cost:

$$\text{OT}_{\varepsilon}(\alpha, \beta) = \langle \alpha, f \rangle + \langle \beta, g \rangle = \langle \alpha, b \rangle + \langle \beta, a \rangle.$$

Optimal Transport = Hausdorff + mass spreading constraint

The optimality conditions read:

$$f(x_i) = b(x) = \min_{y \sim \beta}^{\varepsilon} [C(x,y) - a(y)] \quad ,$$

$$g(y_j) = a(y) = \min_{x \sim \alpha}^{\varepsilon} [C(x,y) - b(x)] \quad .$$

Final cost:

$$\text{OT}_{\varepsilon}(\alpha, \beta) = \langle \alpha, f \rangle + \langle \beta, g \rangle = \langle \alpha, b \rangle + \langle \beta, a \rangle.$$


Discrete, 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 .


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 $\sqrt{\varepsilon} = .1$




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 $\sqrt{\varepsilon} = .1$



Fact 2 : if ε is too small, the Sinkhorn algorithm does not converge

If $\varepsilon = 0$: the Sinkhorn loop gets **stuck** after two iterations.

Fact 2 : if ε is too small, the Sinkhorn algorithm does not converge

If $\varepsilon = 0$: the Sinkhorn loop gets **stuck** after two iterations.

If $\varepsilon > 0$: it is a fixed-point iteration that converges linearly...

Fact 2 : if ε is too small, the Sinkhorn algorithm does not converge

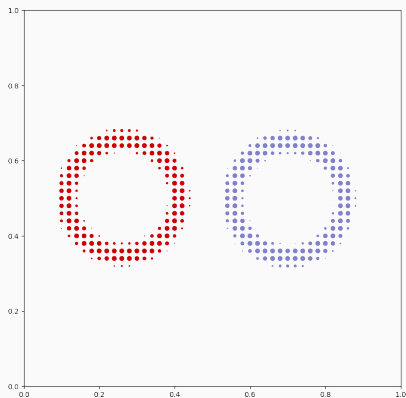
If $\varepsilon = 0$: the Sinkhorn loop gets **stuck** after two iterations.

If $\varepsilon > 0$: it is a fixed-point iteration that converges linearly...

But even in simple cases, we only converge in $\mathbf{O}((1 - \varepsilon)^n)$:
Computing a true Wasserstein distance OT_0 is out-of-reach.


Fact 3 : if $\varepsilon > 0$, OT_ε is not a valid divergence

Registrating circles, $C(x,y) = \|x - y\|^2$, $\sqrt{\varepsilon} = 0.1$:




Fact 3 : if $\varepsilon > 0$, OT_ε is not a valid divergence

Registrating circles, $C(x,y) = \|x - y\|^2$, $\sqrt{\varepsilon} = 0.1$:




Fact 3 : if $\varepsilon > 0$, OT_ε is not a valid divergence

Registrating circles, $C(x,y) = \|x - y\|^2$, $\sqrt{\varepsilon} = 0.1$:




Fact 3 : if $\varepsilon > 0$, OT_ε is not a valid divergence

Registrating circles, $C(x,y) = \|x - y\|^2$, $\sqrt{\varepsilon} = 0.2$:




Fact 3 : if $\varepsilon > 0$, OT_ε is not a valid divergence

Registrating circles, $C(x,y) = \|x - y\|^2$, $\sqrt{\varepsilon} = 0.2$:




Fact 3 : if $\varepsilon > 0$, OT_ε is not a valid divergence

Registrating circles, $C(x,y) = \|x - y\|^2$, $\sqrt{\varepsilon} = 0.2$:



Fact 3 : if $\varepsilon > 0$, OT_ε is not a valid divergence


Registrating circles, $C(x,y) = \|x - y\|^2$, $\sqrt{\varepsilon} = 0.2$:



Bad news: for $0 < \varepsilon \leq +\infty$, we converge towards α such that


$$OT_\varepsilon(\alpha, \beta) < OT_\varepsilon(\beta, \beta).$$

Standard solution: use an annealing scheme



TPS-RPM algorithm, Chui and Rangarajan, CVPR 2000

Standard solution: use an annealing scheme



TPS-RPM algorithm, Chui and Rangarajan, CVPR 2000

⇒ **Expensive** and cumbersome workaround,
with parameters to tune.

A new idea in 2017 : un-biased Sinkhorn divergences

$$\text{OT}_\varepsilon(\alpha, \beta) = \min_{\pi} \langle \pi, \mathbf{C} \rangle + \varepsilon \text{KL}(\pi, \alpha \otimes \beta) \longrightarrow \text{Fuzzy assignment}$$

s.t. $\pi \mathbf{1} = \alpha, \quad \pi^T \mathbf{1} = \beta$

A new idea in 2017 : un-biased Sinkhorn divergences

$$\text{OT}_\varepsilon(\alpha, \beta) = \min_{\pi} \langle \pi, \mathbf{C} \rangle + \varepsilon \text{KL}(\pi, \alpha \otimes \beta) \longrightarrow \text{Fuzzy assignment}$$

s.t. $\pi \mathbf{1} = \alpha, \quad \pi^T \mathbf{1} = \beta$

$$\text{OT}_\varepsilon(\alpha, \beta) \xrightarrow{\varepsilon \rightarrow +\infty} \langle \alpha \otimes \beta, \mathbf{C} \rangle = \langle \alpha, \mathbf{C} \star \beta \rangle$$

A new idea in 2017 : un-biased Sinkhorn divergences

$$\text{OT}_\varepsilon(\alpha, \beta) = \min_{\pi} \langle \pi, \mathbf{C} \rangle + \varepsilon \text{KL}(\pi, \alpha \otimes \beta) \longrightarrow \text{Fuzzy assignment}$$

s.t. $\pi \mathbf{1} = \alpha, \quad \pi^T \mathbf{1} = \beta$

$$\text{OT}_\varepsilon(\alpha, \beta) \xrightarrow{\varepsilon \rightarrow +\infty} \langle \alpha \otimes \beta, \mathbf{C} \rangle = \langle \alpha, \mathbf{C} \star \beta \rangle$$

Define the **Sinkhorn divergence** [Raudas et al., 2017]:

$$S_\varepsilon(\alpha, \beta) = \text{OT}_\varepsilon(\alpha, \beta) - \frac{1}{2} \text{OT}_\varepsilon(\alpha, \alpha) - \frac{1}{2} \text{OT}_\varepsilon(\beta, \beta)$$

A new idea in 2017 : un-biased Sinkhorn divergences

$$\text{OT}_\varepsilon(\alpha, \beta) = \min_{\pi} \langle \pi, C \rangle + \varepsilon \text{KL}(\pi, \alpha \otimes \beta) \longrightarrow \text{Fuzzy assignment}$$

s.t. $\pi \mathbf{1} = \alpha, \quad \pi^T \mathbf{1} = \beta$

$$\text{OT}_\varepsilon(\alpha, \beta) \xrightarrow{\varepsilon \rightarrow +\infty} \langle \alpha \otimes \beta, C \rangle = \langle \alpha, C \star \beta \rangle$$

Define the **Sinkhorn divergence** [Raudas et al., 2017]:

$$S_\varepsilon(\alpha, \beta) = \text{OT}_\varepsilon(\alpha, \beta) - \frac{1}{2} \text{OT}_\varepsilon(\alpha, \alpha) - \frac{1}{2} \text{OT}_\varepsilon(\beta, \beta)$$

$$\text{Wasserstein}_{+C}(\alpha, \beta) \xleftarrow{\varepsilon \rightarrow 0} S_\varepsilon(\alpha, \beta) \xrightarrow{\varepsilon \rightarrow +\infty} \text{Kernel}_{-C}(\alpha, \beta)$$

A new idea in 2017 : un-biased Sinkhorn divergences

$$\text{OT}_\varepsilon(\alpha, \beta) = \min_{\pi} \langle \pi, C \rangle + \varepsilon \text{KL}(\pi, \alpha \otimes \beta) \longrightarrow \text{Fuzzy assignment}$$

s.t. $\pi \mathbf{1} = \alpha, \quad \pi^T \mathbf{1} = \beta$

$$\text{OT}_\varepsilon(\alpha, \beta) \xrightarrow{\varepsilon \rightarrow +\infty} \langle \alpha \otimes \beta, C \rangle = \langle \alpha, C \star \beta \rangle$$

Define the **Sinkhorn divergence** [Raudas et al., 2017]:

$$S_\varepsilon(\alpha, \beta) = \text{OT}_\varepsilon(\alpha, \beta) - \frac{1}{2} \text{OT}_\varepsilon(\alpha, \alpha) - \frac{1}{2} \text{OT}_\varepsilon(\beta, \beta)$$

$$\text{Wasserstein}_{+C}(\alpha, \beta) \xleftarrow{\varepsilon \rightarrow 0} S_\varepsilon(\alpha, \beta) \xrightarrow{\varepsilon \rightarrow +\infty} \text{Kernel}_{-C}(\alpha, \beta)$$

In practice, S_ε is “good enough” for ML applications

[Genevay et al., 2018, Salimans et al., 2018, Sanjabi et al., 2018].

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

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

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

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

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

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

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

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

$\alpha \mapsto S_\varepsilon(\alpha, \beta)$ is convex and differentiable

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

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

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

$\alpha \mapsto S_\varepsilon(\alpha, \beta)$ is convex and differentiable

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

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

In our paper: theoretical guarantees

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


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

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


$\alpha \mapsto S_\varepsilon(\alpha, \beta)$ is convex and differentiable

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

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



Loss = OT_ε



Loss = S_ε

Conclusion

A significant result

The **true** OT_0 problem is hard.

A significant result

The **true** OT_0 problem is hard.

Approximating it with **subsampling** or **smoothing** is easy:
this is what **SoftAssign** is all about.

A significant result


The **true** OT_0 problem is hard.

Approximating it with **subsampling** or **smoothing** is easy:
this is what **SoftAssign** is all about.

Remarkably, $S_\varepsilon(\alpha, \beta)$ is a cheap approximation of $OT_0(\alpha, \beta)$
that defines a **positive definite** cost between the **full samples**.
It is the first known way of doing so.

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.

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.)

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

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

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

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

- Try using $k(x,y) = -\|x - y\|$!
- Remove the **entropic bias** from the SoftAssign algorithm!

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

- Try using $k(x,y) = -\|x - y\|$!
- Remove the **entropic bias** from the SoftAssign algorithm!
- 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/AiStats2019 :
 - link with statistics and **computer graphics**
 - reference **implementation** on sparse data
 - theoretical **guarantees**

Our work:

- Miccai2017 : proof of concept
- ShapeMi2018/AiStats2019 :
 - link with statistics and **computer graphics**
 - reference **implementation** on sparse data
 - theoretical **guarantees**
- 2019 :
 - unbalanced formulation, gestion of **outliers**
 - **evaluation** in varied settings
 - separable **volumetric** implementation

Open questions:

- **Combinatorial** interpretation of the Sinkhorn iterations?

Open questions:

- **Combinatorial** interpretation of the Sinkhorn iterations?
- Link between S_ε and Sobolev **distances**?

Open questions:

- **Combinatorial** interpretation of the Sinkhorn iterations?
- Link between S_ε and Sobolev **distances**?
- What about **Octrees**?

Open questions:


- **Combinatorial** interpretation of the Sinkhorn iterations?
- Link between S_ϵ and Sobolev **distances**?
- What about **Octrees**?
- Interest in the **CVPR/SIGGRAPH** communities?

Thank you for your attention.

Any questions ?


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 $b^k - a^k$

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

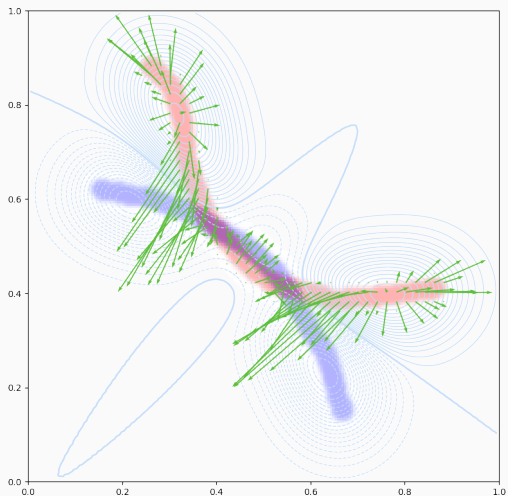


$$d_k(\alpha, \beta) = \frac{1}{2} \langle \alpha - \beta \mid k \star (\alpha - \beta) \rangle$$

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


The Energy Distance is scale-invariant, robust

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




The Energy Distance is scale-invariant, robust

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




An idea from computer graphics: Hausdorff distances

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.



Genevay, A., Peyre, G., and Cuturi, M. (2018).

Learning generative models with sinkhorn divergences.




In Storkey, A. and Perez-Cruz, F., editors, *Proceedings of the Twenty-First International Conference on Artificial Intelligence and Statistics*, volume 84 of *Proceedings of Machine Learning Research*, pages 1608–1617. PMLR.



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.
-  Ramdas, A., Trillos, N. G., and Cuturi, M. (2017).
On wasserstein two-sample testing and related families of nonparametric tests.
Entropy, 19(2).
-  Salimans, T., Zhang, H., Radford, A., and Metaxas, D. (2018).
Improving GANs using optimal transport.
arXiv preprint arXiv:1803.05573.



Sanjabi, M., Ba, J., Razaviyayn, M., and Lee, J. D. (2018).
**On the convergence and robustness of training GANs with
regularized optimal transport.**
arXiv preprint arXiv:1802.08249.