New results on quantitative stability of optimal transport

Cyril Letrouit
Joint work with Quentin Mérigot

CNRS – Université Paris-Saclay
Laboratoire de mathématiques d'Orsay

February 17th, 2025

The problem

Theorem (Brenier)

Let $c(x, y) = \|x - y\|^2$ on $\mathbb{R}^d$ and assume that $\rho \in \mathcal{P}_2(\mathbb{R}^d)$ ***has a density****, and $\mu \in \mathcal{P}_2(\mathbb{R}^d)$. Then there exists a unique optimal transport $T : \mathbb{R}^d \to \mathbb{R}^d$. Moreover, $T = \nabla\phi$ with $\phi : \mathbb{R}^d \to \mathbb{R}$ convex.*

The problem

Theorem (Brenier)

*Let $c(x, y) = \|x - y\|^2$ on $\mathbb{R}^d$ and assume that $\rho \in \mathcal{P}_2(\mathbb{R}^d)$ **has a density**, and $\mu \in \mathcal{P}_2(\mathbb{R}^d)$. Then there exists a unique optimal transport $T : \mathbb{R}^d \to \mathbb{R}^d$. Moreover, $T = \nabla\phi$ with $\phi : \mathbb{R}^d \to \mathbb{R}$ convex.*

- $\rho \in \mathcal{P}_2(\mathbb{R}^d)$ is **fixed** and **has a density**.

The problem

Theorem (Brenier)

Let $c(x, y) = \|x - y\|^2$ on $\mathbb{R}^d$ and assume that $\rho \in \mathcal{P}_2(\mathbb{R}^d)$ ***has a density****, and $\mu \in \mathcal{P}_2(\mathbb{R}^d)$. Then there exists a unique optimal transport $T : \mathbb{R}^d \to \mathbb{R}^d$. Moreover, $T = \nabla\phi$ with $\phi : \mathbb{R}^d \to \mathbb{R}$ convex.*

- $\rho \in \mathcal{P}_2(\mathbb{R}^d)$ is **fixed** and **has a density**.
- The **Brenier map** from ρ to μ is written T_μ.
- The **Brenier potential** from ρ to μ is the unique $\phi_\mu \in L^2(\rho)$ such that $T_\mu = \nabla\phi_\mu$ and $\int_{\mathcal{X}} \phi_\mu d\rho = 0$ ($\mathcal{X}$ is the support of ρ, assumed connected).

The problem

Theorem (Brenier)

*Let $c(x, y) = \|x - y\|^2$ on $\mathbb{R}^d$ and assume that $\rho \in \mathcal{P}_2(\mathbb{R}^d)$ **has a density**, and $\mu \in \mathcal{P}_2(\mathbb{R}^d)$. Then there exists a unique optimal transport $T : \mathbb{R}^d \to \mathbb{R}^d$. Moreover, $T = \nabla\phi$ with $\phi : \mathbb{R}^d \to \mathbb{R}$ convex.*

- $\rho \in \mathcal{P}_2(\mathbb{R}^d)$ is **fixed** and **has a density**.
- The **Brenier map** from ρ to μ is written T_μ.
- The **Brenier potential** from ρ to μ is the unique $\phi_\mu \in L^2(\rho)$ such that $T_\mu = \nabla\phi_\mu$ and $\int_{\mathcal{X}} \phi_\mu d\rho = 0$ ($\mathcal{X}$ is the support of ρ, assumed connected).
- The map $\mu \mapsto T_\mu$ is continuous: if $(\mu_n)_n$ converges to μ in $(\mathcal{P}_2(\mathbb{R}^d), W_2)$, then T_{μ_n} converges to T_μ in $L^2(\rho, \mathbb{R}^d)$. But non-quantitative.

The problem

Theorem (Brenier)

Let $c(x, y) = \|x - y\|^2$ on $\mathbb{R}^d$ and assume that $\rho \in \mathcal{P}_2(\mathbb{R}^d)$ ***has a density****, and $\mu \in \mathcal{P}_2(\mathbb{R}^d)$. Then there exists a unique optimal transport $T : \mathbb{R}^d \to \mathbb{R}^d$. Moreover, $T = \nabla\phi$ with $\phi : \mathbb{R}^d \to \mathbb{R}$ convex.*

- $\rho \in \mathcal{P}_2(\mathbb{R}^d)$ is **fixed** and **has a density**.
- The **Brenier map** from ρ to μ is written T_μ.
- The **Brenier potential** from ρ to μ is the unique $\phi_\mu \in L^2(\rho)$ such that $T_\mu = \nabla\phi_\mu$ and $\int_{\mathcal{X}} \phi_\mu d\rho = 0$ ($\mathcal{X}$ is the support of ρ, assumed connected).
- The map $\mu \mapsto T_\mu$ is continuous: if $(\mu_n)_n$ converges to μ in $(\mathcal{P}_2(\mathbb{R}^d), W_2)$, then T_{μ_n} converges to T_μ in $L^2(\rho, \mathbb{R}^d)$. But non-quantitative.
- Our question: "quantify this continuity". If μ and ν are close, how close are T_μ and T_ν?
- We look for an inequality $\|T_\mu - T_\nu\|_{L^2(\rho)} \leqslant CW_2(\mu, \nu)^\alpha$ for some C, α depending on ρ but not on μ, ν. And similar inequality for $\|\phi_\mu - \phi_\nu\|_{L^2(\rho)}$.

The problem

Theorem (Brenier)

*Let $c(x, y) = \|x - y\|^2$ on $\mathbb{R}^d$ and assume that $\rho \in \mathcal{P}_2(\mathbb{R}^d)$ **has a density**, and $\mu \in \mathcal{P}_2(\mathbb{R}^d)$. Then there exists a unique optimal transport $T : \mathbb{R}^d \to \mathbb{R}^d$. Moreover, $T = \nabla\phi$ with $\phi : \mathbb{R}^d \to \mathbb{R}$ convex.*

- $\rho \in \mathcal{P}_2(\mathbb{R}^d)$ is **fixed** and **has a density**.
- The **Brenier map** from ρ to μ is written T_μ.
- The **Brenier potential** from ρ to μ is the unique $\phi_\mu \in L^2(\rho)$ such that $T_\mu = \nabla\phi_\mu$ and $\int_{\mathcal{X}} \phi_\mu d\rho = 0$ ($\mathcal{X}$ is the support of ρ, assumed connected).
- The map $\mu \mapsto T_\mu$ is continuous: if $(\mu_n)_n$ converges to μ in $(\mathcal{P}_2(\mathbb{R}^d), W_2)$, then T_{μ_n} converges to T_μ in $L^2(\rho, \mathbb{R}^d)$. But non-quantitative.
- Our question: "quantify this continuity". If μ and ν are close, how close are T_μ and T_ν?
- We look for an inequality $\|T_\mu - T_\nu\|_{L^2(\rho)} \leqslant CW_2(\mu, \nu)^\alpha$ for some C, α depending on ρ but not on μ, ν. And similar inequality for $\|\phi_\mu - \phi_\nu\|_{L^2(\rho)}$.
- Remark: a reverse inequality always holds

$$\forall \mu, \nu \in \mathcal{P}_2(\mathbb{R}^d), \quad \|T_\mu - T_\nu\|_{L^2(\rho, \mathbb{R}^d)} \geqslant W_2(\mu, \nu).$$

Linearized optimal transport

- Can we replace the Wasserstein distance by the distance

$$W_{2,\rho}(\mu,\nu) = \| T_\mu - T_\nu \|_{L^2(\rho,\mathbb{R}^d)}, \quad ?$$

How much do $(\mathcal{P}_2(\mathbb{R}^d), W_2)$ and the Hilbert space $L^2(\rho,\mathbb{R}^d)$ look like?

Linearized optimal transport

- Can we replace the Wasserstein distance by the distance

$$W_{2,\rho}(\mu,\nu) = \|T_\mu - T_\nu\|_{L^2(\rho,\mathbb{R}^d)}, \quad ?$$

How much do $(\mathcal{P}_2(\mathbb{R}^d), W_2)$ and the Hilbert space $L^2(\rho,\mathbb{R}^d)$ look like?
Can we "replace" computations with μ by computations with T_μ?

Linearized optimal transport

- Can we replace the Wasserstein distance by the distance

$$W_{2,\rho}(\mu,\nu) = \| T_\mu - T_\nu \|_{L^2(\rho,\mathbb{R}^d)}, \quad ?$$

  How much do $(\mathcal{P}_2(\mathbb{R}^d), W_2)$ and the Hilbert space $L^2(\rho, \mathbb{R}^d)$ look like? Can we "replace" computations with μ by computations with T_μ?

- This is **Linearized optimal transport** (LOT).
- Goal: use classical Hilbertian statistical tools on families of probability measures while keeping some features of the Wasserstein geometry.
- Interest in **numerical analysis and in statistics:** $\mu \in \mathcal{P}_2(\mathbb{R}^d)$ is often approximated by a sequence of finitely supported measures $(\mu_n)_n$.
- Applications to image processing: pattern recognition, detection of differences in images, generative modelling of images, improving resolution of images, **computation of Wasserstein barycenters**...

Linearized optimal transport

- Can we replace the Wasserstein distance by the distance

$$W_{2,\rho}(\mu,\nu) = \| T_\mu - T_\nu \|_{L^2(\rho,\mathbb{R}^d)}, \quad ?$$

  How much do $(\mathcal{P}_2(\mathbb{R}^d), W_2)$ and the Hilbert space $L^2(\rho,\mathbb{R}^d)$ look like?
  Can we "replace" computations with μ by computations with T_μ?
- This is **Linearized optimal transport** (LOT).
- Goal: use classical Hilbertian statistical tools on families of probability measures while keeping some features of the Wasserstein geometry.
- Interest in **numerical analysis and in statistics:** $\mu \in \mathcal{P}_2(\mathbb{R}^d)$ is often approximated by a sequence of finitely supported measures $(\mu_n)_n$.
- Applications to image processing: pattern recognition, detection of differences in images, generative modelling of images, improving resolution of images, **computation of Wasserstein barycenters**...
- But good practical behavior of LOT not justified mathematically.

Our results

Quantitative stability **holds** if ρ is

Our results

Quantitative stability **holds** if ρ is

- e^{-U-F} with $D^2 U \geqslant \kappa \mathrm{Id}$, $\kappa > 0$ and $F \in L^\infty(\mathbb{R}^d)$;

Our results

Quantitative stability **holds** if ρ is

- e^{-U-F} with $D^2U \geqslant \kappa \mathrm{Id}$, $\kappa > 0$ and $F \in L^\infty(\mathbb{R}^d)$;
- $\rho(x) = (1+|x|)^{-\beta}$ in $\mathbb{R}^d$, for $\beta > d+2$.

Our results

Quantitative stability **holds** if ρ is

- e^{-U-F} with $D^2U \geqslant \kappa \mathrm{Id}$, $\kappa > 0$ and $F \in L^\infty(\mathbb{R}^d)$;
- $\rho(x) = (1+|x|)^{-\beta}$ in $\mathbb{R}^d$, for $\beta > d+2$.
- bounded above and below on John (e.g. bounded Lipschitz) domain, or finite union thereof. Previously: bounded convex [Delalande-Mérigot].

Our results

Quantitative stability **holds** if ρ is

- e^{-U-F} with $D^2U \geqslant \kappa \mathrm{Id}$, $\kappa > 0$ and $F \in L^\infty(\mathbb{R}^d)$;
- $\rho(x) = (1+|x|)^{-\beta}$ in $\mathbb{R}^d$, for $\beta > d+2$.
- bounded above and below on John (e.g. bounded Lipschitz) domain, or finite union thereof. Previously: bounded convex [Delalande-Mérigot].
- ρ is the spherical uniform distribution $\rho(x) = \frac{c_d}{|x|^{d-1}}$ on the unit ball of $\mathbb{R}^d$;
- ρ has compact support and blows-up polynomially at the boundary.

Our results

Quantitative stability **holds** if ρ is

- e^{-U-F} with $D^2 U \geqslant \kappa \mathrm{Id}$, $\kappa > 0$ and $F \in L^\infty(\mathbb{R}^d)$;
- $\rho(x) = (1 + |x|)^{-\beta}$ in $\mathbb{R}^d$, for $\beta > d + 2$.
- bounded above and below on John (e.g. bounded Lipschitz) domain, or finite union thereof. Previously: bounded convex [Delalande-Mérigot].
- ρ is the spherical uniform distribution $\rho(x) = \frac{c_d}{|x|^{d-1}}$ on the unit ball of $\mathbb{R}^d$;
- ρ has compact support and blows-up polynomially at the boundary.

Exponents of stability are **sharp for potentials** in the first two cases.

Our results

Quantitative stability **holds** if ρ is

- e^{-U-F} with $D^2U \geqslant \kappa\mathrm{Id}$, $\kappa > 0$ and $F \in L^\infty(\mathbb{R}^d)$;
- $\rho(x) = (1+|x|)^{-\beta}$ in $\mathbb{R}^d$, for $\beta > d+2$.
- bounded above and below on John (e.g. bounded Lipschitz) domain, or finite union thereof. Previously: bounded convex [Delalande-Mérigot].
- ρ is the spherical uniform distribution $\rho(x) = \frac{c_d}{|x|^{d-1}}$ on the unit ball of $\mathbb{R}^d$;
- ρ has compact support and blows-up polynomially at the boundary.

Exponents of stability are **sharp for potentials** in the first two cases.
Quantitative stability **does not hold** for any exponents, in some non-John domains, e.g. for the uniform distribution on “room-and-passage" domains:

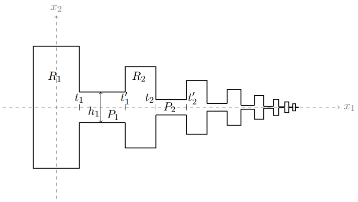

Our results

Quantitative stability **holds** if ρ is

- e^{-U-F} with $D^2 U \geqslant \kappa \mathrm{Id}$, $\kappa > 0$ and $F \in L^\infty(\mathbb{R}^d)$;
- $\rho(x) = (1 + |x|)^{-\beta}$ in $\mathbb{R}^d$, for $\beta > d + 2$.
- bounded above and below on John (e.g. bounded Lipschitz) domain, or finite union thereof. Previously: bounded convex [Delalande-Mérigot].
- ρ is the spherical uniform distribution $\rho(x) = \frac{c_d}{|x|^{d-1}}$ on the unit ball of $\mathbb{R}^d$;
- ρ has compact support and blows-up polynomially at the boundary.

Exponents of stability are **sharp for potentials** in the first two cases.
Quantitative stability **does not hold** for any exponents, in some non-John domains, e.g. for the uniform distribution on “room-and-passage" domains:

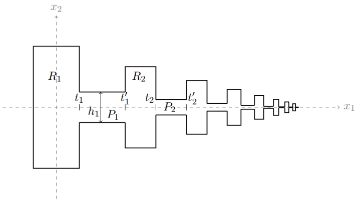

Proof technique: gluing (spectral). Nearly finished: ρ on Riemannian manifolds.