

Global divergences between measures: from Hausdorff distance to Optimal Transport

Jean Feydy Alain Trouvé

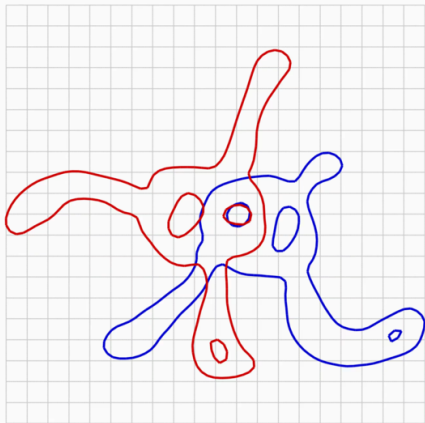
Curves and Surfaces, Arcachon – 2 juillet 2018

Écoles Normales Supérieures de Paris et Paris-Saclay

Collaboration with B. Charlier, J. Glaunès, F.-X. Vialard, G. Peyré

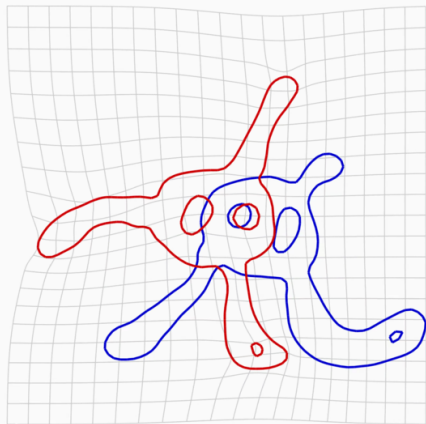
Today: focus on shape registration

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



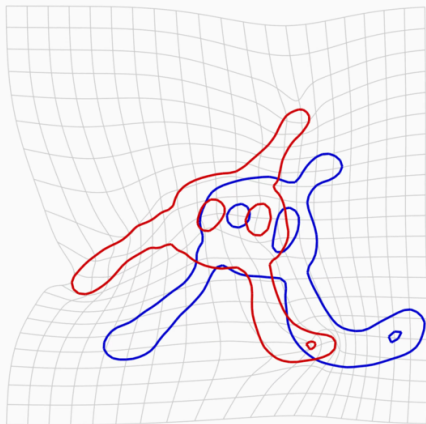
Today: focus on shape registration

Source A , target B , mapping φ



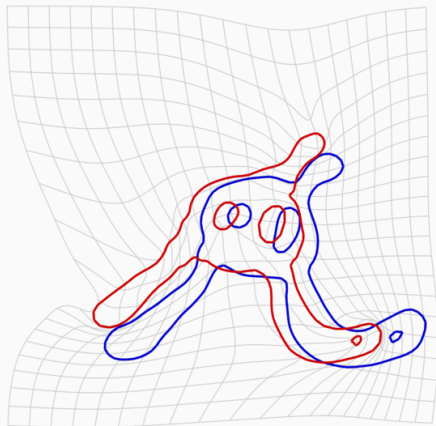
Today: focus on shape registration

Source A , target B , mapping φ



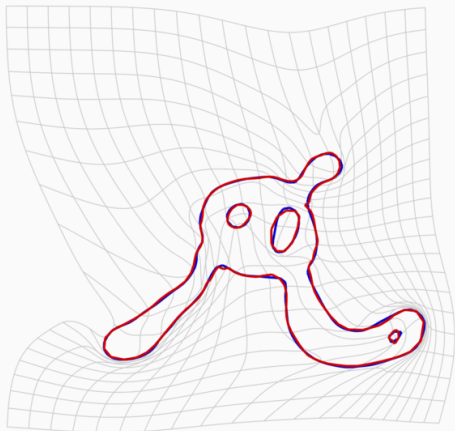
Today: focus on shape registration

Source **A**, target **B**, mapping φ



Today: focus on shape registration

Source A , target B , mapping φ



In practice: gradient descent on the deformation

$$\text{Cost}(\varphi) = \underbrace{\text{Reg}(\varphi)}_{\text{regularization}} + \underbrace{d(\varphi(A), B)}_{\text{fidelity}}$$

In practice: gradient descent on the deformation

$$\text{Cost}(\varphi) = \underbrace{\text{Reg}(\varphi)}_{\text{regularization}} + \underbrace{d(\varphi(\mathbf{A}), \mathbf{B})}_{\text{fidelity}}$$

If \mathbf{A} and \mathbf{B} are labeled vectors of $\mathbb{R}^{N \times D}$, you can use

Affine registration: $\text{Cost}(\varphi) = 1_{\text{affine}}(\varphi) + \|\varphi(\mathbf{A}) - \mathbf{B}\|_2^2$

Thin Plate Splines: $\text{Cost}(\varphi) = \lambda \|\Delta\varphi\|_2^2 + \|\varphi(\mathbf{A}) - \mathbf{B}\|_2^2$

In practice: gradient descent on the deformation

$$\text{Cost}(\varphi) = \underbrace{\text{Reg}(\varphi)}_{\text{regularization}} + \underbrace{d(\varphi(A), B)}_{\text{fidelity}}$$

Iterative Matching Algorithm

- 1: $\varphi \leftarrow \text{Id}$
- 2: **while** updates $>$ tol **do**
- 3: “ $\varphi \leftarrow \varphi - \alpha \cdot (\nabla_{\varphi} \text{Reg}(\varphi) + \nabla_{\varphi} [d(\varphi(A), B)])$ ”
- 4: **return** φ

Output: matching transformation φ .

In practice: gradient descent on the deformation

$$\text{Cost}(\varphi) = \underbrace{\text{Reg}(\varphi)}_{\text{regularization}} + \underbrace{d(\varphi(A), B)}_{\text{fidelity}}$$

Iterative Matching Algorithm

- 1: $\varphi \leftarrow \text{Id}$
- 2: **while** updates $>$ tol **do**
- 3: “ $\varphi \leftarrow \varphi - \alpha \cdot (\nabla_{\varphi} \text{Reg}(\varphi) + \nabla_{\varphi} [d(\varphi(A), B)])$ ”
- 4: **return** φ

Output: matching transformation φ .

\Rightarrow The fidelity's gradient **drives** the registration

Encoding unlabeled shapes as measures

Let's enforce sampling invariance:

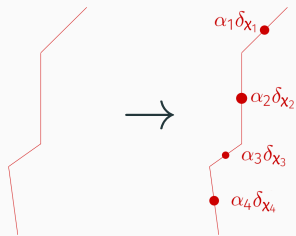
$$A \longrightarrow \alpha = \sum_{i=1}^N \alpha_i \delta_{x_i}, \quad B \longrightarrow \beta = \sum_{j=1}^M \beta_j \delta_{y_j}.$$

Encoding unlabeled shapes as measures

Let's enforce sampling invariance:

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

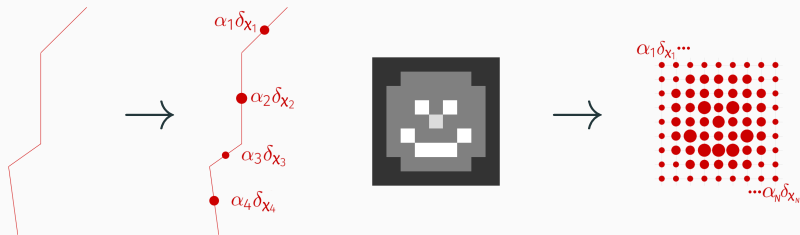
$$B \longrightarrow \beta = \sum_{j=1}^M \beta_j \delta_{y_j}.$$



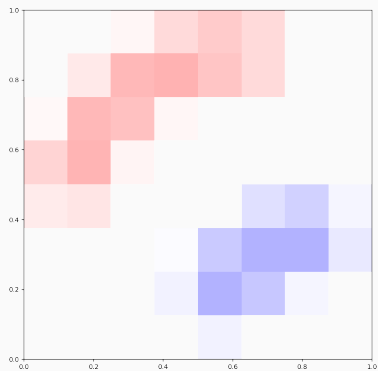
Encoding unlabeled shapes as measures

Let's enforce sampling invariance:

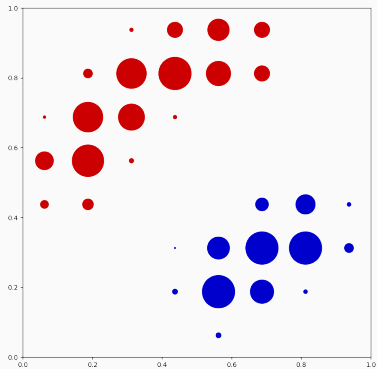
$$A \longrightarrow \alpha = \sum_{i=1}^N \alpha_i \delta_{x_i}, \quad B \longrightarrow \beta = \sum_{j=1}^M \beta_j \delta_{y_j}.$$



A baseline setting: density registration

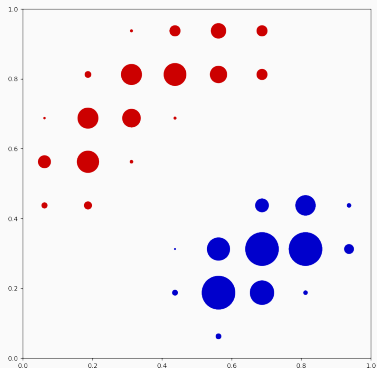


A baseline setting: density registration



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

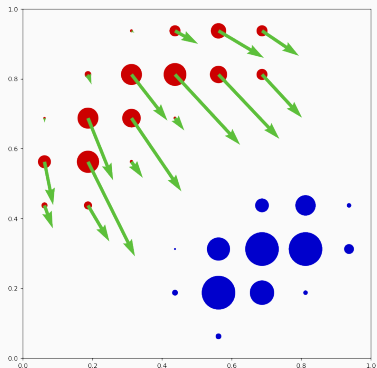
A baseline setting: density registration



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

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

A baseline setting: density registration

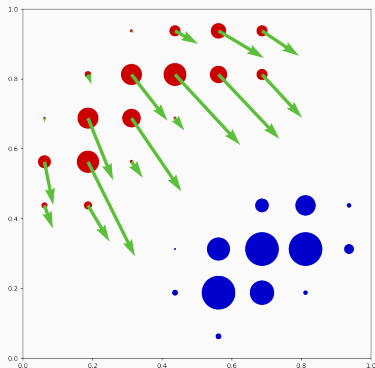


$$\alpha = \sum_{i=1}^N \alpha_i \delta_{x_i}, \quad \beta = \sum_{j=1}^M \beta_j \delta_{y_j}.$$

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

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].

Computing fidelities between **measures**:

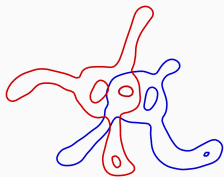
1. **Statistics**: kernel distances
2. **Computer graphics**: Hausdorff distances
3. **Optimal Transport**: Sinkhorn divergences

Computing fidelities between **measures**:

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

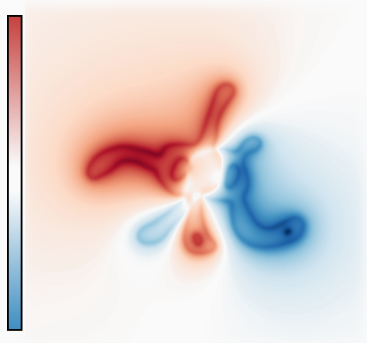
**An idea from statistics:
Kernel distances**

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



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

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

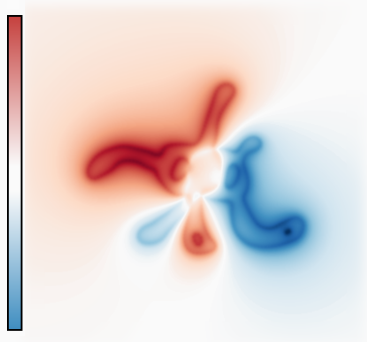


Blurred signal $g \star (\alpha - \beta)$.

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 $d(\alpha, \beta)$

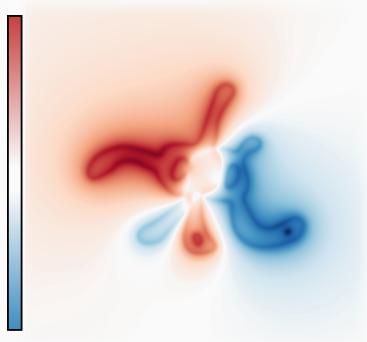


Blurred signal $g \star (\alpha - \beta)$.

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 $d(\alpha, \beta)$

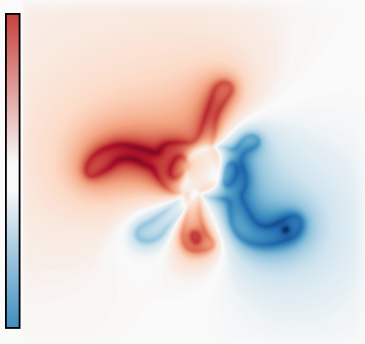


Blurred signal $g \star (\alpha - \beta)$.

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

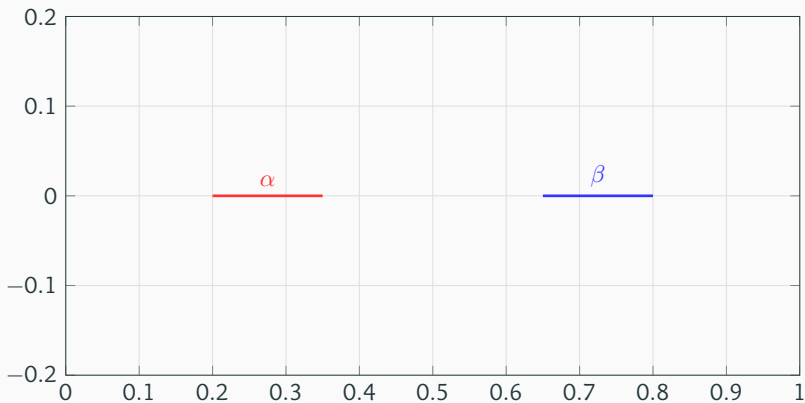
Choose a symmetric blurring function g , a **kernel** $k = g \star g$:

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

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

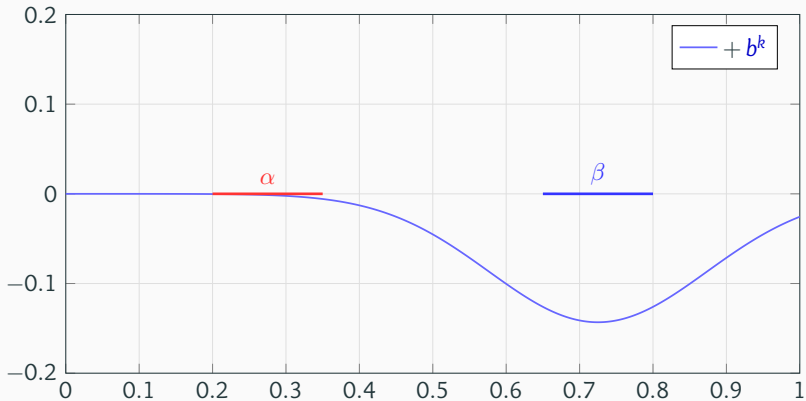
The registration flows along the gradient of $b^k - a^k$

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



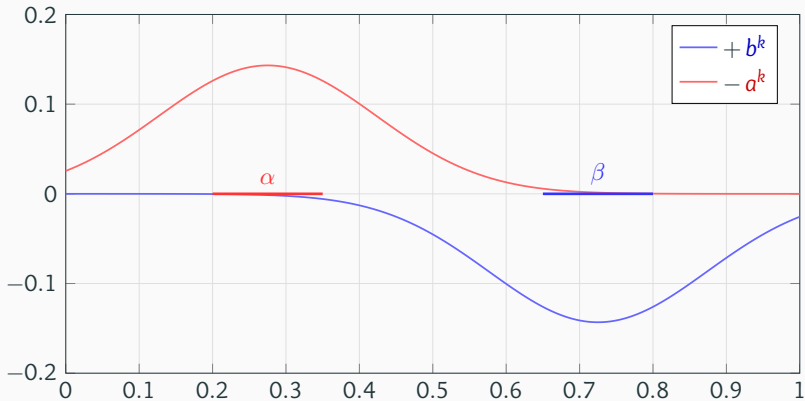
The registration flows along the gradient of $b^k - a^k$

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



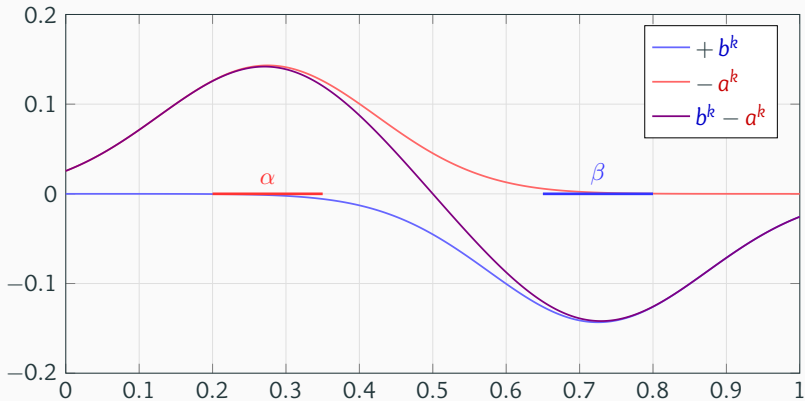
The registration flows along the gradient of $b^k - a^k$

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

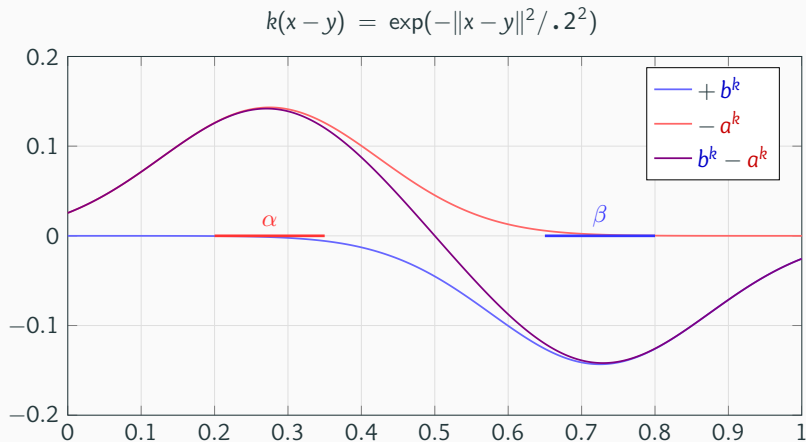


The registration flows along the gradient of $b^k - a^k$

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



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

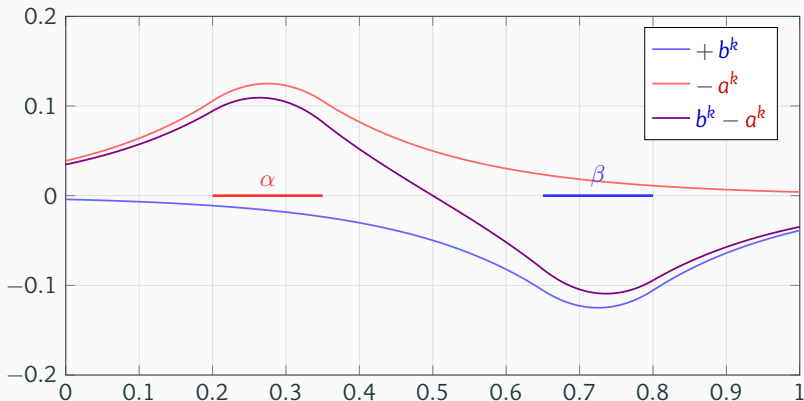


$$d_k(\alpha, \beta) = \langle \alpha - \beta \mid k \star (\alpha - \beta) \rangle$$

$$\frac{1}{2} \nabla_{x_i} d_k(\alpha, \beta) = \nabla [k \star (\alpha - \beta)](x_i) = \nabla b^k(x_i) - \nabla a^k(x_i)$$

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

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

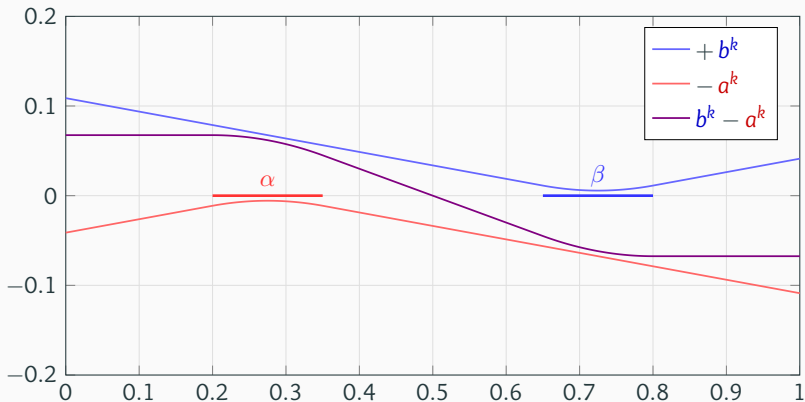


$$d_k(\alpha, \beta) = \langle \alpha - \beta \mid k \star (\alpha - \beta) \rangle$$

$$\frac{1}{2} \nabla_{x_i} d_k(\alpha, \beta) = \nabla [k \star (\alpha - \beta)](x_i) = \nabla b^k(x_i) - \nabla a^k(x_i)$$

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

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

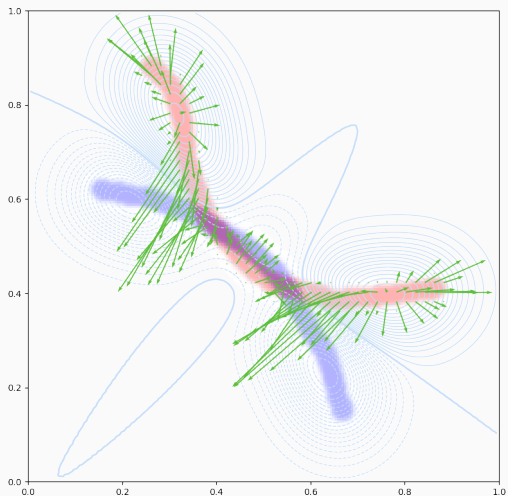


$$d_k(\alpha, \beta) = \langle \alpha - \beta \mid k \star (\alpha - \beta) \rangle$$

$$\frac{1}{2} \nabla_{x_i} d_k(\alpha, \beta) = \nabla [k \star (\alpha - \beta)](x_i) = \nabla b^k(x_i) - \nabla a^k(x_i)$$

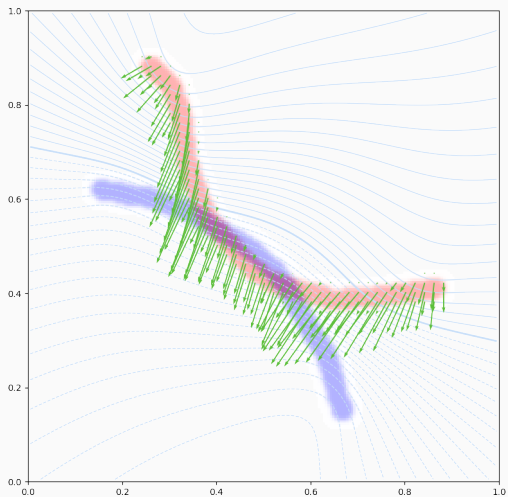
The Energy Distance is scale-invariant, robust

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



The Energy Distance is scale-invariant, robust

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



An idea from computer graphics:
Hausdorff distances

Can we go further?

$$\begin{matrix} & \beta_1 & \beta_2 & \dots & \beta_M \\ \alpha_1 & \left(\begin{array}{cccc} \|\mathbf{x}_1 - \mathbf{y}_1\| & \|\mathbf{x}_1 - \mathbf{y}_2\| & \dots & \|\mathbf{x}_1 - \mathbf{y}_M\| \\ \|\mathbf{x}_2 - \mathbf{y}_1\| & \|\mathbf{x}_2 - \mathbf{y}_2\| & \dots & \|\mathbf{x}_2 - \mathbf{y}_M\| \\ \vdots & \vdots & \ddots & \vdots \\ \|\mathbf{x}_N - \mathbf{y}_1\| & \|\mathbf{x}_N - \mathbf{y}_2\| & \dots & \|\mathbf{x}_N - \mathbf{y}_M\| \end{array} \right) \\ \alpha_2 & & & & \\ \vdots & & & & \\ \alpha_N & & & & \end{matrix}$$

Can we go further?

$$\begin{matrix} & \beta_1 & \beta_2 & \dots & \beta_M \\ \alpha_1 & \left(\begin{array}{cccc} \|\mathbf{x}_1 - \mathbf{y}_1\| & \|\mathbf{x}_1 - \mathbf{y}_2\| & \dots & \|\mathbf{x}_1 - \mathbf{y}_M\| \\ \|\mathbf{x}_2 - \mathbf{y}_1\| & \|\mathbf{x}_2 - \mathbf{y}_2\| & \dots & \|\mathbf{x}_2 - \mathbf{y}_M\| \\ \vdots & \vdots & \ddots & \vdots \\ \|\mathbf{x}_N - \mathbf{y}_1\| & \|\mathbf{x}_N - \mathbf{y}_2\| & \dots & \|\mathbf{x}_N - \mathbf{y}_M\| \end{array} \right) \\ \alpha_2 & & & & \\ \vdots & & & & \\ \alpha_N & & & & \end{matrix}$$

Can we go further?

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

$$\text{Energy Distance} \quad : \quad \sum_j \beta_j \|x_i - y_j\| = b^k(x_i)$$

Can we go further?

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

$$\text{Energy Distance} \quad : \quad \sum_j \beta_j \|x_i - y_j\| = b^k(x_i)$$

$$\text{Hausdorff Distance} \quad : \quad \min_j \|x_i - y_j\| = d(x_i, \text{supp}(\beta))$$

Can we go further?

$$\begin{array}{ccccccc} & \beta_1 & \beta_2 & \dots & \beta_M & & \\ \alpha_1 & \left(\begin{array}{cccc} \|x_1 - y_1\| & \|x_1 - y_2\| & \dots & \|x_1 - y_M\| \end{array} \right) & \rightarrow & \text{Smin}_{\varepsilon, y \sim \beta} \|x_1 - y\| \\ \alpha_2 & \left(\begin{array}{cccc} \|x_2 - y_1\| & \|x_2 - y_2\| & \dots & \|x_2 - y_M\| \end{array} \right) & \rightarrow & \text{Smin}_{\varepsilon, y \sim \beta} \|x_2 - y\| \\ \vdots & \left(\begin{array}{cccc} \vdots & \vdots & \ddots & \vdots \end{array} \right) & \rightarrow & \vdots \\ \alpha_N & \left(\begin{array}{cccc} \|x_N - y_1\| & \|x_N - y_2\| & \dots & \|x_N - y_M\| \end{array} \right) & \rightarrow & \text{Smin}_{\varepsilon, y \sim \beta} \|x_N - y\| \end{array}$$

$$\begin{array}{lcl} \text{Energy Distance} & : & \sum_j \beta_j \|x_i - y_j\| = b^k(x_i) \\ \varepsilon\text{-SoftMin} & : & \text{Smin}_{\varepsilon, y \sim \beta} \|x_i - y\| = b^\varepsilon(x_i) \\ \text{Hausdorff Distance} & : & \min_j \|x_i - y_j\| = d(x_i, \text{supp}(\beta)) \end{array}$$

The log-sum-exp trick

$$\log(e^{c_1} + e^{c_2}) = c_{\max} + \log(\underbrace{e^{c_1 - c_{\max}} + e^{c_2 - c_{\max}}}_{\in [1,2]})$$

The log-sum-exp trick

$$\log(e^{c_1} + e^{c_2}) = c_{\max} + \log(\underbrace{e^{c_1 - c_{\max}} + e^{c_2 - c_{\max}}}_{\in [1,2]})$$

Building on this, we define

$$b^\varepsilon(\mathbf{x}) = \text{Smin}_{\varepsilon, y \sim \beta} \|\mathbf{x} - \mathbf{y}\|$$

The log-sum-exp trick

$$\log(e^{c_1} + e^{c_2}) = c_{\max} + \underbrace{\log(e^{c_1 - c_{\max}} + e^{c_2 - c_{\max}})}_{\in [1,2]}$$

Building on this, we define

$$\begin{aligned} b^\varepsilon(\mathbf{x}) &= \text{Smin}_{\varepsilon, \mathbf{y} \sim \beta} \|\mathbf{x} - \mathbf{y}\| \\ &= -\varepsilon \log \sum_{j=1}^M \exp\left(\log(\beta_j) - \frac{1}{\varepsilon} \|\mathbf{x} - \mathbf{y}_j\|\right) \end{aligned}$$

The log-sum-exp trick

$$\log(e^{c_1} + e^{c_2}) = c_{\max} + \underbrace{\log(e^{c_1 - c_{\max}} + e^{c_2 - c_{\max}})}_{\in [1,2]}$$

Building on this, we define

$$\begin{aligned} b^\varepsilon(\mathbf{x}) &= \text{Smin}_{\varepsilon, y \sim \beta} \|\mathbf{x} - \mathbf{y}\| \\ &= -\varepsilon \log \sum_{j=1}^M \exp\left(\log(\beta_j) - \frac{1}{\varepsilon} \|\mathbf{x} - \mathbf{y}_j\|\right) \\ &\xrightarrow{\varepsilon \rightarrow +\infty} \sum_{j=1}^M \beta_j \|\mathbf{x} - \mathbf{y}_j\| \end{aligned}$$

The log-sum-exp trick

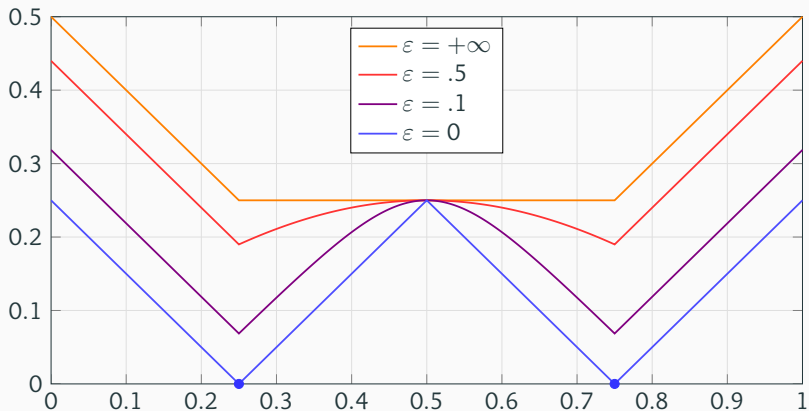
$$\log(e^{c_1} + e^{c_2}) = c_{\max} + \underbrace{\log(e^{c_1 - c_{\max}} + e^{c_2 - c_{\max}})}_{\in [1,2]}$$

Building on this, we define

$$\begin{aligned} b^\varepsilon(\mathbf{x}) &= \text{Smin}_{\varepsilon, \mathbf{y} \sim \beta} \|\mathbf{x} - \mathbf{y}\| \\ &= -\varepsilon \log \sum_{j=1}^M \exp\left(\log(\beta_j) - \frac{1}{\varepsilon} \|\mathbf{x} - \mathbf{y}_j\|\right) \\ &\xrightarrow{\varepsilon \rightarrow +\infty} \sum_{j=1}^M \beta_j \|\mathbf{x} - \mathbf{y}_j\| \\ &\xrightarrow{\varepsilon \rightarrow 0} \min_j \|\mathbf{x} - \mathbf{y}_j\| \end{aligned}$$

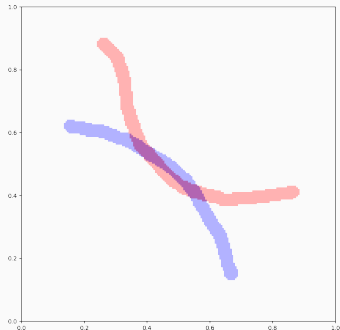
$Smin_{\epsilon}$ interpolates between a sum and a minimum

$x \mapsto Smin_{\epsilon, y \sim \beta} |x - y|$, with $\beta = \frac{1}{2}\delta_{.25} + \frac{1}{2}\delta_{.75}$



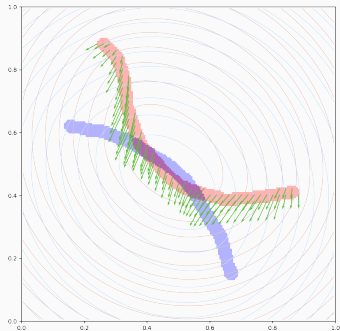
The ε -SoftMin fidelity interpolates between ED and Hausdorff

$$d_{\varepsilon\text{-SoftMin}}(\alpha, \beta) = \langle \alpha - \beta, b^\varepsilon - a^\varepsilon \rangle$$



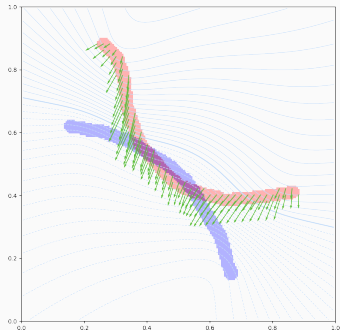
The ε -SoftMin fidelity interpolates between ED and Hausdorff

$$d_{\varepsilon\text{-SoftMin}}(\alpha, \beta) = \langle \alpha - \beta, b^\varepsilon - a^\varepsilon \rangle$$
$$\xrightarrow{\varepsilon \rightarrow +\infty} d_{\text{ED}}(\alpha, \beta) = \|\alpha - \beta\|_{-|\cdot|}^2$$



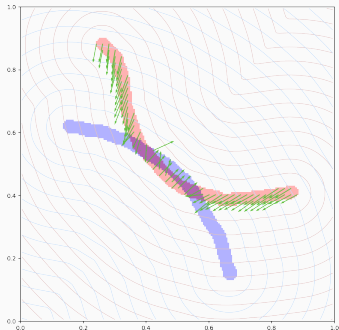
The ε -SoftMin fidelity interpolates between ED and Hausdorff

$$d_{\varepsilon\text{-SoftMin}}(\alpha, \beta) = \langle \alpha - \beta, b^\varepsilon - a^\varepsilon \rangle$$
$$\xrightarrow{\varepsilon \rightarrow +\infty} d_{\text{ED}}(\alpha, \beta) = \|\alpha - \beta\|_{-|\cdot|}$$



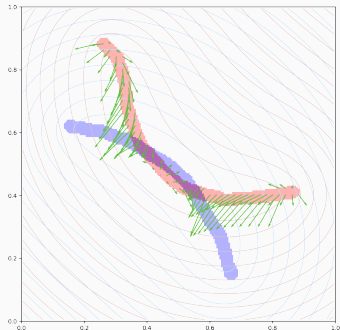
The ε -SoftMin fidelity interpolates between ED and Hausdorff

$$\begin{aligned}d_{\varepsilon\text{-SoftMin}}(\alpha, \beta) &= \langle \alpha - \beta, b^\varepsilon - a^\varepsilon \rangle \\ \xrightarrow{\varepsilon \rightarrow +\infty} d_{\text{ED}}(\alpha, \beta) &= \|\alpha - \beta\|_{-|\cdot|}^2 \\ \xrightarrow{\varepsilon \rightarrow 0} & \sum_i \alpha_i \min_{y \sim \beta} \|x_i - y\| + \sum_j \beta_j \min_{x \sim \alpha} \|x - y_j\|\end{aligned}$$



The ε -SoftMin fidelity interpolates between ED and Hausdorff

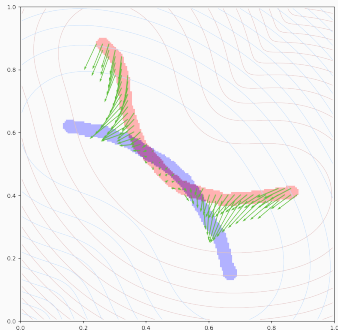
$$\begin{aligned}d_{\varepsilon\text{-SoftMin}}(\alpha, \beta) &= \langle \alpha - \beta, b^\varepsilon - a^\varepsilon \rangle \\ \xrightarrow{\varepsilon \rightarrow +\infty} d_{\text{ED}}(\alpha, \beta) &= \|\alpha - \beta\|_{-|\cdot|}^2 \\ \xrightarrow{\varepsilon \rightarrow 0} & \sum_i \alpha_i \min_{y \sim \beta} \|x_i - y\| + \sum_j \beta_j \min_{x \sim \alpha} \|x - y_j\|\end{aligned}$$



The ε -SoftMin fidelity interpolates between ED and Hausdorff

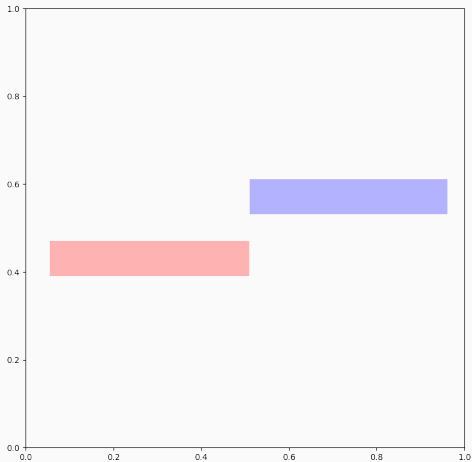
$$\begin{aligned}d_{\varepsilon\text{-SoftMin}}(\alpha, \beta) &= \langle \alpha - \beta, b^\varepsilon - a^\varepsilon \rangle \\ \xrightarrow{\varepsilon \rightarrow +\infty} d_{\text{ED}}(\alpha, \beta) &= \|\alpha - \beta\|_{-|\cdot|}^2 \\ \xrightarrow{\varepsilon \rightarrow 0} & \sum_i \alpha_i \min_{y \sim \beta} \|x_i - y\| + \sum_j \beta_j \min_{x \sim \alpha} \|x - y_j\|\end{aligned}$$

You can also use it with $C(x, y) = \|x - y\|^2$, etc.



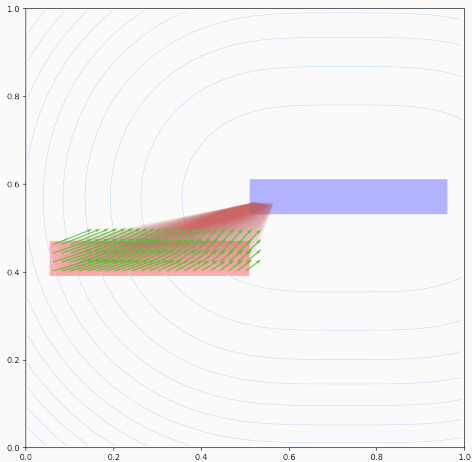
Shape registration isn't *always* about naive projections...

$$d_{\varepsilon\text{-SoftMin}}(\alpha, \beta) = \langle \alpha, b^{\varepsilon} - a^{\varepsilon} \rangle + \langle \beta, a^{\varepsilon} - b^{\varepsilon} \rangle$$



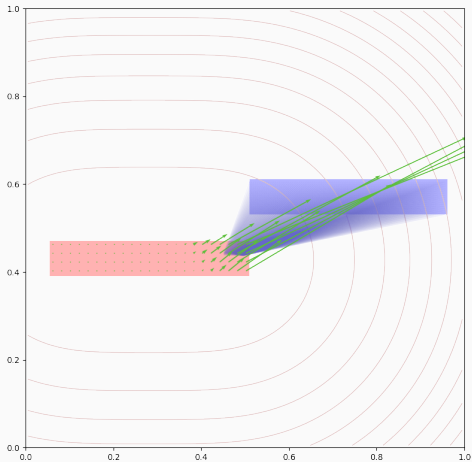
Shape registration isn't *always* about naive projections...

$$d_{\varepsilon\text{-SoftMin}}(\alpha, \beta) = \langle \alpha, b^{\varepsilon} - a^{\varepsilon} \rangle + \langle \beta, a^{\varepsilon} - b^{\varepsilon} \rangle$$



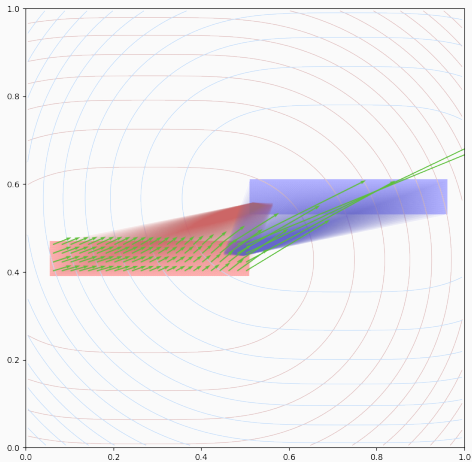
Shape registration isn't *always* about naive projections...

$$d_{\varepsilon\text{-SoftMin}}(\alpha, \beta) = \langle \alpha, b^{\varepsilon} - a^{\varepsilon} \rangle + \langle \beta, a^{\varepsilon} - b^{\varepsilon} \rangle$$



Shape registration isn't *always* about naive projections...

$$d_{\varepsilon\text{-SoftMin}}(\alpha, \beta) = \langle \alpha, b^{\varepsilon} - a^{\varepsilon} \rangle + \langle \beta, a^{\varepsilon} - b^{\varepsilon} \rangle$$



An idea from optimal transport theory:
Sinkhorn and Wasserstein divergences

Computational Optimal Transport

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

Enforce a **mass repartition** constraint through alternating projections onto α and β .

Computational Optimal Transport

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

Enforce a **mass repartition** constraint through alternating projections onto α and β .

Baseline algorithm:

Use this **Sinkhorn** loop to compute $a^{\alpha \rightarrow \beta}$ and $b^{\beta \rightarrow \alpha}$.

Define $W_\varepsilon(\alpha, \beta) = \langle \alpha, b^{\beta \rightarrow \alpha} \rangle + \langle \beta, a^{\alpha \rightarrow \beta} \rangle$

Computational Optimal Transport

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

Enforce a **mass repartition** constraint through alternating projections onto α and β .

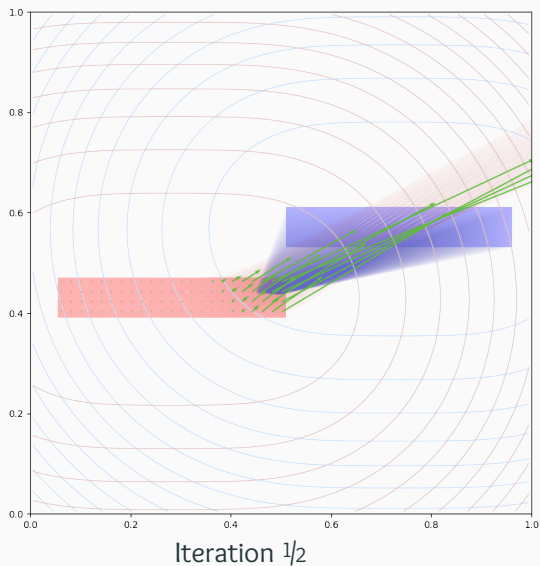
Baseline algorithm:

Use this **Sinkhorn** loop to compute $a^{\alpha \rightarrow \beta}$ and $b^{\beta \rightarrow \alpha}$.

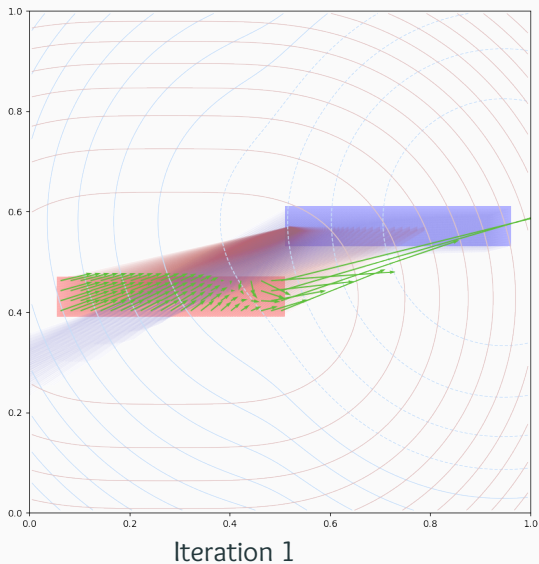
Define $W_\varepsilon(\alpha, \beta) = \langle \alpha, b^{\beta \rightarrow \alpha} \rangle + \langle \beta, a^{\alpha \rightarrow \beta} \rangle$

The core operation is still Smin_ε , for some $\varepsilon > 0$.

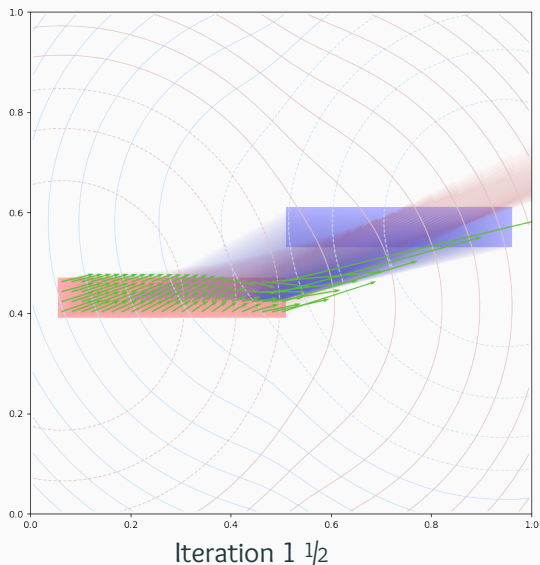
The Sinkhorn algorithm, in practice; $C(x,y) = \|x - y\|^2$, $\sqrt{\epsilon} = .1$



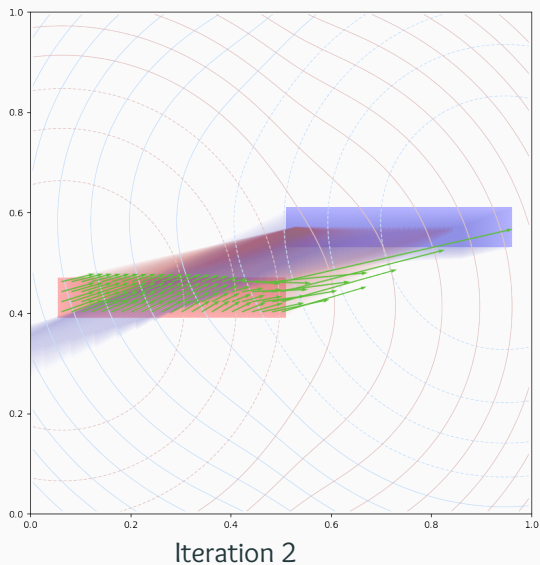
The Sinkhorn algorithm, in practice; $C(x,y) = \|x - y\|^2$, $\sqrt{\epsilon} = .1$



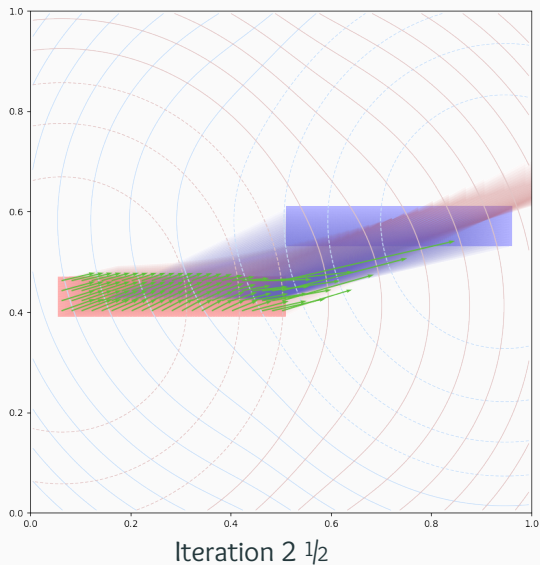
The Sinkhorn algorithm, in practice; $C(x,y) = \|x - y\|^2$, $\sqrt{\epsilon} = .1$



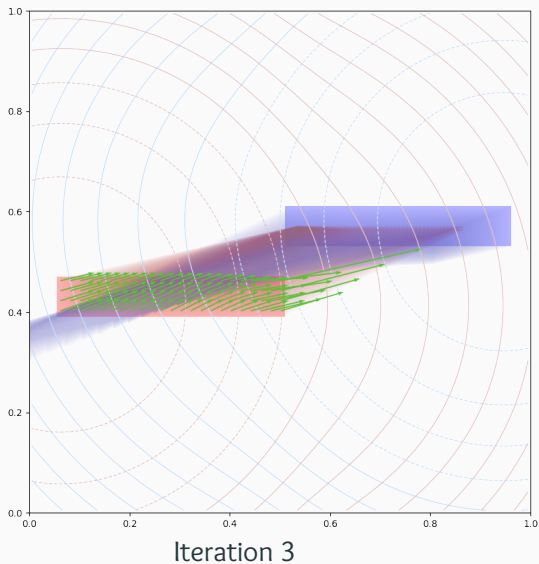
The Sinkhorn algorithm, in practice; $C(x,y) = \|x - y\|^2$, $\sqrt{\epsilon} = .1$



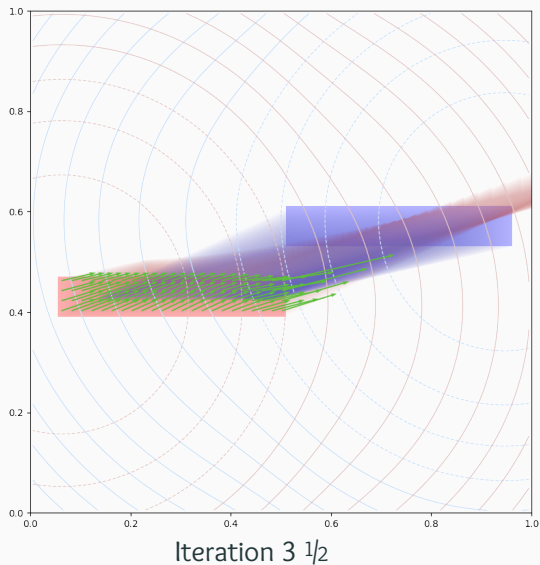
The Sinkhorn algorithm, in practice; $C(x,y) = \|x - y\|^2$, $\sqrt{\epsilon} = .1$



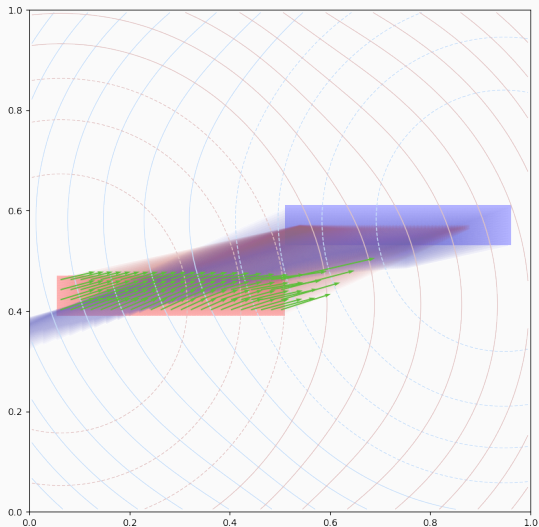
The Sinkhorn algorithm, in practice; $C(x,y) = \|x - y\|^2$, $\sqrt{\epsilon} = .1$



The Sinkhorn algorithm, in practice; $C(x,y) = \|x - y\|^2$, $\sqrt{\epsilon} = .1$

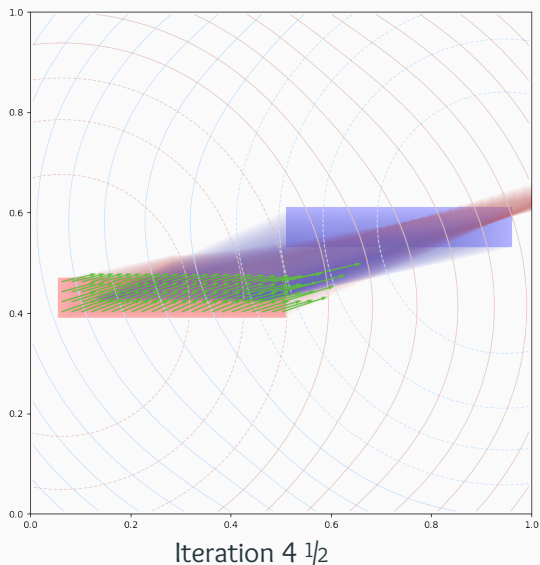


The Sinkhorn algorithm, in practice; $C(x,y) = \|x - y\|^2$, $\sqrt{\epsilon} = .1$

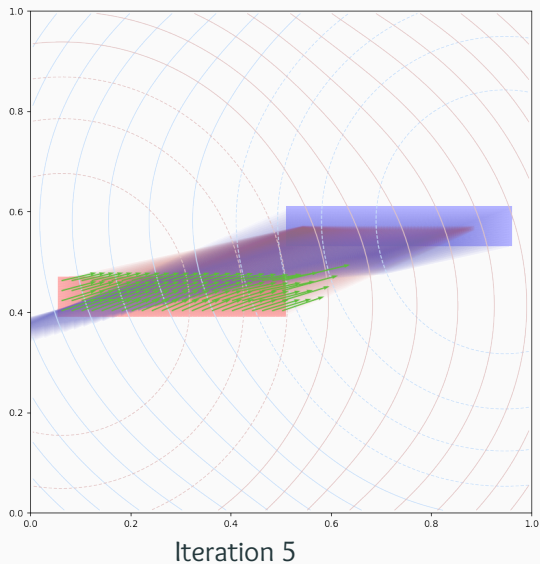


Iteration 4

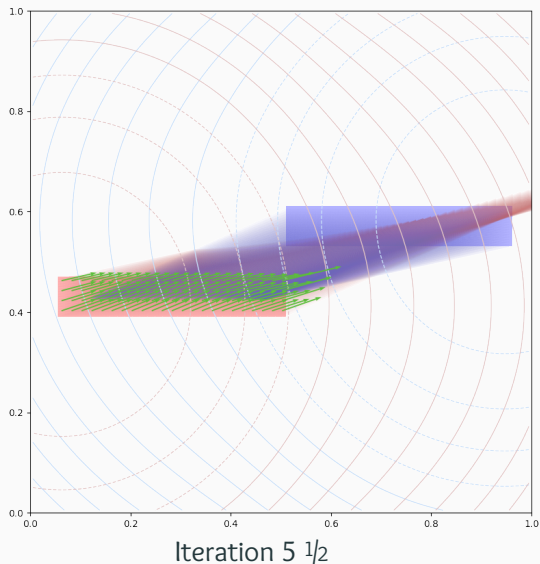
The Sinkhorn algorithm, in practice; $C(x,y) = \|x - y\|^2$, $\sqrt{\epsilon} = .1$



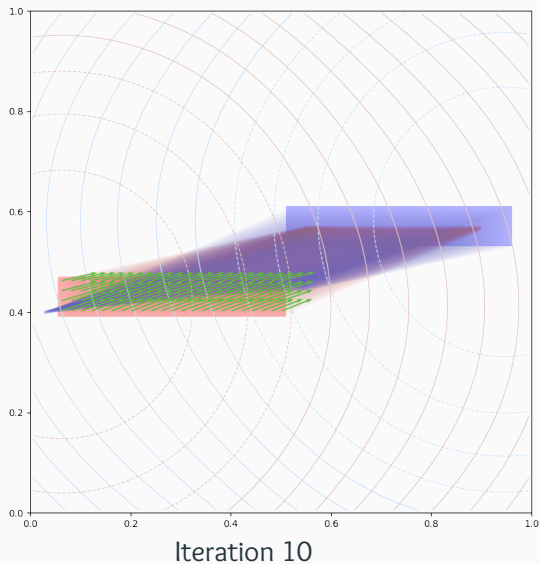
The Sinkhorn algorithm, in practice; $C(x,y) = \|x - y\|^2$, $\sqrt{\epsilon} = .1$



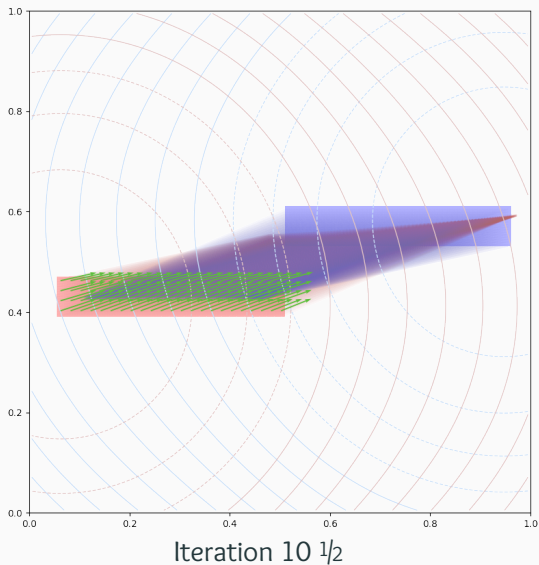
The Sinkhorn algorithm, in practice; $C(x,y) = \|x - y\|^2$, $\sqrt{\epsilon} = .1$



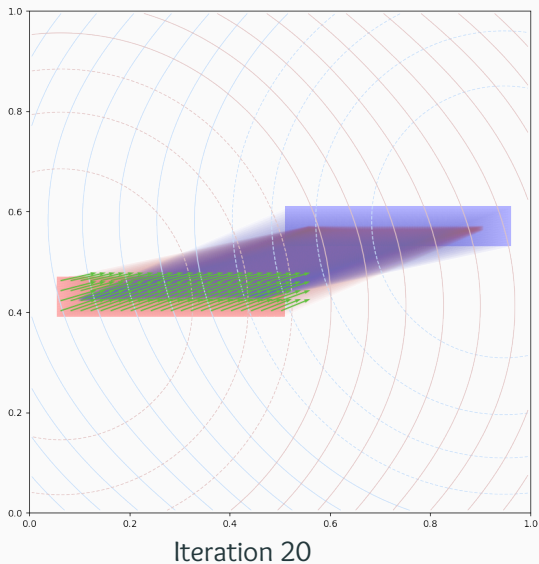
The Sinkhorn algorithm, in practice; $C(x,y) = \|x - y\|^2$, $\sqrt{\epsilon} = .1$



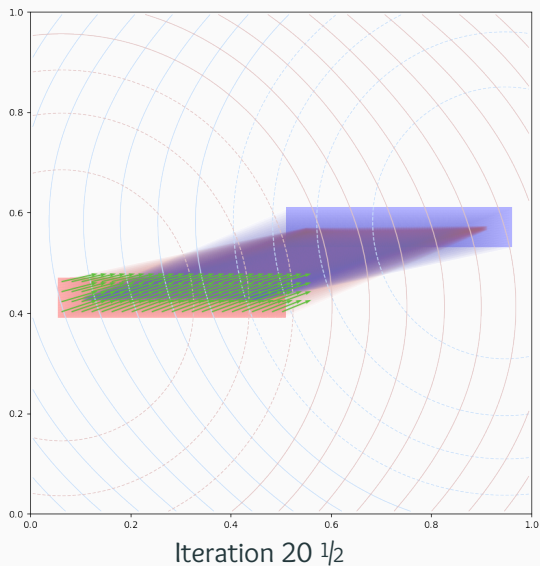
The Sinkhorn algorithm, in practice; $C(x,y) = \|x - y\|^2$, $\sqrt{\epsilon} = .1$



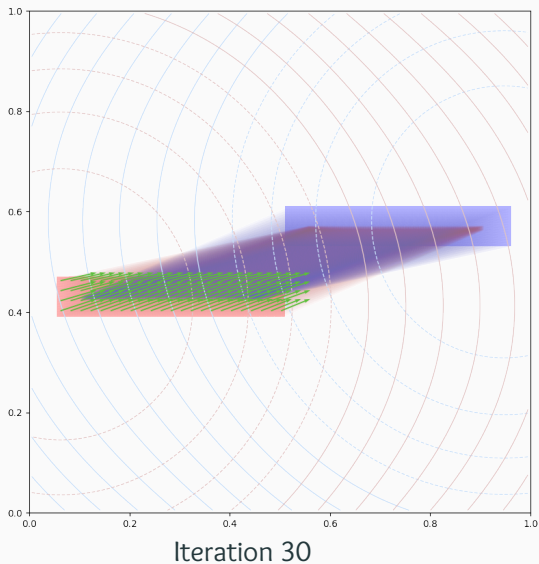
The Sinkhorn algorithm, in practice; $C(x,y) = \|x - y\|^2$, $\sqrt{\epsilon} = .1$



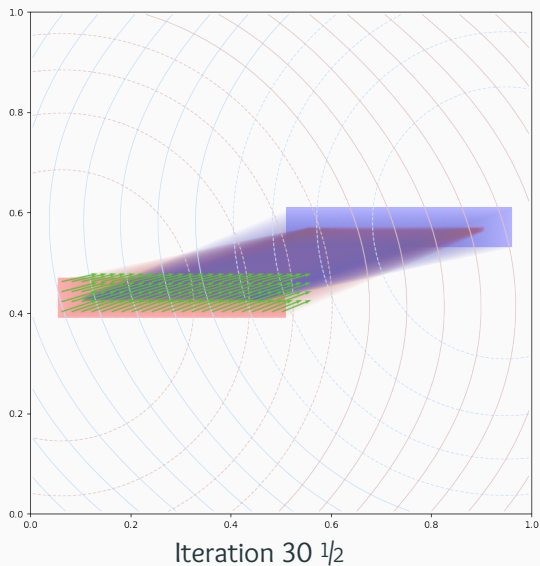
The Sinkhorn algorithm, in practice; $C(x,y) = \|x - y\|^2$, $\sqrt{\epsilon} = .1$



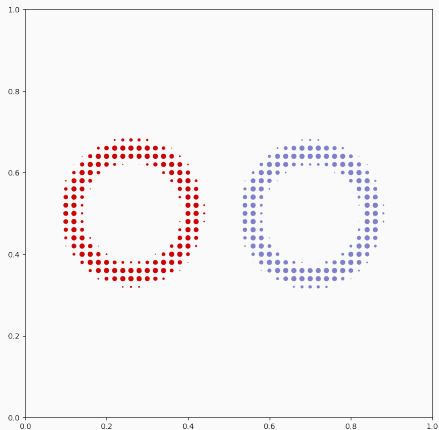
The Sinkhorn algorithm, in practice; $C(x,y) = \|x - y\|^2$, $\sqrt{\epsilon} = .1$



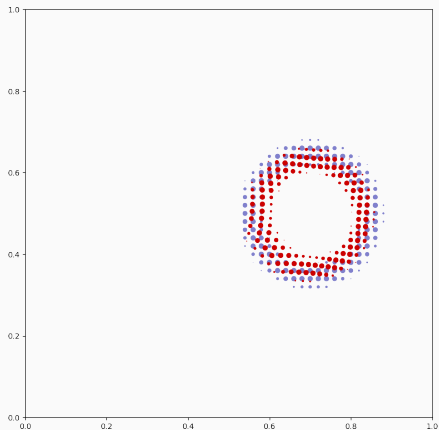
The Sinkhorn algorithm, in practice; $C(x,y) = \|x - y\|^2$, $\sqrt{\epsilon} = .1$



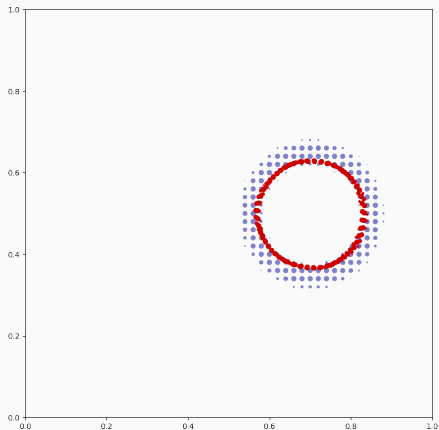
Registrating circles; $C(x,y) = \|x - y\|^2$, $\sqrt{\varepsilon} = 0.1$



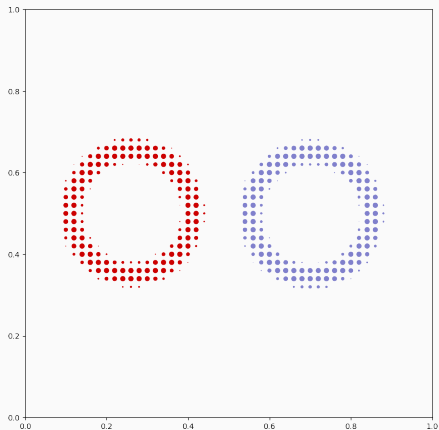
Registrating circles; $C(x,y) = \|x - y\|^2$, $\sqrt{\varepsilon} = 0.1$



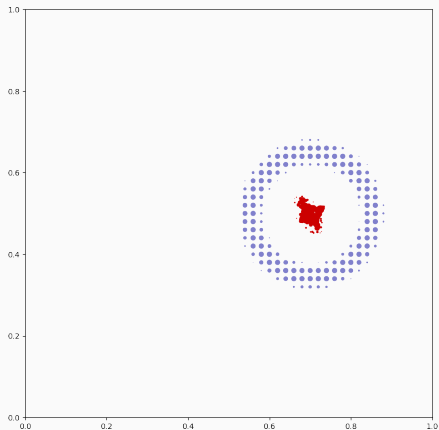
Registrating circles; $C(x,y) = \|x - y\|^2$, $\sqrt{\varepsilon} = 0.1$



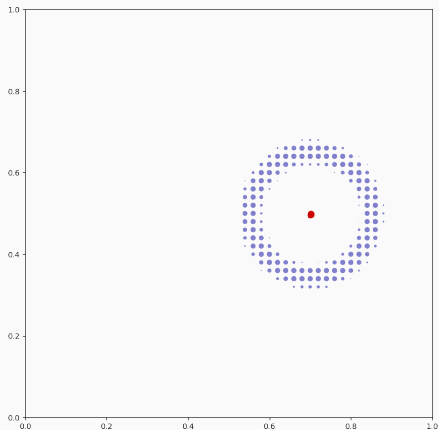
Registrating circles; $C(x,y) = \|x - y\|^2$, $\sqrt{\varepsilon} = 0.2$



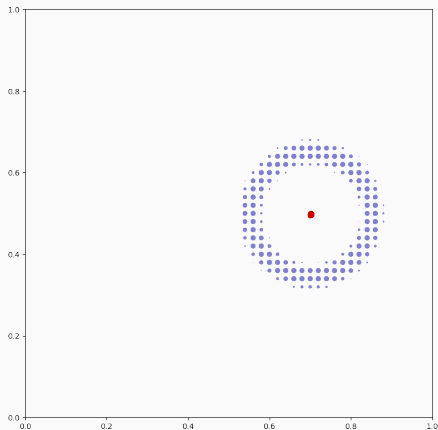
Registrating circles; $C(x,y) = \|x - y\|^2$, $\sqrt{\varepsilon} = 0.2$



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



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



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

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

Solution: Use an unbiased divergence [Genevay et al., 2018]

$$d_{\varepsilon\text{-Sinkhorn}}(\alpha, \beta) = w_{\varepsilon}(\alpha, \beta) - \frac{1}{2}w_{\varepsilon}(\alpha, \alpha) - \frac{1}{2}w_{\varepsilon}(\beta, \beta).$$

In our paper: theoretical guarantees

Solution: Use an unbiased divergence [Genevay et al., 2018]

$$d_{\varepsilon\text{-Sinkhorn}}(\alpha, \beta) = W_{\varepsilon}(\alpha, \beta) - \frac{1}{2}W_{\varepsilon}(\alpha, \alpha) - \frac{1}{2}W_{\varepsilon}(\beta, \beta).$$

Theorem (Positivity ; F., Vialard)

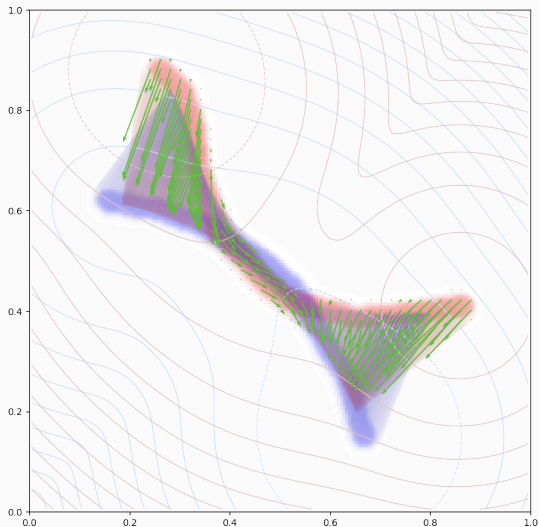
We define $d_{\varepsilon\text{-Hausdorff}}(\alpha, \beta) \simeq d_{\varepsilon\text{-SoftMin}}(\alpha, \beta)$. Then, if

$$k_{\varepsilon}(\mathbf{x}, \mathbf{y}) = \exp\left(-\frac{1}{\varepsilon}C(\mathbf{x}, \mathbf{y})\right)$$

defines a positive kernel, we have

$$0 \leq d_{\varepsilon\text{-Hausdorff}}(\alpha, \beta) \leq d_{\varepsilon\text{-Sinkhorn}}(\alpha, \beta).$$

The ε -Sinkhorn divergence; $C(x,y) = \|x - y\|^2$, $\sqrt{\varepsilon} = .1$



A high-quality gradient.

Conclusion

How do I implement this on real shapes?

$$\begin{array}{cccc} & \beta_1 & \beta_2 & \dots & \beta_M \\ \alpha_1 & \left(\begin{array}{cccc} \|\mathbf{x}_1 - \mathbf{y}_1\| & \|\mathbf{x}_1 - \mathbf{y}_2\| & \cdots & \|\mathbf{x}_1 - \mathbf{y}_M\| \\ \|\mathbf{x}_2 - \mathbf{y}_1\| & \|\mathbf{x}_2 - \mathbf{y}_2\| & \cdots & \|\mathbf{x}_2 - \mathbf{y}_M\| \\ \vdots & \vdots & \ddots & \vdots \\ \|\mathbf{x}_N - \mathbf{y}_1\| & \|\mathbf{x}_N - \mathbf{y}_2\| & \cdots & \|\mathbf{x}_N - \mathbf{y}_M\| \end{array} \right) & \rightarrow & \text{Smin}_{\varepsilon, \mathbf{y} \sim \beta} \|\mathbf{x}_1 - \mathbf{y}\| \\ \alpha_2 & & & & \rightarrow & \text{Smin}_{\varepsilon, \mathbf{y} \sim \beta} \|\mathbf{x}_2 - \mathbf{y}\| \\ \vdots & & & & \vdots & \\ \alpha_N & & & & \rightarrow & \text{Smin}_{\varepsilon, \mathbf{y} \sim \beta} \|\mathbf{x}_N - \mathbf{y}\| \end{array}$$

How do I implement this on real shapes?

$$\begin{array}{cccc} & \beta_1 & \beta_2 & \dots & \beta_M \\ \alpha_1 & \left(\begin{array}{cccc} \|x_1 - y_1\| & \|x_1 - y_2\| & \dots & \|x_1 - y_M\| \\ \|x_2 - y_1\| & \|x_2 - y_2\| & \dots & \|x_2 - y_M\| \\ \vdots & \vdots & \ddots & \vdots \\ \|x_N - y_1\| & \|x_N - y_2\| & \dots & \|x_N - y_M\| \end{array} \right) & \rightarrow & \text{Smin}_{\epsilon, y \sim \beta} \|x_1 - y\| \\ \alpha_2 & & & & \rightarrow & \text{Smin}_{\epsilon, y \sim \beta} \|x_2 - y\| \\ \vdots & & & & \vdots & \\ \alpha_N & & & & \rightarrow & \text{Smin}_{\epsilon, y \sim \beta} \|x_N - y\| \end{array}$$

Huge 100,000-by-100,000 matrices just don't fit into GPU memories.

How do I implement this on real shapes?

$$\begin{array}{cccccc} & \beta_1 & \beta_2 & \dots & \beta_M & \\ \alpha_1 & \left(\begin{array}{cccc} \|x_1 - y_1\| & \|x_1 - y_2\| & \dots & \|x_1 - y_M\| \end{array} \right) & \rightarrow & \text{Smin}_{\epsilon, y \sim \beta} \|x_1 - y\| \\ \alpha_2 & \left(\begin{array}{cccc} \|x_2 - y_1\| & \|x_2 - y_2\| & \dots & \|x_2 - y_M\| \end{array} \right) & \rightarrow & \text{Smin}_{\epsilon, y \sim \beta} \|x_2 - y\| \\ \vdots & \left(\begin{array}{cccc} \vdots & \vdots & \ddots & \vdots \end{array} \right) & & \vdots \\ \alpha_N & \left(\begin{array}{cccc} \|x_N - y_1\| & \|x_N - y_2\| & \dots & \|x_N - y_M\| \end{array} \right) & \rightarrow & \text{Smin}_{\epsilon, y \sim \beta} \|x_N - y\| \end{array}$$

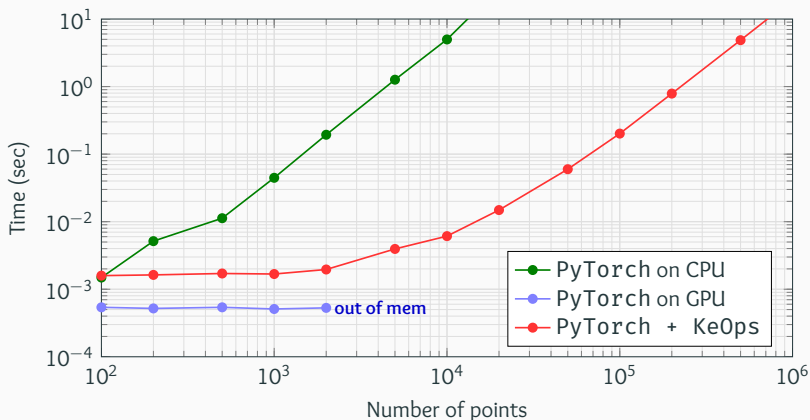
Huge 100,000-by-100,000 matrices just don't fit into GPU memories.

We need online map-reduce routines.

Kernel OPERATIONS, with autodiff, without memory overflows

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

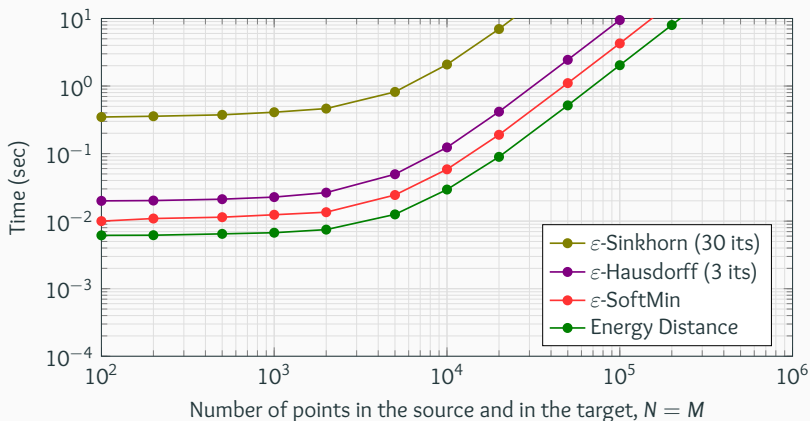
Time needed to compute an N-by-N Gaussian convolution in \mathbb{R}^3



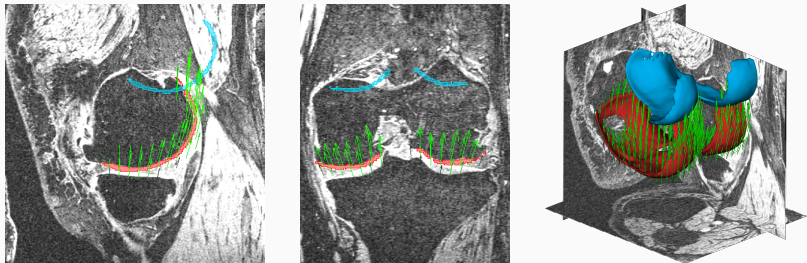
Kernel OPERATIONS, with autodiff, without memory overflows

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

Time needed to compute a fidelity and its gradient, using KeOps



On real data, from the OsteoArthritis Initiative



Gradient of the Energy Distance, computed in 0.5s on my laptop.
(52,319 and 34,966 voxels – out of a 192-192-160 volume)

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

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

Our contributions:

- Statistics \longleftrightarrow Computer Graphics \longleftrightarrow Optimal Transport

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

Our contributions:

- Statistics \longleftrightarrow Computer Graphics \longleftrightarrow Optimal Transport
- New ε -Hausdorff fidelity, metric entropy.

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

Our contributions:

- Statistics \longleftrightarrow Computer Graphics \longleftrightarrow Optimal Transport
- New ε -**Hausdorff** fidelity, metric entropy.
- First proof that the ε -Sinkhorn divergence (from ML) is **positive**.

Conclusion

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

Our contributions:

- Statistics \longleftrightarrow Computer Graphics \longleftrightarrow Optimal Transport
- New ε -**Hausdorff** fidelity, metric entropy.
- First proof that the ε -Sinkhorn divergence (from ML) is **positive**.
- **KeOps**: efficient online map-reduce routines

CUDA + Matlab, numpy, PyTorch

Our code is available:

`plmlab.math.cnrs.fr/jeanfeydy/shapes_toolbox`

Our code is available:

`plmlab.math.cnrs.fr/jeanfeydy/shapes_toolbox`

Our papers:

- *Global divergences between measures: from Hausdorff distance to Optimal Transport*, F., Trouvé, 2018 (uploaded yesterday on HAL)

Our code is available:

`plmlab.math.cnrs.fr/jeanfeydy/shapes_toolbox`

Our papers:

- *Global divergences between measures: from Hausdorff distance to Optimal Transport*, F., Trouvé, 2018 (uploaded yesterday on HAL)
- *Sinkhorn entropies and divergences*, F., Séjourné, Vialard, Amari, Trouvé, Peyré, 2018 (soon)

Our code is available:

`plmlab.math.cnrs.fr/jeanfeydy/shapes_toolbox`

Our papers:

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

Thank you for your attention.

Any questions ?



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.