Optimal Transport for Data Analysis

Bernhard Schmitzer

2017-06-13

1 Introduction

1.1 Reminders on Measure Theory

Reference: Ambrosio, Fusco, Pallara: Functions of Bounded Variation and Free Discontinuity Problems, Chapters 1 & 2, [Ambrosio et al., 2000].

Definition 1.1 (σ-algebra)**.** A collection $\mathcal{E}$ of subsets of a set X is called σ-algebra if

(i) $\emptyset \in \mathcal{E}$; $[A \in \mathcal{E}] \Rightarrow [X \setminus A \in \mathcal{E}]$;

(ii) for a sequence $A_n \in \mathcal{E} \Rightarrow \bigcup_{n=0}^{\infty} A_n \in \mathcal{E}$.

Comment: Closed under finite unions, intersections and countable intersections. $A \cap B = X \setminus ((X \setminus A) \cup (X \setminus B))$.

Comment: Elements of $\mathcal{E}$: 'measurable sets'. Pair $(X, \mathcal{E})$: 'measure space'.

Example 1.2. Borel algebra: smallest σ algebra containing all open sets of a topological space.

Comment: Intersection of two σ-algebras is again σ-algebra. 'smallest' is well-defined.

Definition 1.3 (Positive measure and vector measure)**.** For measure space $(X, \mathcal{E})$ a function $\mu : \mathcal{E} \mapsto [0, +\infty]$ is called 'positive measure' if

(i) $\mu(\emptyset) = 0$;

(ii) for pairwise disjoint sequence $A_n \in \mathcal{E} \Rightarrow \mu\left(\bigcup_{n=0}^{\infty} A_n\right) = \sum_{n=0}^{\infty} \mu(A_n)$

For measure space $(X, \mathcal{E})$ and $\mathbb{R}^m$, $m \geq 1$, a function $\mu : \mathcal{E} \mapsto \mathbb{R}^m$ is called 'measure' if μ satisfies (i) and (ii) with absolute convergence.

Comment: Measures are vector space, measures are finite, positive measures may be infinite.

Example 1.4. Examples for measures:

1. counting measure: $\#(A) = |A|$ if A finite, $+\infty$ else.
2. Dirac measure: $\delta_x(A) = 1$ if $x \in A$, 0 else.
3. Lebesgue measure $\mathcal{L}([a, b]) = b - a$ for $b \geq a$.
4. Scaled measures: positive measure μ, function $f \in L^1(\mu)$, new measure $\nu = f \cdot \mu$. $\nu(A) \stackrel{\text{def.}}{=} \int_A f(x)\, \mathrm{d}\mu(x)$.

5. Weak gradient of discontinuous function f, $\mu = D\,f$.
$$\int_\Omega \varphi(x) \cdot \mathrm{d}\mu(x) = -\int_\Omega \operatorname{div}\varphi(x)\, f(x)\, \mathrm{d}x$$
for $\varphi \in C^1(\Omega)$.

Definition 1.5 (Total variation)**.** For measure μ on $(X, \mathcal{E})$ the total variation $|\mu|$ of $A \in \mathcal{E}$ is
$$|\mu|(A) = \sup\left\{\sum_{n=0}^{\infty} |\mu(A_n)| \middle| A_n \in \mathcal{E}, \text{ pairwise disjoint}, \bigcup_{n=0}^{\infty} A_n = A\right\}.$$
$|\mu|$ is finite, positive measure on $(X, \mathcal{E})$.

Comment: Careful with nomenclature in image analysis.

Definition 1.6. A set $N \subset X$ is μ-negligible if $\exists\ A \in \mathcal{E}$ with $N \subset A$ and $\mu(A) = 0$. Two functions f, $g : X \to Y$ are identical 'μ-almost everywhere' when $\{x \in X | f(x) \neq g(x)\}$ is μ-negligible.

Example 1.7. Null sets are Lebesgue-negligible sets.

Definition 1.8 (Measurable functions, push-forward)**.** Let $(X, \mathcal{E})$, $(Y, \mathcal{F})$ be measurable spaces. A function $f : X \to Y$ is 'measurable' if $f^{-1}(A) \in \mathcal{E}$ for $A \in \mathcal{F}$.
For measure μ on $(X, \mathcal{E})$ the 'push-forward' of μ under f to $(Y, \mathcal{F})$, we write $f_\sharp\mu$, is defined by $f_\sharp\mu(A) = \mu(f^{-1}(A))$ for $A \in \mathcal{F}$.
Change of variables formula:
$$\int_X g(f(x))\, \mathrm{d}\mu(x) = \int_Y g(y)\, \mathrm{d}f_\sharp\mu(y)$$

Sketch: Varying densities.

Example 1.9 (Marginal)**.** Let $\mathrm{proj}_i : X \times X \to X$, $\mathrm{proj}_i(x_0, x_1) = x_i$. Marginals of measure γ on $X \times X$:
$$\mathrm{proj}_{0\,\sharp}\gamma(A) = \gamma(A \times X)\,, \qquad \mathrm{proj}_{1\,\sharp}\gamma(A) = \gamma(X \times A)\,.$$

Sketch: Discuss pre-images of proj_i.

Definition 1.10 (Absolute continuity, singularity)**.** Let μ be positive measure, ν measure on measurable space $(X, \mathcal{E})$. ν is 'absolutely continuous' w.r.t. μ, we write $\nu \ll \mu$, if $[\mu(A) = 0] \Rightarrow [\nu(A) = 0]$.

Sketch: Density $\ll$ Lebesgue, density $\not\ll$ density when support different, Dirac measures $\not\ll$ Lebesgue, mixed measures $\not\ll$ density, mixed measures $\ll$ mixed measures when Diracs coincide.

Positive measures μ, ν are 'mutually singular', we write $\mu \perp \nu$, if $\exists\ A \in \mathcal{E}$ such that $\mu(A) = 0$, $\mu(X \setminus A) = 0$. For general measures replace μ, ν by $|\mu|$, $|\nu|$.

Definition 1.11 (σ-finite)**.** A positive measure μ is called σ-finite if $X = \bigcup_{n=0}^\infty A_n$ for sequence $A_n \in \mathcal{E}$ with $\mu(A_n) < +\infty$.

Example 1.12. Lebesgue measure is not finite but σ-finite.

Theorem 1.13 (Radon–Nikodym, Lebesgue decomposition [Ambrosio et al., 2000, Theorem 1.28])**.** Let μ be σ-finite positive measure. ν general measure.
Radon–Nikodym: For $\nu \ll \mu$ there is a function $f \in L^1(\mu)$ such that $\nu = f \cdot \mu$. f is unique μ-almost everywhere. It is called 'density of ν with respect to μ' and usually denoted by $f = \frac{d\nu}{d\mu}$.
Lebesgue decomposition: there exist unique measures ν_a, ν_s such that

$$\nu = \nu_a + \nu_s, \qquad \nu_a \ll \mu, \qquad \nu_s \perp \mu\,.$$

Note: $\nu_a = f \cdot \mu$ for some $f \in L^1(\mu)$.

Corollary 1.14. A real-valued measure ν can be decomposed into $\nu = \nu_+ - \nu_-$ with ν_+, ν_- mutually singular positive measures.

Proof. Since $\nu \ll |\nu|$ there exists $f \in L^1(|\nu|)$ with $\nu = f \cdot |\nu|$. Set $A_+ = f^{-1}((0, +\infty))$, $A_- = f^{-1}((-\infty, 0))$ and set $\nu_\pm(B) = |\nu(B \cap A_\pm)|$. □

Comment: f is only unique $|\nu|$-almost everywhere.

1.2 Duality

References: Kurdila, Zabarankin: Convex functional analysis [Kurdila and Zabarankin, 2005]. For Hilbert spaces: Bauschke, Combettes: Convex Analysis and Monotone Operator Theory in Hilbert Spaces [Bauschke and Combettes, 2011]

Definition 1.15 (Dual space)**.** For normed vector space $(X, \|\cdot\|_X)$ its topological dual space is given by

$$X^* = \{y : X \to \mathbb{R} \,|\, y\,\text{linear, continous, i.e.}\ \exists\, C < \infty,\ |y(x)| \leq C\,\|x\|_X\,\forall\, x \in X\}\,.$$

Norm on X^*:

$$\|y\|_{X^*} = \sup\,\{|y(x)||x \in X,\ \|x\|_X \leq 1\}$$

$(X^*, \|\cdot\|_{X^*})$ is Banach space. For $y(x)$ one often writes $\langle y, x\rangle$ or $\langle y, x\rangle_{X^*,X}$.

Comment: Linear not necessarily continuous in infinite dimensions. Dual norm is operator norm.

Definition 1.16 (Weak convergence)**.** A sequence x_n in X converges weakly to $x \in X$ if $y(x_n) \to y(x)$ for all $y \in X^*$. We write $x_n \rightharpoonup x$.

Definition 1.17 (Weak* convergence)**.** A sequence y_n in X^* converges weakly to $y \in X^*$ if $y_n(x) \to y(x)$ for all $x \in X$. We write $y_n \overset{*}{\rightharpoonup} y$.

Application to measures:

Definition 1.18 (Radon measures)**.** Let (X, d) be compact metric space, let $\mathcal{E}$ be Borel-σ-algebra. A finite measure (positive or vector valued) is called a 'Radon measure'. Write:

- $\mathcal{M}_+(X)$: positive Radon measures,
- $\mathcal{P}(X) \subset \mathcal{M}_+(X)$: Radon probability measures (total mass $= 1$),
- $\mathcal{M}(X)^m$: (vector valued) Radon measures.

Theorem 1.19 (Regularity [Ambrosio et al., 2000, Proposition 1.43])**.** For positive Radon measures on $(X, \mathcal{E})$ one has for $A \in \mathcal{E}$

$$\mu(A) = \sup\{\mu(B) \,|\, B \in \mathcal{E},\, B \subset A,\, B \text{ compact}\} = \inf\{\mu(B) \,|\, B \in \mathcal{E},\, A \subset B,\, B \text{ open}\} .$$

Theorem 1.20 (Duality [Ambrosio et al., 2000, Theorem 1.54])**.** Let (Ω, d) be compact metric space. Let $C(\Omega)^m$ be space of continuous functions from Ω to $\mathbb{R}^m$, equipped with sup-norm. The topological dual of $C(\Omega)^m$ can be identified with the space $\mathcal{M}(\Omega)^m$ equipped with the total variation norm $\|\mu\|_{\mathcal{M}} \stackrel{\text{def.}}{=} |\mu|(\Omega)$. Duality pairing for $\mu \in \mathcal{M}(\Omega)^m$, $f \in C(\Omega)^m$:

$$\mu(f) = \langle \mu, f \rangle_{\mathcal{M},C} = \int_\Omega f(x)\, \mathrm{d}\mu(x)$$

Corollary 1.21. Two measures μ, $\nu \in \mathcal{M}(\Omega)^m$ with $\mu(f) = \nu(f)$ for all $f \in C(\Omega)^m$ coincide.

Theorem 1.22 (Banach–Alaoglu [Kurdila and Zabarankin, 2005, Theorem 2.4.4])**.** Let X be a separable normed space. Any bounded sequence in X^* has a weak$*$ convergent subsequence.

Comment: Since $C(\Omega)$ is separable, any bounded sequence in $\mathcal{M}(\Omega)$ has a weak$*$ convergent subsequence.

1.3 Monge formulation of optimal transport

Comment: Gaspard Monge: French mathematician and engineer, 18^{th} century. Studied problem of optimal allocation of resources to minimize transport cost.

Sketch: Bakeries and cafes

Example 1.23 (According to Villani)**.** Every morning in Paris bread must be transported from bakeries to cafes for consumption. Every bakery produces prescribed amount of bread, every cafe orders prescribed amount. Assume: total amounts identical. Look for most economical way to distribute bread.

Mathematical model:

- $\Omega \subset \mathbb{R}^2$: area of Paris
- $\mu \in \mathcal{P}(\Omega)$: distribution of bakeries and produced amount of bread,
- $\nu \in \mathcal{P}(\Omega)$: distribution of cafes and consumed amount of bread
- Cost function $c : \Omega \times \Omega \to \mathbb{R}_+$. $c(x, y)$ gives cost of transporting 1 unit of bread from bakery at x to cafe at y.
- Describe transport by map $T : \Omega \to \Omega$. Bakery at x will deliver bread to cafe at $T(x)$. Consistency condition: $T_\sharp\mu = \nu$.

Comment: Each cafe receives precisely ordered amount of bread.

- Total cost of transport map

$$C_M(T) = \int_\Omega c(x, T(x))\, \mathrm{d}\mu(x)$$

Comment: For bakery at location x pay $c(x, T(x)) \cdot \mu(x)$. Sum (i.e. integrate) over all bakeries.

Definition 1.24. Monge optimal transport problem: find T that minimizes C_M.

Problems:

- Do maps T with $T_\sharp\mu = \nu$ exist? Can not split mass.

Sketch: Splitting of mass.

- Does minimal T exist? Non-linear, non-convex constraint and objective.

Comment: $\Rightarrow$ problem remained unsolved for long time.

1.4 Kantorovich formulation of optimal transport

Comment: Leonid Kantorovich: Russian mathematician, 20^{th} century. Founding father of linear programming, proposed modern formulation of optimal transport. (Nobel prize in economics 1975.)

Do not describe transport by map T, but by positive measure $\pi \in \mathcal{M}_+(\Omega \times \Omega)$.

Definition 1.25 (Coupling / Transport Plan)**.** Let $\mu, \nu \in \mathcal{P}(\Omega)$. Set of 'couplings' or 'transport plans' $\Pi(\mu, \nu)$ is given by

$$\Pi(\mu, \nu) = \left\{\pi \in \mathcal{P}(\Omega \times \Omega) \,\middle|\, {\mathrm{proj}_0}_\sharp \pi = \mu,\ {\mathrm{proj}_1}_\sharp \pi = \nu\right\} .$$

Example 1.26. $\Pi(\mu, \nu) \neq \emptyset$, contains at least product measure $\mu \otimes \nu \in \Pi(\mu, \nu)$. $(\mu \otimes \nu)(A \times B) = \mu(A) \cdot \nu(B)$ for measurable A, $B \subset \Omega$.

Definition 1.27. For compact metric space (Ω, d), μ, $\nu \in \mathcal{P}(\Omega)$, $c \in C(\Omega \times \Omega)$ the Kantorovich optimal transport problem is given by

$$\mathcal{C}(\mu, \nu) = \inf \left\{ \int_{\Omega \times \Omega} c(x, y)\, \mathrm{d}\pi(x, y) \,\middle|\, \pi \in \Pi(\mu, \nu) \right\} \tag{1}$$

Comment: Linear (continuous) objective, affine constraint set.

Comment: Language of measures covers finite dimensional and infinite dimensional case.

Theorem 1.28. Minimizers of (1) exist.

Proof.

- Let π_n be minimizing sequence. Since $\pi_n \in \mathcal{P}(\Omega \times \Omega)$ have $\|\pi_n\|_{\mathcal{M}} = 1$. By Banach-Alaoglu (Theorem 1.22) $\exists$ converging subsequence. After extraction of subsequence have convergent minimizing sequence $\pi_n \stackrel{*}{\rightharpoonup} \pi$.

- Positivity: π is a positive measure. Otherwise find function $\phi \in C(\Omega \times \Omega)$ with $\int_{\Omega \times \Omega} \phi\, \mathrm{d}\pi < 0$ (use Corollary 1.14 and Theorem 1.19 for construction) which contradicts weak$*$ convergence.

- Unit mass: $\pi(\Omega \times \Omega) = \int_{\Omega \times \Omega} 1\, \mathrm{d}\pi = \lim_{n \to \infty} \int_{\Omega \times \Omega} 1\, \mathrm{d}\pi_n = \pi_n(\Omega \times \Omega) = 1$

- Marginal constraint: For every $\phi \in C(\Omega)$

$$\begin{aligned} \int_\Omega \phi\, \mathrm{d}{\mathrm{proj}_0}_\sharp \pi &= \int_{\Omega \times \Omega} \phi \circ \mathrm{proj}_0\, \mathrm{d}\pi \\ &= \lim_{n \to \infty} \int_{\Omega \times \Omega} \phi \circ \mathrm{proj}_0\, \mathrm{d}\pi_n = \lim_{n \to \infty} \int_\Omega \phi\, \mathrm{d}{\mathrm{proj}_0}_\sharp \pi_n = \int_\Omega \phi\, \mathrm{d}\mu \end{aligned}$$

So ${\mathrm{proj}_0}_\sharp \pi = \mu$. Analogous: ${\mathrm{proj}_1}_\sharp \pi = \nu$.

- So: $\pi \in \Pi(\mu,\nu)$.
- Since $c \in C(\Omega \times \Omega)$ and $\pi_n \overset{*}{\rightharpoonup} \pi$ have

$$\int_{\Omega\times\Omega} c\,\mathrm{d}\pi = \lim_{n\to\infty} \int_{\Omega\times\Omega} c\,\mathrm{d}\pi_n\,.$$

Therefore, π is minimizer.

□

Comment: For proof under more general conditions see for instance [Villani, 2009, Chapter 4]

Proof two additional useful results to get some practice and intuition.

Proposition 1.29 (Restriction [Villani, 2009, Theorem 4.6])**.** Let μ, $\nu \in \mathcal{P}(\Omega)$, $c \in C(\Omega\times\Omega)$, let π be optimizer for $\mathcal{C}(\mu,\nu)$. Let $\tilde{\pi} \in \mathcal{M}_+(\Omega \times \Omega)$, $\tilde{\pi}(\Omega \times \Omega) > 0$, $\tilde{\pi}(A) \leq \pi(A)$ for all measurable $A \subset \Omega \times \Omega$. Set $\pi' = \frac{\tilde{\pi}}{\tilde{\pi}(\Omega\times\Omega)}$, $\pi' \in \mathcal{P}(\Omega \times \Omega)$. Let $\mu' = \mathrm{proj}_{0\,\sharp}\pi'$, $\nu' = \mathrm{proj}_{1\,\sharp}\pi'$. Then π' is minimal for $\mathcal{C}(\mu',\nu')$.

Example 1.30. $\tilde{\pi}(A) \overset{\text{def.}}{=} \pi(A \cap (\Omega_0 \times \Omega_1))$ for Ω_0, $\Omega_1 \subset \Omega$.

Sketch: Restriction to subset. More general restriction.

Proof.
- Assume π' is not optimal. Then there is a measure $\pi'' \in \Pi(\mu',\nu')$ with strictly better cost.
- Consider the measure $\hat{\pi} = \pi - \tilde{\pi} + \tilde{\pi}(\Omega \times \Omega) \cdot \pi''$. $\hat{\pi}$ is a positive measure since $\tilde{\pi} \leq \pi$. $\hat{\pi} \in \mathcal{P}(\Omega \times \Omega)$ since $\pi'' \in \mathcal{P}(\Omega \times \Omega)$.

$$\begin{aligned} \mathrm{proj}_{0\,\sharp}\hat{\pi} &= \mathrm{proj}_{0\,\sharp}\pi - \mathrm{proj}_{0\,\sharp}\tilde{\pi} + \tilde{\pi}(\Omega \times \Omega) \cdot \mathrm{proj}_{0\,\sharp}\pi'' \\ &= \mu - \tilde{\pi}(\Omega \times \Omega) \cdot \left(\mu' - \mu'\right) = \mu \end{aligned}$$

So $\hat{\pi} \in \Pi(\mu,\nu)$.
- $\hat{\pi}$ has lower transport cost than π:

$$\int_{\Omega\times\Omega} c\,\mathrm{d}\hat{\pi} = \int_{\Omega\times\Omega} c\,\mathrm{d}\pi - \tilde{\pi}(\Omega\times\Omega) \int_{\Omega\times\Omega} c\,\mathrm{d}\pi' + \tilde{\pi}(\Omega\times\Omega) \int_{\Omega\times\Omega} c\,\mathrm{d}\pi'' < \int_{\Omega\times\Omega} c\,\mathrm{d}\pi$$

- So π is not optimal which is a contradiction. Therefore π' must be optimal.

□

Proposition 1.31 (Convexity [Villani, 2009, Theorem 4.8])**.** The function $\mathcal{P}(\Omega)^2 \to \mathbb{R}$, $(\mu,\nu) \mapsto \mathcal{C}(\mu,\nu)$ is convex.

Proof.
- Let μ_0, μ_1, ν_0, $\nu_1 \in \mathcal{P}(\Omega)$. Let π_i be corresponding minimizers in $\mathcal{C}(\mu_i,\nu_i)$, $i \in \{0,1\}$.
- For $\lambda \in (0,1)$ set

$$\hat{\mu} = (1-\lambda)\,\mu_0 + \lambda\,\mu_1\,, \qquad \hat{\nu} = (1-\lambda)\,\nu_0 + \lambda\,\nu_1\,, \qquad \hat{\pi} = (1-\lambda)\,\pi_0 + \lambda\,\pi_1\,.$$

- $\hat{\pi} \in \Pi(\hat{\mu},\hat{\nu})$ since

$$\mathrm{proj}_{0\,\sharp}\hat{\pi} = (1-\lambda)\,\mathrm{proj}_{0\,\sharp}\pi_0 + \lambda\,\mathrm{proj}_{0\,\sharp}\pi_1 = (1-\lambda)\,\mu_0 + \lambda\,\mu_1 = \hat{\mu}\,.$$

- Convexity:

$$\mathcal{C}(\hat{\mu},\hat{\nu}) \leq \int_{\Omega\times\Omega} c\,d\hat{\pi} = (1-\lambda) \int_{\Omega\times\Omega} c\,d\pi_0 + \lambda \int_{\Omega\times\Omega} c\,d\pi_1 = (1-\lambda)\,\mathcal{C}(\mu_0,\nu_0) + \lambda\,\mathcal{C}(\mu_1,\nu_1)$$

□

2 Kantorovich duality

2.1 More duality

Definition 2.1 (Topologically paired spaces)**.** Two vector spaces X, X^* with locally convex Hausdorff topology are called *topologically paired spaces* if all continuous linear functionals on one space can be identified with all elements of the other.

Example 2.2. Let (Ω, d) be a compact metric space. $C(\Omega)$ and $\mathcal{M}(\Omega)$ with the sup-norm topology and the weak-$*$ topology are topologically paired spaces.
Any continuous linear functional on $C(\Omega)$ can be identified with an element in $\mathcal{M}(\Omega)$ by construction. If Φ is a weak-$*$ continuous linear functional on $\mathcal{M}(\Omega)$ it can be identified with the continuous function $\varphi : x \mapsto \Phi(\delta_x)$.

Definition 2.3 (Fenchel–Legendre conjugates)**.** Let X, X^* be topologically paired spaces. Let $f : X \to \mathbb{R} \cup \{\infty\}$. Its Fenchel–Legendre conjugate $f^* : X^* \to \mathbb{R} \cup \{\infty\}$ is given by

$$f^*(y) = \sup\{\langle y, x\rangle - f(x) | x \in X\}\,.$$

f^* is convex, lsc on X^*. Likewise, for $g : X^* \to \mathbb{R} \cup \{\infty\}$ define conjugate g^*. If f, g convex, lsc then $f = f^{**}$, $g = g^{**}$.

Comment: Lsc: lower semicontinuous, $[x_n \to x] \Rightarrow [f(x) \leq \liminf_{n\to\infty} f(x_n)]$

Theorem 2.4 (Fenchel–Rockafellar [Rockafellar, 1967])**.** Let (X, X^*), (Y, Y^*) be two pairs of topologically paired spaces. Let $f : X \to \mathbb{R} \cup \{\infty\}$, $g : Y \to \mathbb{R} \cup \{\infty\}$, f, g convex, $A : X \to Y$ linear, continuous. Assume $\exists\, x \in X$ such that f finite at x, g finite and continuous at $A\,x$. Then

$$\inf\{f(x) + g(A\,x) | x \in X\} = \max\{-f^*(-A^*z) - g^*(z) | z \in Y^*\}\,.$$

In particular a maximizer of the problem on the right exists. $A^* : Y^* \to X^*$ is adjoint of A defined by $\langle z, A\,x\rangle_{Y^*,Y} = \langle A^*z, x\rangle_{X^*,X}$.

Comment: Can sometimes be used 'in both directions' to establish existence of both primal and dual problem.

2.2 Dual Kantorovich problem

Theorem 2.5. Given the setting of Definition 1.27 one finds

$$\mathcal{C}(\mu, \nu) = \sup\left\{\int_\Omega \alpha\, \mathrm{d}\mu + \int_\Omega \beta\, \mathrm{d}\nu \middle| \alpha,\, \beta \in C(\Omega), \alpha(x) + \beta(y) \leq c(x, y) \text{ for all } (x, y) \in \Omega^2\right\} \quad (2)$$

Proof.
- Problem (2) can be written as

$$\mathcal{C}(\mu, \nu) = -\inf\left\{f(\alpha, \beta) + g(A(\alpha, \beta)) \middle| (\alpha, \beta) \in C(\Omega)^2\right\}$$

with

$$\begin{aligned}
&f : C(\Omega)^2 \to \mathbb{R}, && (\alpha, \beta) \mapsto -\int_\Omega \alpha\, \mathrm{d}\mu - \int_\Omega \beta\, \mathrm{d}\nu \\
&g : C(\Omega^2) \to \mathbb{R} \cup \{\infty\}, && \psi \mapsto \begin{cases} 0 & \text{if } \psi(x, y) \leq c(x, y) \text{ for all } (x, y) \in \Omega^2 \\ +\infty & \text{else.} \end{cases} \\
&A : C(\Omega)^2 \to C(\Omega^2), && [A(\alpha, \beta)](x, y) = \alpha(x) + \beta(y)\,.
\end{aligned}$$

- f, g are convex, lsc. A is bounded, linear.
- Let (α, β) be two constant, finite functions with $\alpha(x) + \beta(y) < \min\{c(x', y') | (x', y') \in \Omega^2\}$. Then $f(\alpha, \beta) < \infty$, $g(A(\alpha, \beta)) < \infty$ and g is continuous at $A(\alpha, \beta)$. Thus, with Theorem 2.4 (and Example 2.2)

$$\mathcal{C}(\mu, \nu) = \min \left\{ f^*(-A^*\pi) + g^*(\pi) \middle| \pi \in \mathcal{M}(\Omega^2) \right\} .$$

- One obtains:

$$\begin{aligned} f^*(-\rho, -\sigma) &= \sup \left\{ -\int_\Omega \alpha \, \mathrm{d}\rho - \int_\Omega \beta \, \mathrm{d}\sigma + \int_\Omega \alpha \, \mathrm{d}\mu + \int_\Omega \beta \, \mathrm{d}\nu \middle| (\alpha, \beta) \in C(\Omega)^2 \right\} \\ &= \begin{cases} 0 & \text{if } \rho = \mu, \, \sigma = \nu \,, \\ +\infty & \text{else.} \end{cases} \end{aligned}$$

(Reasoning similar than for positivity of limit π in proof of Theorem 1.28.)

$$\begin{aligned} g^*(\pi) &= \sup \left\{ \int_{\Omega^2} \psi \, \mathrm{d}\pi \middle| \psi \in C(\Omega^2), \psi(x, y) \leq c(x, y) \text{ for all } (x, y) \in \Omega^2 \right\} \\ &= \begin{cases} \int_{\Omega^2} c \, \mathrm{d}\pi & \text{if } \pi \in \mathcal{M}_+(\Omega^2), \\ +\infty & \text{else.} \end{cases} \end{aligned}$$

□

So far we have not yet proven existence of dual maximizers. For this we need some additional arguments. We follow the presentation in [Santambrogio, 2015, Section 1.2].

Definition 2.6 (c-transform)**.** For $\psi \in C(\Omega)$ define its c-transform $\psi^c \in C(\Omega)$ by

$$\psi^c(y) = \inf \{c(x, y) - \psi(x) | x \in \Omega\}$$

and its $\bar{c}$-transform $\psi^{\bar{c}} \in C(\Omega)$ by

$$\psi^{\bar{c}}(x) = \inf \{c(x, y) - \psi(y) | y \in \Omega\}$$

A function ψ is called $\bar{c}$-concave if it can be written as $\psi = \phi^c$ for some $\phi \in C(\Omega)$. Analogously, ψ is c-concave if it can be written as $\psi = \phi^{\bar{c}}$.

Comment: Setting $\beta = \alpha^c$ (or $\alpha = \beta^{\bar{c}}$) in (2) corresponds to optimization over β for fixed α (and vice versa). In general alternating optimization of (2) in α and β does not yield an optimal solution.

Lemma 2.7. The set of c-concave and $\bar{c}$-concave functions are equicontinuous.

Proof.
- Since $c \in C(\Omega \times \Omega)$ and (Ω, d) compact there is a continuous function $\omega : \mathbb{R}_+ \to \mathbb{R}_+$ with $\omega(0) = 0$ such that $|c(x, y) - c(x, y')| \leq \omega(d(y, y'))$.
- Let $\psi = \phi^c$. Set $\phi_x : y \mapsto c(x, y) - \phi(x)$. For every $x \in \Omega$ have $|\phi_x(y) - \phi_x(y')| \leq \omega(d(y, y'))$. One finds

$$\psi(y) \leq \phi_x(y) \leq \phi_x(y') + \omega(d(y, y'))$$

for all $x, y, y' \in \Omega$. Taking the infimum over x one gets $\psi(y) \leq \psi(y') + \omega(d(y, y'))$ and by symmetry $|\psi(y) - \psi(y')| \leq \omega(d(y, y'))$. This implies equicontinuity of $\bar{c}$-concave functions.

- Argument for $\phi^{\overline{c}}$ analogous.

$\square$

Theorem 2.8 (Arzelà-Ascoli [Rudin, 1986, Thm. 11.28])**.** If (Ω, d) is a compact metric space and $(f_n)_n$ is a sequence of uniformly bounded, equicontinuous functions in $C(\Omega)$ then $(f_n)_n$ has a uniformly converging subsequence.

Theorem 2.9 ([Santambrogio, 2015, Prop. 1.11])**.** Maximizers of (2) exist. For an optimal pair (α, β) one finds $\beta = \alpha^c$, $\alpha = \beta^{\overline{c}}$.

Proof.

- For feasible (α, β) with finite score in (2) one can always replace β by α^c and subsequently α by $(\alpha^c)^{\overline{c}}$ which are still feasible and do not decrease the functional value. Hence, we may impose the additional constraint that (α, β) in (2) are $(c, \overline{c})$-concave and their respective $(c, \overline{c})$-transforms.

- Replacing feasible (α, β) in (2) by

$$(x \mapsto \alpha(x) - C, y \mapsto \beta(y) + C) \qquad \text{with } C = \min_{x' \in \Omega} \alpha(x')$$

  does not change the functional value or affect the constraints.

- Arguing as in Lemma 2.7 one finds for c-concave α with $\min_x \alpha(x) = 0$ that $\alpha(x) \in [0, \omega(\operatorname{diam} \Omega)]$ and for the corresponding $\beta = \alpha^c$ that $\beta(y) \in [\min c - \omega(\operatorname{diam} \Omega), \max c]$.

- Hence, we may consider maximizing sequences of (2) that are uniformly bounded and equicontinuous. By the Arzelà-Ascoli Theorem there exists a uniformly converging subsequence. Since the objective (and the constraints) of (2) are upper semicontinuous (see proof of Theorem 2.5), the limit must be a maximizer.

$\square$

Corollary 2.10 (Primal-dual optimality condition)**.** π solves (1) and (α, β) solve (2) if and only if $\alpha(x) + \beta(y) = c(x, y)$ π-almost everywhere.

Proof.

- $\Rightarrow$: Assume π, (α, β) are primal and dual optimal then:

$$\int_{\Omega \times \Omega} c(x, y)\, \mathrm{d}\pi(x, y) = \int_\Omega \alpha\, \mathrm{d}\mu + \int_\Omega \beta\, \mathrm{d}\nu = \int_{\Omega \times \Omega} [\alpha(x) + \beta(y)] \mathrm{d}\pi(x, y)$$

  And $\alpha(x) + \beta(y) \leq c(x, y)$ for all $(x, y) \in \Omega^2$. Therefore $\alpha(x) + \beta(y) = c(x, y)$ π-a.e..

- $\Leftarrow$: Assume $\alpha(x) + \beta(y) = c(x, y)$ π-a.e..

$$\int_\Omega \alpha\, \mathrm{d}\mu + \int_\Omega \beta\, \mathrm{d}\nu = \int_{\Omega \times \Omega} [\alpha(x) + \beta(y)] \mathrm{d}\pi(x, y) = \int_{\Omega \times \Omega} c(x, y)\, \mathrm{d}\pi(x, y)$$

$\square$

Remark 2.11 (Economic Interpretation of Kantorovich Duality)**.** Bakeries and cafes hire a third-party company to do the transportation and agree to split the transport cost. When picking up bread at bakery x in the morning, the company charges an advance payment $\alpha(x)$ per unit of bread for the transport. Upon delivery at a cafe at y it charges a final payment $\beta(y)$ per unit of bread from the cafe.

The total payment to the company will be $\int_\Omega \alpha \, \mathrm{d}\mu + \int_\Omega \beta \, \mathrm{d}\nu$. It is left to the company to decide which bread to deliver where. And they will want to minimize the total transport cost, i.e. to find the *global minimum* of $\int_{\Omega\times\Omega} c \, \mathrm{d}\pi$.
But it can never charge more than $c(x, y) - \alpha(x)$ when dropping of bread from x at y, otherwise the cafe y may complain and try to hire another company to get bread from bakery x at a lower price. When every cafe receives bread from its 'subjectively cheapest' bakery (and similarly each bakery delivers to its 'subjectively cheapest' cafe), the transport plan is said to be at equilibrium: no party will attempt to change its partner in a local attempt to reduce its costs.
Kantorovich duality states that for the optimal transport model the global minimum and equilibrium coincide.

A useful application of duality is the following result which is also the foundation for the numerical approximation of the Kantorovich problem.

Proposition 2.12 (Stability of optimal plans)**.** Let $(\mu_n)_n$, $(\nu_n)_n$ be sequences in $\mathcal{P}(\Omega)$ converging weak$*$ to μ and ν respectively. Let $(\pi_n)_n$ be a corresponding sequence of optimal plans. Then any cluster point of $(\pi_n)_n$ is an optimal coupling for $\mathcal{C}(\mu, \nu)$.

Comment: $(\pi_n)_n$ will always have cluster points due to Theorem 1.22.

Proof.

- Let π be a cluster point of $(\pi_n)_n$. Without changing notation let $(\pi_n)_n$ be a subsequence converging weak$*$ to π. Then $\pi \in \Pi(\mu, \nu)$ as for any $\phi \in C(\Omega)$:

$$\int_\Omega \phi \mathrm{d}(\mathrm{proj}_{0\,\sharp}\pi) = \int_{\Omega\times\Omega} \phi \circ \mathrm{proj}_0 \, \mathrm{d}\pi = \lim_{n\to\infty} \int_{\Omega\times\Omega} \phi \circ \mathrm{proj}_0 \, \mathrm{d}\pi_n = \lim_{n\to\infty} \int_\Omega \phi \, \mathrm{d}\mu_n = \int_\Omega \phi \, \mathrm{d}\mu$$

- Since π is feasible for $\mathcal{C}(\mu, \nu)$, for this converging subsequence:

$$\mathcal{C}(\mu, \nu) \leq \int_{\Omega\times\Omega} c \, \mathrm{d}\pi = \lim_{n\to\infty} \int_{\Omega\times\Omega} c \, \mathrm{d}\pi_n = \lim_{n\to\infty} \mathcal{C}(\mu_n, \nu_n)$$

- Since this is true for any converging sub-sequence, get:

$$\mathcal{C}(\mu, \nu) \leq \liminf_{n\to\infty} \mathcal{C}(\mu_n, \nu_n)$$

- Let $\varepsilon > 0$.

- Let $(\alpha_n, \beta_n)_n$ be a sequence of dual optimizers for $\mathcal{C}(\mu_n, \nu_n)$. Arguing as in the proof of Theorem 2.9 we can extract a converging subsequence $(\alpha_n, \beta_n)_n$, converging uniformly to some (α, β). Note that $\alpha(x) + \beta(y) \leq c(x, y)$ for all $(x, y) \in \Omega \times \Omega$.

- There is some $N \in \mathbb{N}$ such that $|\alpha - \alpha_n| \leq \varepsilon/2$, $|\beta - \beta_n| \leq \varepsilon/2$ for all $n \geq N$. So:

$$\int_\Omega \alpha \, \mathrm{d}\mu_n + \int_\Omega \beta \, \mathrm{d}\nu_n \geq \int_\Omega \alpha_n \, \mathrm{d}\mu_n + \int_\Omega \beta_n \, \mathrm{d}\nu_n - \varepsilon$$

- This is true for all converging subsequences $(\alpha_n, \beta_n)_n$. Taking the supremum over the limit superior for all such subsequences, we get

$$\mathcal{C}(\mu, \nu) \geq \int_\Omega \alpha \, \mathrm{d}\mu + \int_\Omega \beta \, \mathrm{d}\nu \geq \limsup_{n\to\infty} \mathcal{C}(\mu, \nu) - \varepsilon$$

- Since this is true for any $\varepsilon > 0$ we find

$$\mathcal{C}(\mu, \nu) \geq \limsup_{n\to\infty} \mathcal{C}(\mu, \nu)$$

and thus π is optimal for $\mathcal{C}(\mu, \nu)$. (And $\mathcal{C}(\mu_n, \nu_n)$ converges.)

□

Comment: The proof can be extended to cover a sequence of changing cost functions $(c_n)_n$ in $C(\Omega \times \Omega)$, where c_n is used for $\mathcal{C}(\mu_n, \nu_n)$ if $(c_n)_n$ converges uniformly to a limit $c \in C(\Omega \times \Omega)$.

Comment: For treatment of duality in more general regularity setting see for instance [Villani, 2009, Chapter 5]. A preview of the required concepts is given in subsection below.

2.3 c-cyclical monotonicity and duality

The proof for Proposition 2.12 relies on the uniform convergence of the dual potentials (α_n, β_n). This is not available in less regular settings and a fundamentally different argument has to be used relying on the following property:

Definition 2.13 (c-cyclical monotonicity [Santambrogio, 2015, Def. 1.36])**.** Let $c \in C(\Omega \times \Omega)$. A set $\Gamma \subset \Omega \times \Omega$ is c-cyclical monotone (short: c-CM) if for every $n \in \mathbb{N}$ and every tuple of points $((x_1, y_1), \ldots, (x_n, y_n)) \in \Gamma^n$ one has

$$\sum_{i=1}^{n} c(x_i, y_i) \leq \sum_{i=1}^{n} c(x_i, y_{i-1})$$

with convention $y_0 := y_n$.

Comment: Intuition behind this: assume bakery x_i delivers to cafe y_i, $(x_i, y_i) \in \Gamma \subset \Omega \times \Omega$ and Γ is c-CM. Now assume bakery x_1 decides to deliver to cafe y_n instead, which then rejects bread from bakery x_n. This bakery has now reroute its bread to cafe y_{n-1} and so forth, until eventually a cycle occurs and bakery x_2 reroutes its bread to cafe y_1. The fact that Γ is c-CM implies that such a cyclic rerouting can never improve the transport cost.

Definition 2.14 (Support of measure)**.** Let (Ω, d) be a compact metric space with its Borel σ-algebra and $\mu \in \mathcal{M}_+(\Omega)$. The support of μ, denoted $\operatorname{spt}\mu$ is the smallest closed set $A \subset \Omega$ such that $\mu(A) = \mu(\Omega)$. For $x \in \operatorname{spt}\mu$ one has $\mu(B_r(x)) > 0$ for any $r > 0$.

From the above comment we deduce intuitively: if π is optimal transport plan one must have $\operatorname{spt}\pi$ is c-CM. Otherwise a cyclic rerouting of mass, as above, could yield an improved plan. Formal statement:

Proposition 2.15 ([Santambrogio, 2015, Thm. 1.38])**.** If π is an optimal transport plan for $\mathcal{C}(\mu, \nu)$ then $\operatorname{spt}\pi$ is c-CM.

Proof.

- Assume $\operatorname{spt}\pi$ is not c-CM. Then there is $n \in \mathbb{N}$, and a tuple $((x_1, y_1), \ldots, (x_n, y_n)) \in (\operatorname{spt}\pi)^n$ with

$$\sum_{i=1}^{n} c(x_i, y_i) - \sum_{i=1}^{n} c(x_i, y_{i-1}) = \varepsilon > 0\,.$$

(and all x_i, y_i different).

- Select small open environments $U_i \subset \Omega$ of x_i, and $V_i \subset \Omega$ of y_i such that

$$|c(x,y) - c(x',y')| \leq \tfrac{\varepsilon}{4n} \quad \text{for all} \quad x, x' \in U_i,\, y, y' \in V_j$$

  for all $i, j \in \{1, \dots, n\}$. (and U_i, V_i pairwise disjoint).

- Set $\delta = \min_i \pi(U_i \times V_i) > 0$ since $(x_i, y_i) \in \operatorname{spt} \pi$.

- Set:

$$\pi_i = \frac{\pi \llcorner (U_i \times V_i)}{\pi(U_i \times V_i)}, \qquad \mu_i = {\operatorname{proj}_0}_\sharp \pi_i, \qquad \nu_i = {\operatorname{proj}_1}_\sharp \pi_i, \qquad \hat{\pi}_i = \mu_i \otimes \nu_{i-1}\,.$$

- Let:

$$\hat{\pi} = \pi - \delta \sum_{i=1}^{n} \pi_i + \delta \sum_{i=1}^{n} \hat{\pi}_i$$

  Note: $\hat{\pi} \geq 0$, $\hat{\pi} \in \Pi(\mu, \nu)$.

- New cost:

$$\begin{aligned}
\int c \,\mathrm{d}\hat{\pi} &= \int c \,\mathrm{d}\pi + \delta \sum_{i=1}^{n} \left(\underbrace{\int c \,\mathrm{d}\hat{\pi}_i}_{\leq c(x_i, y_{i-1}) + \frac{\varepsilon}{4n}} - \underbrace{\int c \,\mathrm{d}\pi_i}_{\geq c(x_i, y_i) - \frac{\varepsilon}{4n}} \right) \\
&\leq \int c \,\mathrm{d}\pi + \delta \left[\underbrace{\sum_{i=1}^{n} (c(x_i, y_{i-1}) - c(x_i, y_i))}_{=-\varepsilon} + \tfrac{\varepsilon}{2} \right] \\
&= \int c \,\mathrm{d}\pi - \tfrac{\delta \varepsilon}{2}\,.
\end{aligned}$$

  So π cannot be optimal.

□

The converse implication is much less clear: if $\operatorname{spt} \pi$ is c-CM, is π an optimal transport plan? While there are no cyclical rearrangements, possibly there is a more complicated way to improve the plan. With duality we can show that c-CM is indeed sufficient for optimality.

Proposition 2.16 ([Santambrogio, 2015, Thm. 1.37])**.** Let $c \in C(\Omega \times \Omega)$, $\Gamma \subset \Omega^2$, $\Gamma \neq \emptyset$, Γ is c-CM. Then there exists a c-concave function $\alpha \in C(\Omega)$ such that

$$\alpha(x) + \alpha^c(y) = c(x,y) \quad \text{for all} \quad (x,y) \in \Gamma\,.$$

For proof use small auxiliary Lemma.

Lemma 2.17. Let $c \in C(\Omega \times \Omega)$, $\beta : \Omega \to \mathbb{R} \cup \{-\infty\}$, β bounded from above, β not identical $-\infty$. Set

$$\alpha(x) := \inf\{c(x,y) - \beta(y) \,|\, y \in \Omega\}.$$

Then $\alpha \in C(\Omega)$, $(\alpha^c)^{\bar{c}} = \alpha$ which also implies that α is c-concave.

Proof of Proposition 2.16. • β is bounded from above, and finite at least at one point: family of functions $(x \mapsto c(x,y) - \beta(y))_{y\in\Omega:\beta(y)>-\infty}$ is non-empty and uniformly bounded from below. $\Rightarrow \alpha$ is pointwise infimum over non-empty family of equicontinuous functions uniformly bounded from below. $\Rightarrow \alpha \in C(\Omega)$.

- Now:

$$\begin{aligned}\alpha(x) &= \inf_y \{c(x,y) - \beta(y)\}\\ \alpha^c(y') &= \inf_x \sup_y \{c(x,y') - c(x,y) + \beta(y)\}\\ (\alpha^c)^{\overline{c}}(x') &= \inf_{y'} \sup_x \inf_y \{c(x',y') - c(x,y') + c(x,y) - \beta(y)\}\end{aligned}$$

- By setting $x = x'$ in supremum get: $(\alpha^c)^{\overline{c}}(x') \geq \alpha(x')$.
- By setting $y = y'$ in inner infimum get: $(\alpha^c)^{\overline{c}}(x') \leq \alpha(x')$.

$\square$

Proof. • Pick $(x_1, y_1) \in \Gamma$. For $y \in \Omega$ set

$$\beta(y) = \sup\left\{\sum_{i=1}^n c(x_i,y_i) - \sum_{i=2}^n c(x_i,y_{i+1}) \middle| n \in \mathbb{N}, (x_i,y_i) \in \Gamma \text{ for } i = 1,\ldots,n,\ y_n = y\right\}$$

- For $y \notin \mathrm{proj}_1(\Gamma)$ find $\beta(y) = -\infty$ (supremum over empty set).
- For $y \in \mathrm{proj}_1(\Gamma)$ use c-CM of Γ:

$$\beta(y) = \sup\left\{\underbrace{\sum_{i=1}^n c(x_i,y_i) - \sum_{i=2}^n c(x_i,y_{i+1})}_{\leq c(x_1,y_n)} \middle| n \in \mathbb{N}, (x_i,y_i) \in \Gamma \text{ for } i = 1,\ldots,n,\ y_n = y\right\}$$

  So β is bounded from above.
- For $y = y_1$ get by setting $n = 1$:

$$\beta(y_1) \geq c(x_1, y_1)$$

  So β is not identical to $-\infty$.
- Now for $x \in \Omega$ set

$$\alpha(x) = \inf\left\{c(x,y) - \beta(y) | y \in \Omega\right\}.$$

  By Lemma 2.17: $\alpha \in C(\Omega)$, $(\alpha^c)^{\overline{c}} = \alpha$.
- Now let $(x,y) \in \Gamma$. We need: $\alpha(x) + \alpha^c(y) = c(x,y)$. Since $\alpha^c \geq \beta$ and $\alpha(x) + \alpha^c(y) \leq c(x,y)$ a sufficient condition is $\alpha(x) + \beta(y) \geq c(x,y)$.

- For every $\varepsilon > 0$ there is some $\hat{y}$ such that

$$\alpha(x) \geq c(x, \hat{y}) - \beta(\hat{y}) - \varepsilon\,.$$

  And since $\alpha(x) \in \mathbb{R}$ have $\hat{y} \in \mathrm{proj}_1(\Gamma)$.

- For β get recursive formula:

$$\beta(y) = \sup\left\{c(x, y) - c(x, \hat{y}) + \beta(\hat{y}) | (x, y) \in \Gamma, \hat{y} \in \mathrm{proj}_1(\Gamma)\right\}$$

  So:

$$\beta(y) \geq c(x, y) - c(x, \hat{y}) + \beta(\hat{y})\,.$$

- Combining lower bounds for $\alpha(x)$ and $\beta(y)$ we get:

$$\alpha(x) + \beta(y) \geq c(x, \hat{y}) - \beta(\hat{y}) - \varepsilon + c(x, y) - c(x, \hat{y}) + \beta(\hat{y}) = c(x, y) - \varepsilon\,.$$

  Since this is true for any $\varepsilon > 0$ it is true for $\varepsilon = 0$.

□

Application: sufficient condition for optimality of transport plans:

Corollary 2.18. If $\pi \in \Pi(\mu, \nu)$ and $\mathrm{spt}\,\pi$ is c-CM, then π is an optimal coupling for $\mathcal{C}(\mu, \nu)$.

Proof. Use $(\alpha, \beta = \alpha^c)$ for function α provided by Proposition 2.16 as dual feasible candidates.

□

Another application: alternative proof for stability result Proposition 2.12. Need a few ingredients.

Definition 2.19. For metric space (Ω, d) define Hausdorff distance for subsets A, B of Ω:

$$d_H(A, B) = \max\{\max\{d(x, B) \,|\, x \in A\}, \max\{d(A, y) \,|\, y \in B\}\}$$

Theorem 2.20 (Blaschke [Ambrosio and Tilli, 2004, Thm. 4.4.15])**.** For a compact metric space (Ω, d) the set of compact subsets of Ω with the distance d_H is a compact metric space.

Lemma 2.21. Let A_n be a sequence of compact subsets of Ω, $A \subset \Omega$ compact, let $A_n \to A$ in the Hausdorff distance. Then for every $x \in A$ there is a sequence $(x_n)_n$, $x_n \in A_n$ such that $x_n \to x$.

Proof.
- Let $x \in A$. For any $\varepsilon > 0$ there is some N such that $d_H(A, A_n) \leq \varepsilon$ for $n \geq N$. Then:

$$d(x, A_n) \leq d_H(A, A_n) \leq \varepsilon$$

  So there is some $x_n \in A_n$ such that $d(x, x_n) \leq \varepsilon$.

□

Lemma 2.22. Compact metric space (Ω, d). Hausdorff convergent sequence of compact subsets A_n to A. Weak$*$ convergent sequence of measures $(\mu_n)_n$ in $\mathcal{P}(\Omega)$ to μ. $\mathrm{spt}\,\mu_n \subset A_n$.
Then $\mathrm{spt}\,\mu \subset A$.

Proof.

- Consider sequence of functions $f_n \in C(\Omega)$, $f_n : x \mapsto d(x, A_n)$ and set $f : x \mapsto d(x, A)$. Then $f_n \to f$ uniformly (with Lemma 2.21).
- So for every $\varepsilon > 0$ there is some $N < \infty$ such that $|f_n - f| \leq \varepsilon$ for $n \geq N$.
- Then:

$$\int f_n \, \mathrm{d}\mu_n = 0$$

$$\int f \, \mathrm{d}\mu = \lim_{n\to\infty} \int f \, \mathrm{d}\mu_n \leq \lim_{n\to\infty} \int f_n \, \mathrm{d}\mu_n + \varepsilon = \varepsilon$$

□

Alternative proof for Proposition 2.12.

- As in first proof: let $(\pi_n)_n$ be sequence of optimizers. For any converging subsequence the limit π is in $\Pi(\mu, \nu)$.
- With Theorem 2.20 the sequence $(\operatorname{spt} \pi_n)_n$ has a convergent subsequence in the Hausdorff metric. Denote limit set by Γ.
- By Lemma 2.22 have $\operatorname{spt} \pi \subset \Gamma$ for any cluster point π of $(\pi_n)_n$.
- Every $\operatorname{spt} \pi_n$ is c-CM by Proposition 2.15. Therefore, so is Γ. Indeed: let $n \in \mathbb{N}$, $((x_1, y_1), \ldots, (x_n, y_n)) \in \Gamma^n$. For $i = 1, \ldots, n$ let $((x_{i,k}, y_{i,k}))_k$ be sequence with $(x_{i,k}, y_{i,k}) \in \operatorname{spt} \pi_k$, with $(x_{i,k}, y_{i,k}) \to (x_i, y_i)$. For every k find by c-CM of $\operatorname{spt} \pi_k$:

$$\sum_{i=1}^{n} c(x_{i,k}, y_{i,k}) \leq \sum_{i=1}^{n} c(x_{i,k}, y_{i-1,k})$$

  Hence this is also true in limit.
- So $\operatorname{spt} \pi \subset \Gamma$ is c-CM. With Corollary 2.18 π is optimal for $\mathcal{C}(\mu, \nu)$.

□

2.4 Solution to the Monge problem

We now consider a special case for which the Monge problem has a solution. Duality will be an important ingredient in the proof.
First we briefly discuss that the Kantorovich formulation of optimal transport, Definition 1.27, can be interpreted as a relaxation of the Monge formulation, Definition 1.24.

Proposition 2.23 (Kantorovich is a relaxation of the Monge problem)**.** Assume $T : \Omega \to \Omega$ is a feasible transport map for the Monge problem between μ and ν, Definition 1.24. In particular $T_\sharp \mu = \nu$.
Let

$$(\mathrm{id}, T) : \Omega \to \Omega \times \Omega, \qquad x \mapsto (x, T(x)) \,.$$

Then $\pi = (\mathrm{id}, T)_\sharp \mu \in \Pi(\mu, \nu)$ and

$$\int_{\Omega \times \Omega} c \, \mathrm{d}\pi = \int_\Omega c(x, T(x)) \, \mathrm{d}\mu(x) \,.$$

Proof.

- Clearly $\pi \in \mathcal{P}(\Omega \times \Omega)$.
- $\text{proj}_0 \circ T = \text{id}$. Hence
$$\text{proj}_{0\sharp}\pi = \text{proj}_{0\sharp}T_\sharp\mu = \mu\,.$$
- Similarly $\text{proj}_1 \circ T = T$. Hence
$$\text{proj}_{1\sharp}\pi = \text{proj}_{1\sharp}T_\sharp\mu = T_\sharp\mu = \nu\,.$$
- Equality of cost follows from change of variables under (id, T).

□

This implies in particular that $\mathcal{C}(\mu,\nu) \le C_M(\mu,\nu)$ since every feasible Monge map induces a Kantorovich coupling of equal cost.
The converse inequality is in general not true but we will now prove it for a special case.

Theorem 2.24 (Solution to the Monge problem)**.** Let $\Omega \subset \mathbb{R}^d$ be compact, let the cost function c be given by $c(x,y) = h(x-y)$ for a strictly convex function $h : \mathbb{R}^d \to \mathbb{R}$. Let μ be Lebesgue-absolutely continuous and let $\partial\Omega$ be μ-negligible.
Then the optimal transport plan π is supported on the graph of a transport map $T : \Omega \to \Omega$.

Proof.

- h is convex and finite. Hence it is continuous and locally Lipschitz. Therefore, c is continuous and Lipschitz (since Ω is compact).
- Therefore Theorem 1.28 and Theorem 2.9 apply and provide existence of primal and dual optimizers π and (α,β).
- From the proof of Theorem 2.9 know: $\beta = \alpha^c$, $\alpha = \beta^{\bar{c}}$. Analogous to Lemma 2.7 this implies that α and β are Lipschitz.
- By Rademacher's theorem (see e.g. [Ziemer, 1989, Theorem 2.2.1]) α is Lebesgue-almost everywhere differentiable in $\text{int}\,\Omega$. And consequently μ-almost everywhere on Ω (since $\partial\Omega$ is μ-negligible).
- From Corollary 2.10: $\alpha(x)+\beta(y) = c(x,y)$ π-almost everywhere. For (x_0,y_0) with $\alpha(x_0)+\beta(y_0) = c(x_0,y_0)$ we find
$$x \mapsto c(x,y_0) - \alpha(x)$$
is minimal at x_0 (since $\beta(y_0) = \inf_x\{c(x,y_0)-\alpha(x)\} = c(x_0,y_0)-\alpha(x_0)$). If α is differentiable at x_0 (which it is μ-a.e., i.e. for (x,y) π-a.e.), then $\nabla\alpha(x_0) \in \partial h(x_0-y_0)$.
- For a strictly convex function ∂h is 'invertible'. That is, for every $v \in \mathbb{R}^d$ there is a unique $w \in \mathbb{R}^d$ such that $v \in \partial h(w)$. We denote this map by ∂h^{-1} and find
$$x_0 - y_0 = \partial h^{-1}(\nabla\alpha(x_0))\,.$$
- This relation is still true π-almost everywhere. Set $T(x) = x - \partial h^{-1}(\nabla\alpha(x))$. Then $y = T(x)$ π-almost everywhere.

- Equality of cost:

$$\int_{\Omega\times\Omega} c(x,y)\,\mathrm{d}\pi(x,y) = \int_{\Omega\times\Omega} c(x,T(x))\,\mathrm{d}\pi(x,y) = \int_{\Omega} c(x,T(x))\,\mathrm{d}\mu(x)$$

- Push-forward condition:

$$\int_{\Omega} \phi(y)\,\mathrm{d}T_{\sharp}\mu(y) = \int_{\Omega} \phi(T(x))\,\mathrm{d}\mu(x) = \int_{\Omega\times\Omega} \phi(T(x))\,\mathrm{d}\pi(x,y)$$
$$= \int_{\Omega\times\Omega} \phi(y)\,\mathrm{d}\pi(x,y) = \int_{\Omega} \phi(y)\,\mathrm{d}\nu(y)$$

□

Example 2.25 (Quadratic case: $c(x,y) = \frac{1}{2}\|x-y\|^2$)**.** This corresponds to $h(x) = \frac{1}{2}\|x\|^2$. Consequently, $\partial h(x) = \{x\}$ and $\partial h^{-1}(x) = x$. Moreover since $\alpha = \beta^{\bar{c}}$:

$$\begin{aligned}\alpha(x) &= \inf\left\{\tfrac{1}{2}\|x-y\|^2 - \beta(y)\middle|y\in\Omega\right\} = \tfrac{1}{2}\|x|^2 + \inf\left\{-\langle x,y\rangle + \tfrac{1}{2}\|y\|^2 - \beta(y)\middle|y\in\Omega\right\}\\ &= \tfrac{1}{2}\|x\|^2 - \underbrace{\sup\left\{\langle x,y\rangle - g(y)\middle|y\in\mathbb{R}^d\right\}}_{:=\phi(x):\text{ convex}}\end{aligned}$$

Therefore,

$$T(x) = x - \nabla\alpha(x) = x - (x - \nabla\phi(x)) = \nabla\phi(x)\,.$$

So T is almost everywhere the gradient of a convex function. This is part of the famous polar factorization theorem by [Brenier, 1991].

3 Wasserstein spaces

3.1 Definition and basic properties

Definition 3.1 (Wasserstein distance)**.** Let (Ω,d) be a compact metric space. For $p\in[1,\infty)$ let $W_p : \mathcal{P}(\Omega)\times\mathcal{P}(\Omega)\to\mathbb{R}$,

$$W_p(\mu,\nu) = \left(\inf\left\{\int_{\Omega\times\Omega} d(x,y)^p\,\mathrm{d}\pi(x,y)\middle|\pi\in\Pi(\mu,\nu)\right\}\right)^{1/p}$$

Comment: On non-compact spaces one usually restricts the Wasserstein space to measures with finite moment of order p, i.e., $\int_\Omega d(x,x_0)^p\,\mathrm{d}\mu < +\infty$ for some arbitrary reference point $x_0\in\Omega$. This is a sufficient condition to keep W_p finite.

Example 3.2. Dirac measures are isometric embedding of Ω into $\mathcal{P}(\Omega)$: $W_p(\delta_x,\delta_y) = d(x,y)$, since $\Pi(\delta_x,\delta_y) = \{\delta_{(x,y)}\}$.

To prove that W_p is indeed a distance we will rely on the following powerful theorem which is often useful to dissect and reassemble measures with certain sought-after properties.

Theorem 3.3 (Disintegration [Ambrosio et al., 2005, Theorem 5.3.1])**.** Let $\tilde{\Omega}$, Ω be compact metric spaces, let $f : \tilde{\Omega} \to \Omega$ be measurable and $\pi \in \mathcal{P}(\tilde{\Omega})$. Set $\mu = f_\sharp \pi \in \mathcal{P}(\Omega)$. Then there is a family $(\pi_y)_{y \in \Omega}$ in $\mathcal{P}(\tilde{\Omega})$, unique μ-a.e., such that $\pi_y(f^{-1}(\{y\})) = 1$ and for $\phi \in C(\tilde{\Omega})$ one has

$$\int_{\tilde{\Omega}} \phi \, \mathrm{d}\pi = \int_{\Omega'} \left(\int_{\tilde{\Omega}} \phi \, \mathrm{d}\pi_y \right) \mathrm{d}\mu(y) \, .$$

Sketch: Table, disintegration.

Comment: Disintegration formalizes the notion of conditional probability. It is easiest to visualize in a discrete case when $\tilde{\Omega} = \Omega \times \Omega$ and $f = \mathrm{proj}_0$. Then π can be interpreted as table and any π_y will be the restriction of π to row y, renormalized to mass 1 (if the row is non-empty). π_y gives the probabilities of picking a given column under the condition that row y has already been selected.

Example 3.4 (Disintegration of transport plan)**.** Let $\pi \in \Pi(\mu, \nu)$. Let $(\gamma_x)_{x \in \Omega}$ be the disintegration of π with respect to proj_0. That is, for any $\phi \in C(\Omega \times \Omega)$ have

$$\int_{\Omega \times \Omega} \phi(x, y) \, \mathrm{d}\pi(x, y) = \int_{\Omega} \left(\int_{\Omega} \phi(x, y) \, \mathrm{d}\gamma_x(y) \right) \mathrm{d}\mu(x) \, .$$

γ_x can be interpreted as describing where mass particles starting in x are going. Note that it is only uniquely defined μ-a.e..

Comment: By the disintegration theorem γ_x would be in $\mathcal{P}(\Omega \times \Omega)$. But since $\gamma_x(\mathrm{proj}_0^{-1}(\{x\})) = \gamma_x(\{x\} \times \Omega) = 1$ we can interpret γ_x as element of $\mathcal{P}(\Omega)$.

Theorem 3.5. W_p is a metric on $\mathcal{P}(\Omega)$.

Proof.

- W_p is non-negative (since $d(x, y)^p \geq 0$), symmetric (since $d(x, y)^p$ is symmetric) and finite (since Ω is compact, i.e., d is bounded).
- Let $T : \Omega \to \Omega \times \Omega$, $T(x) = (x, x)$ be the 'diagonal' embedding of Ω into $\Omega \times \Omega$. $W_p(\mu, \mu) = 0$, since $\pi = T_\sharp \mu \in \Pi(\mu, \mu)$ and $\int d^p \, \mathrm{d}\pi = 0$: Note that $(\mathrm{proj}_i \circ T)(x) = x$ and that $f_\sharp(g_\sharp \rho) = (f \circ g)_\sharp \rho$. Hence, $\mathrm{proj}_{i\sharp} T_\sharp \mu = \mu$. Further,

$$\int_{\Omega \times \Omega} d^p \, \mathrm{d}\pi = \int_{\Omega \times \Omega} d^p \, \mathrm{d}(T_\sharp \mu) = \int_{\Omega} d^p \circ T \, \mathrm{d}\mu = 0 \, .$$

- Let $W_p(\mu, \nu) = 0$. Then there must be some $\pi \in \Pi(\mu, \nu)$ with $\int_{\Omega \times \Omega} d(x, y)^p \, \mathrm{d}\pi(x, y) = 0$, which implies $d(x, y) = 0$ π-a.e., i.e., $x = y$ π-a.e.. So for $\phi \in C(\Omega)$

$$\int_{\Omega \times \Omega} \phi(x) \, \mathrm{d}\pi(x, y) = \int_{\Omega \times \Omega} \phi(y) \, \mathrm{d}\pi(x, y)$$

and thus $\mathrm{proj}_{0\sharp} \pi = \mathrm{proj}_{1\sharp} \pi$ which implies $\mu = \nu$.

- Towards triangle inequality: Let $\mu, \nu, \rho \in \mathcal{P}(\Omega)$, let π_{01}, π_{12} be optimal couplings for $W_p(\mu, \nu)$ and $W_p(\nu, \rho)$. Let $(\gamma_{01,y})_{y \in \Omega}$ be the disintegration of π_{01} with respect to proj_1. That is, for any $\phi \in C(\Omega \times \Omega)$ have

$$\int_{\Omega \times \Omega} \phi(x, y) \, \mathrm{d}\pi_{01}(x, y) = \int_{\Omega} \left(\int_{\Omega} \phi(x, y) \, \mathrm{d}\gamma_{01,y}(x) \right) \mathrm{d}\nu(y) \, .$$

Similarly, let $(\gamma_{12,y})_{y \in \Omega}$ be the disintegration of π_{12} with respect to proj_0.

- Define a new measure $\pi \in \mathcal{P}(\Omega \times \Omega)$ via

$$\int_{\Omega\times\Omega} \phi(x,z)\,\mathrm{d}\pi(x,z) = \int_\Omega \left(\int_{\Omega\times\Omega} \phi(x,z)\,\mathrm{d}\gamma_{01,y}(x)\,\mathrm{d}\gamma_{12,y}(z) \right) \mathrm{d}\nu(y)\,.$$

Sketch: Some intuition for π.

- Claim: $\pi \in \Pi(\mu,\rho)$. For $\phi \in C(\Omega)$ get

$$\begin{aligned}\int_{\Omega\times\Omega} \phi(x)\,\mathrm{d}\pi(x,z) &= \int_\Omega \left(\int_{\Omega\times\Omega} \phi(x)\,\mathrm{d}\gamma_{01,y}(x)\,\mathrm{d}\gamma_{12,y}(z) \right) \mathrm{d}\nu(y) \\ &= \int_\Omega \left(\int_\Omega \phi(x)\,\mathrm{d}\gamma_{01,y}(x) \right) \mathrm{d}\nu(y) = \int_{\Omega\times\Omega} \phi(x)\,\mathrm{d}\pi_{01}(x,y) = \int_\Omega \phi\,\mathrm{d}\mu\end{aligned}$$

- Triangle inequality:

$$\begin{aligned}W_p(\mu,\rho) &\leq \left(\int_{\Omega\times\Omega} d(x,z)^p\,\mathrm{d}\pi(x,z) \right)^{1/p} = \left(\int_\Omega \left(\int_{\Omega\times\Omega} d(x,z)^p\,\mathrm{d}\gamma_{01,y}(x)\,\mathrm{d}\gamma_{12,y}(z) \right) \mathrm{d}\nu(y) \right)^{1/p} \\ &\leq \left(\int_\Omega \left(\int_{\Omega\times\Omega} \big(d(x,y)+d(y,z)\big)^p\,\mathrm{d}\gamma_{01,y}(x)\,\mathrm{d}\gamma_{12,y}(z) \right) \mathrm{d}\nu(y) \right)^{1/p} \\ &\overset{\text{Minkowski ineq.}}{\leq} \left(\int_\Omega \left(\int_{\Omega\times\Omega} d(x,y)^p\,\mathrm{d}\gamma_{01,y}(x)\,\mathrm{d}\gamma_{12,y}(z) \right) \mathrm{d}\nu(y) \right)^{1/p} \\ &\qquad + \left(\int_\Omega \left(\int_{\Omega\times\Omega} d(y,z)^p\,\mathrm{d}\gamma_{01,y}(x)\,\mathrm{d}\gamma_{12,y}(z) \right) \mathrm{d}\nu(y) \right)^{1/p} \\ &= \left(\int_{\Omega\times\Omega} d(x,y)^p\,\mathrm{d}\pi_{01}(x,y) \right)^{1/p} + \left(\int_{\Omega\times\Omega} d(y,z)^p\,\mathrm{d}\pi_{12}(x,y) \right)^{1/p} \\ &= W_p(\mu,\nu) + W_p(\nu,\rho)\,.\end{aligned}$$

□

Theorem 3.6 (W_p metrizes weak$*$ convergence)**.** Let (Ω, d) be a compact metric space. W_p metrizes the weak$*$ convergence on $\mathcal{P}(\Omega)$. That is, for a sequence $(\mu_n)_n$ and some μ in $\mathcal{P}(\Omega)$ one has:

$$[W_p(\mu_n,\mu) \to 0] \qquad \Leftrightarrow \qquad [\mu_n \overset{*}{\rightharpoonup} \mu]$$

Proof.
- $\Rightarrow$: assume $W_p(\mu_n,\mu) \to 0$. Let $(\pi_n)_n$ be a corresponding sequence of optimal transport plans. Let $\tilde{\mu}$ be a cluster point of $(\mu_n)_n$ and let $\pi \in \Pi(\tilde{\mu},\mu)$ a corresponding cluster point of $(\pi_n)_n$. As before, denote the converging subsequence also by $(\pi_n)_n$. One has:

$$W_p(\tilde{\mu},\mu) \leq \lim_{n\to\infty} \int_{\Omega\times\Omega} d^p\,\mathrm{d}\pi_n = \lim_{n\to\infty} W_p(\mu_n,\mu) = 0$$

  Since W_p is a metric, $\tilde{\mu} = \mu$. Hence, $\mu_n \overset{*}{\rightharpoonup} \mu$.

- $\Leftarrow$: assume $\mu_n \overset{*}{\rightharpoonup} \mu$. Let π_n be optimal plans for $W_p(\mu_n, \mu)$. Extract a converging subsequence, again denoted by $(\pi_n)_n$. By Proposition 2.12 (stability of optimal plans) any cluster point π of $(\pi_n)_n$ is an optimal coupling for $W_p(\mu, \mu)$. So:

$$0 = W_p(\mu, \mu) = \int_{\Omega\times\Omega} d^p \, \mathrm{d}\pi = \lim_{n\to\infty} \int_{\Omega\times\Omega} d^p \, \mathrm{d}\pi_n = \lim_{n\to\infty} W_p(\mu_n, \mu)$$

□

3.2 Displacement interpolation

An intriguing property of the Wasserstein space $(\mathcal{P}(\Omega), W_p)$ is that it is a length space if (Ω, d) is a length space.

Definition 3.7 (Length space)**.** A metric space (Ω, d) is a length space if for every pair $(x, y) \in \Omega$ there is a continuous map $\gamma_{x,y} \in C([0,1], \Omega)$ with

$$\gamma_{x,y}(0) = x, \qquad \gamma_{x,y}(1) = 1, \qquad d\big(\gamma_{x,y}(s), \gamma_{x,y}(t)\big) = d(x,y) \cdot |s - t|$$

for s, $t \in [0, 1]$.

Theorem 3.8. If (Ω, d) is a length space and the map $(x, y) \mapsto \gamma_{x,y}$ that takes start and endpoint to a shortest path between them is measurable then $(\mathcal{P}(\Omega), W_p)$ is a length space.

Comment: Sufficient conditions for the measurability of $(x, y) \mapsto \gamma_{x,y}$ can be found for instance in [Villani, 2009, Proposition 7.16].

Proof. • Let $(\gamma_{x,y})_{(x,y)\in\Omega^2}$ be the family of maps for (Ω, d) as given by Definition 3.7. For fixed $s, t \in [0, 1]$ let

$$\begin{aligned} &\Gamma_s : \Omega \times \Omega \to \Omega, && (x, y) \mapsto \gamma_{x,y}(s), \\ &\Gamma_{s,t} : \Omega \times \Omega \to \Omega \times \Omega, && (x, y) \mapsto (\gamma_{x,y}(s), \gamma_{x,y}(t)). \end{aligned}$$

Comment: Between γ and Γ the roles of 'index' and 'arguments' of the functions are exchanged. This is formally helpful to use the push-forward of Γ.

- For given $\mu, \nu \in \mathcal{P}(\Omega)$ let π be an optimal coupling for $W_p(\mu, \nu)$. Denote $\rho_s = \Gamma_{s\sharp}\pi$.

Sketch: Interpretation of ρ_s.

- Claim: $s \mapsto \rho_s$ is a geodesic in $(\mathcal{P}(\Omega), W_p)$ between μ and ν. A 'length space map' for $(\mathcal{P}(\Omega), W_p)$ between μ and ν, $\gamma_{\mu,\nu} : [0, 1] \to \mathcal{P}(\Omega)$ is given by $\gamma_{\mu,\nu}(s) = \rho_s$. We will now show this.

- Measurability of Γ_s: By assumption $S : (x, y) \mapsto \gamma_{x,y}$ is measurable. For fixed $t \in [0, 1]$ the map $e_t : C([0, 1], \Omega) \to \Omega$, $\gamma \mapsto \gamma(t)$ is continuous and thus measurable. We find $\Gamma_s = e_s \circ S$. Similarly, $\Gamma_{s,t} = (\Gamma_s, \Gamma_t) = (e_s, e_t) \circ S$ is measurable.

- Claim: $\Gamma_{s,t\sharp}\pi \in \Pi(\rho_s, \rho_t)$.

$$\mathrm{proj}_{0\sharp}\Gamma_{s,t\sharp}\pi = (\mathrm{proj}_0 \circ \Gamma_{s,t})_\sharp\pi = \Gamma_{s\sharp}\pi = \rho_s$$

- Claim: $W_p(\rho_s,\rho_t) = |s-t| \cdot W_p(\mu,\nu)$.

$$\begin{aligned}
W_p(\rho_s,\rho_t)^p &\leq \int_{\Omega\times\Omega} d(x,y)^p \,\mathrm{d}(\Gamma_{s,t\,\sharp}\pi)(x,y) = \int_{\Omega\times\Omega} ((d\circ\Gamma_{s,t})(x,y))^p \,\mathrm{d}\pi(x,y) \\
&= \int_{\Omega\times\Omega} (d(\gamma_{x,y}(s),\gamma_{x,y}(t)))^p \,\mathrm{d}\pi(x,y) = |s-t|^p \int_{\Omega\times\Omega} (d(x,y))^p \,\mathrm{d}\pi(x,y) \\
W_p(\rho_s,\rho_t) &\leq |s-t| \cdot W_p(\mu,\nu)
\end{aligned}$$

So for $0 \leq s \leq t \leq 1$ have

$$W_p(\mu,\rho_s) \leq s \cdot W_p(\mu,\nu), \quad W_p(\rho_s,\rho_t) \leq (t-s)\cdot W_p(\mu,\nu), \quad W_p(\rho_t,\nu) \leq (1-t)\cdot W_p(\mu,\nu)$$

So

$$W_p(\mu,\rho_s) + W_p(\rho_s,\rho_t) + W_p(\rho_t,\nu) \leq W_p(\mu,\nu)$$

and by the triangle inequality

$$W_p(\mu,\rho_s) + W_p(\rho_s,\rho_t) + W_p(\rho_t,\nu) \geq W_p(\mu,\nu)\,.$$

Hence we must have equality and in particular $W_p(\rho_s,\rho_t) = |s-t| \cdot W_p(\mu,\nu)$.

□

4 Wasserstein-1 space

In this section we study in more detail the structure of the Wasserstein distance for $p=1$. Short summary: the dual problem allows a particular simplification, only possible for $p=1$, which then can be reformulated as min cost flow problem (also known as Beckmann's problem). This is particularly suitable for numerical optimization.

4.1 Duality: Kantorovich–Rubinstein formula

Definition 4.1 (1-Lipschitz functions)**.** A function $\alpha : \Omega \to \mathbb{R}$ is called 1-Lipschitz if for any pair $x, y \in \Omega$ one finds

$$|\alpha(x) - \alpha(y)| \leq d(x,y)\,.$$

The set of 1-Lipschitz functions on Ω is denoted by $\mathrm{Lip}(\Omega)$ and evidently $\mathrm{Lip}(\Omega) \subset C(\Omega)$.

Theorem 4.2 (Kantorovich–Rubinstein formula)**.**

$$W_1(\mu,\nu) = \sup\left\{ \int_\Omega \alpha \,\mathrm{d}(\mu-\nu) \middle| \alpha \in \mathrm{Lip}(\Omega) \right\}$$

Proof.
- By Theorem 2.5 (dual Kantorovich problem) we find

$$W_1(\mu,\nu) = \sup\left\{ \int_\Omega \alpha\,\mathrm{d}\mu + \int_\Omega \beta\,\mathrm{d}\nu \middle| \alpha,\, \beta \in C(\Omega), \alpha(x)+\beta(y) \leq d(x,y) \text{ for all } (x,y)\in\Omega^2 \right\}$$

and with Theorem 2.9 (existence of dual optimizers) we can restrict the optimization in the dual problem to pairs (α,β) with $\alpha = \beta^c$, $\beta = \alpha^c$ (note that by symmetry of c we need not distinguish between c- and $\bar{c}$-transform).

- For any potentially optimal pair (α, β) we know:

$$\alpha(x) = \inf \{d(x, y) - \beta(y) | y \in \Omega\} \tag{3}$$

  Since d, β continuous, Ω compact: inf is min. For fixed x let y be minimizer. For different x':

$$\alpha(x') \leq d(x', y) - \beta(y) \leq d(x', x) + d(x, y) - \beta(y) = d(x', x) + \alpha(x)$$

  where we have used the triangle inequality $d(x', y) \leq d(x', x) + d(x, y)$. This holds for all pairs $(x, x') \in \Omega^2$. So $\alpha \in \mathrm{Lip}(\Omega)$ and likewise $\beta \in \mathrm{Lip}(\Omega)$.

- Set $y = x$ in (3) to get $\alpha(x) \leq -\beta(x)$. Use that β is 1-Lipschitz to get $\beta(y) \leq \beta(x) + d(x, y)$ to get

$$\alpha(x) = \inf \{d(x, y) - \beta(y) | y \in \Omega\} \geq \inf \{d(x, y) - \beta(x) - d(x, y) | y \in \Omega\} = -\beta(x).$$

  So for optimal pair have $\alpha = -\beta$.

- So in dual Kantorovich problem can restrict optimization to 1-Lipschitz α and β where $\beta = -\alpha$.

□

Sketch: Interpretation with μ, ν being Dirac measures.

Remark 4.3. From the proof we learn: For $c(x, y) = d(x, y)$ the set of $(c, \overline{c})$-concave functions are 1-Lipschitz functions and if $\alpha \in \mathrm{Lip}(\Omega)$ then $\alpha^c = -\alpha$.

An advantage of the Kantorovich–Rubinstein formula is that $\alpha \in \mathrm{Lip}(\Omega)$ can often be turned into a local constraint (pointwise, or on small neighbourhoods) and thus, the large set of constraints of the original dual Kantorovich problem ($\alpha(x) + \beta(y) \leq d(x, y)$ for all $(x, y) \in \Omega^2$) can be replaced by a smaller number of constraints.

4.2 Beckmann's problem

Throughout this subsection Ω is compact and the closure of an open set of $\mathbb{R}^d$. d is the metric induced by shortest paths in Ω, according to the Euclidean length of paths.
Then the Lipschitz constraint can be approximated by a local constraint on the gradient of the Kantorovich potential α. Going then back to a primal formulation one finds a transport problem where movement of mass is described by a vector field $\omega \in \mathcal{M}(\Omega)^2$ that describes how mass flows from μ to ν.

Definition 4.4 (Weak divergence)**.** A measure $\omega \in \mathcal{M}(\Omega)^2$ is said to have weak divergence $\rho \in \mathcal{M}(\Omega)$, we write $\operatorname{div} \omega = \rho$, if for every $\phi \in C^1(\Omega)$ one has

$$\int_\Omega (\nabla \phi) \cdot \mathrm{d}\omega + \int_\Omega \phi \, \mathrm{d}\rho = 0 \, .$$

Comment: If ω is a smooth vector field can apply integration by parts on first term to get standard definition of divergence.

Theorem 4.5 (Beckmann's problem)**.**

$$W_1(\mu,\nu) = \min\left\{\|\omega\|_{\mathcal{M}(\Omega)^2} \middle| \omega \in \mathcal{M}(\Omega)^2,\ \operatorname{div}\omega = \nu - \mu\right\}$$

Comment: Interpretation: ω is flow field that takes μ to ν. At each point, ω gives orientation and amount of flow.

Comment: After discretization with $|\Omega| = N$ discrete points, the standard Kantorovich dual problem needs $\mathcal{O}(N)$ variables and $\mathcal{O}(N^2)$ constraints. Beckmann's problem only requires $\mathcal{O}(N)$ variables and constraints.

Proof.

- Start with Kantorovich–Rubinstein formula, Theorem 4.2:

$$W_1(\mu,\nu) = \sup\left\{\int_\Omega \alpha\,\mathrm{d}(\mu-\nu) \middle| \alpha \in \mathrm{Lip}(\Omega)\right\}$$

- $C^1(\Omega) \cap \mathrm{Lip}(\Omega)$ is dense in $\mathrm{Lip}(\Omega)$ for the sup-norm.

Comment: The proof is intuitive but somewhat lengthy. It can be found, e.g., in [Schmitzer and Wirth, 2017].

- Since $\alpha \mapsto \int_\Omega \alpha\,\mathrm{d}(\mu-\nu)$ is continuous, by density can restrict optimization of α to $C^1(\Omega) \cap \mathrm{Lip}(\Omega)$.
- Find

$$C^1(\Omega) \cap \mathrm{Lip}(\Omega) = \{\alpha \in C^1(\Omega) \,|\, \|\nabla\alpha(x)\| \le 1 \text{ for all } x \in \Omega\} = \{\alpha \in C^1(\Omega) \,|\, \|\nabla\alpha\|_{C(\Omega)^2} \le 1\}$$

- Get:

$$W_1(\mu,\nu) = \sup\left\{-f(\alpha) - g(\nabla\alpha) \middle| \alpha \in C^1(\Omega)\right\}$$

with

$$f(\alpha) = \int_\Omega \alpha\,\mathrm{d}(\nu-\mu),$$
$$g(v) = \begin{cases} 0 & \text{if } \|v\|_{C(\Omega)^2} \le 1 \\ \infty & \text{else.} \end{cases}$$

- Now apply Fenchel-Rockafellar. f, g are convex, lower semi-continuous (f is linear, g is indicator of unit ball of sup-norm). $X = C^1(\Omega)$, $Y = C(\Omega)^2$, $A = \nabla$. At $\alpha : x \mapsto 0$ have $f(\alpha) < \infty$ and $g(\nabla\alpha) < \infty$ and g is continuous at $\nabla\alpha$.
- Dual problem:

$$W_1(\mu,\nu) = \min\left\{f^*(-\nabla^*\omega) + g^*(\omega) \middle| \omega \in \mathcal{M}(\Omega)^2\right\}$$

- Compute conjugates:

$$
\begin{aligned}
f^*(-\nabla^*\omega) &= \sup\left\{ \langle \alpha, -\nabla^*\omega\rangle_{C^1(\Omega),C^1(\Omega)^*} - f(\alpha) \middle| \alpha \in C^1(\Omega) \right\} \\
&= \sup\left\{ \langle -\nabla\alpha, \omega\rangle_{C(\Omega)^2,\mathcal{M}(\Omega)^2} + \int_\Omega \alpha \, \mathrm{d}(\mu-\nu) \middle| \alpha \in C^1(\Omega) \right\} \\
&= \begin{cases} 0 & \text{if } \operatorname{div}\omega = \nu - \mu \\ +\infty & \text{else.} \end{cases} \\
g^*(\omega) &= \sup\left\{ \langle v, \omega\rangle_{C(\Omega)^2,\mathcal{M}(\Omega)^2} \middle| v \in C(\Omega)^2, \|v\|_{C(\Omega)^2} \le 1 \right\}
\end{aligned}
$$

So g^* is operator-norm of ω (interpreted as linear map on $C(\Omega)^2$). By duality:

$$g^*(\omega) = \|\omega\|_{\mathcal{M}(\Omega)^2}$$

□

Remark 4.6 (Min-cost flow on discrete graph)**.** Assume Ω is a finite, discrete vertex set of a graph with edge set $E \subset \Omega \times \Omega$ and edge lengths $\ell : E \to (0,\infty]$. Let d be metric induced by shortest paths in (Ω, E, ℓ). Then

$$[\alpha \in \mathrm{Lip}(\Omega)] \Leftrightarrow [|\alpha(x) - \alpha(y)| \le \ell(x,y) \text{ for all } (x,y) \in E] \,.$$

In many applications $|E| \ll |\Omega|$ and thus the right condition is easier to check than the left. Can define discrete gradient operator $G : C(\Omega) \to C(E)$, $(G\alpha)(x,y) = (\alpha(x) - \alpha(y))/\ell(x,y)$ and replace Lipschitz constraint by local constraint on gradient, as above. This leads to discrete min-cost flow problem on metric graph. See for instance [Bertsekas and Eckstein, 1988].

5 The Benamou–Brenier formula for W_2

Similar as in Section 4.2: throughout this section Ω is convex, compact and the closure of an open set of $\mathbb{R}^d$. d is the metric induced by shortest paths in Ω, according to the Euclidean length of paths. Due to convexity of Ω, shortest paths are given by straight lines.

5.1 Intuition

In Section 3.2 we have discussed the displacement interpolation: for a given optimal coupling π, mass can be thought of as traveling along geodesics with constant speed. In In Section 4.2 we have given Beckmann's problem as equivalent formulation for W_1: mass transport is described by a (time-independent) flow field ω. Now we study the Benamou-Brenier formulation for W_2 on $\mathbb{R}^d$. Intuitively it is the analogy of Beckmann's formulation for W_2. But here we need a time-dependent flow field.

We start with an intuitive and informal discussion. In the Benamou–Brenier formulation, mass transport between μ and $\nu \in \mathcal{P}(\Omega)$ is described as directly in the displacement interpolation: we are looking for a function $[0,1] \ni t \mapsto \rho_t \in \mathcal{P}(\Omega)$ such that ρ_t is the displacement interpolation between μ and ν. So we want in particular:

$$\rho_0 = \mu, \qquad\qquad \rho_1 = \nu \tag{4a}$$

The temporal change of ρ_t will be described by a time dependent vector field $v_t : \Omega \to \mathbb{R}^d$ where $v_t(x)$ gives the velocity of a mass particle at time t and position x. ρ_t and v_t are connected by the 'continuity equation':

$$\partial_t \rho_t + \operatorname{div}(\rho_t\, v_t) = 0 \tag{4b}$$

Remark 5.1 (Interpretation of continuity equation)**.** Assume ρ_t is Lebesgue-absolutely continuous at all times. By abuse of notation write ρ_t for its Lebesgue density. Let $A \subset \Omega$ be an open set with a smooth surface ∂A. Then formally from the continuity equation we obtain with the divergence theorem:

$$\frac{\mathrm{d}}{\mathrm{d}t}\left(\int_A \rho_t\, \mathrm{d}x\right) = -\int_A \operatorname{div}(\rho_t\, v_t)\, \mathrm{d}x = \int_{\partial A} \rho_t\, v_t \cdot n(s)\mathrm{d}s$$

where n is the outward pointing unit-normal at $s \in \partial A$ and $\mathrm{d}s$ is the surface volume measure of ∂A.
That is, the change of mass in the volume A is related to the flow of $\rho_t\, v_t$ across the surface ∂A.

It will turn out that a pair (ρ_t, v_t) describes the displacement interpolation between μ and ν if it solves the continuity equation with temporal boundary conditions, (4), and minimizes the following functional:

$$\int_{[0,1]} \int_\Omega |v_t(x)|^2\, \mathrm{d}\rho_t(x)\, \mathrm{d}t \tag{5}$$

It is often called the 'action' functional. This can be interpreted as (twice) the integral of the kinetic energy of all particles over the time-interval $[0,1]$ (recall the iconic formula $E_{\text{kinetic}} = \frac{1}{2} m\, v^2$). So the displacement interpolation is the trajectory between μ and ν with the least average kinetic energy of the particles.
This formulation has several problems: when ρ_t is not Lebesgue-absolutely continuous, but contains concentrated Dirac masses, it is difficult to define the divergence condition (4b) as above. Moreover, the functional (5) is not jointly convex in ρ_t and v_t.
In the next sections we will provide a rigorous definition of the Benamou–Brenier formulation and sketch the proof for equivalence with the Kantorovich definition of W_2.

5.2 Rigorous definition

Now we give a rigorous definition of the Benamou–Brenier formulation. ρ will be defined as a measure on $[0,1] \times \Omega$. The velocity field v will be replaced by a momentum field $\omega \in \mathcal{M}_+([0,1] \times \Omega)^d$ where intuitively $\omega(t,x)$ represents $\rho_t(x)\, v_t(x)$. The decomposition of ρ and ω into time-slices will be discussed in Proposition 5.7.

Definition 5.2 (Weak continuity equation)**.** Let μ, $\nu \in \mathcal{P}(\Omega)$. A pair $(\rho, \omega) \in \mathcal{M}([0,1] \times \Omega) \times \mathcal{M}([0,1] \times \Omega)^d$ is said to solve the weak continuity equation with temporal boundary conditions μ and ν if

$$\int_{[0,1]\times\Omega} (\partial_t \phi)\, \mathrm{d}\rho + \int_{[0,1]\times\Omega} \nabla\phi \cdot \mathrm{d}\omega = \int_\Omega \phi(1,\cdot)\, \mathrm{d}\nu - \int_\Omega \phi(0,\cdot)\, \mathrm{d}\mu$$

for all $\phi \in C^1([0,1] \times \Omega)$. We denote the set of solutions by $\mathcal{CE}(\mu, \nu)$.

Remark 5.3. $\mathcal{CE}(\mu,\nu)$ is an affine set. When (ρ_1,ω_1), (ρ_2,ω_2) are two solutions, all points on the entire line spanned by them are solutions as well. This is, due to the left hand side being linear in (ρ,ω). In particular $\mathcal{CE}(\mu,\nu)$ is convex.
Moreover, since $\partial_t\phi \in C([0,1]\times\Omega)$ and $\nabla\phi \in C([0,1]\times\Omega)^d$, $\mathcal{CE}(\mu,\nu)$ is weak$*$-closed.

Remark 5.4 (1-homogeneous functions and integration)**.** A function $f:\mathbb{R}^n\to\mathbb{R}$ is 1-homogeneous if

$$f(\lambda\cdot x)=\lambda\cdot f(x)$$

for all $\lambda\geq 0$, $x\in\mathbb{R}^n$ (with the convention $0\cdot\infty=0$). Let $\mu\in\mathcal{M}(\Omega)^n$ and let $\sigma\in\mathcal{M}_+(\Omega)$ such that $\mu\ll\sigma$. Then can define following integral:

$$\int_\Omega f\left(\tfrac{\mathrm{d}\mu}{\mathrm{d}\sigma}\right)\,\mathrm{d}\sigma$$

It is important to note that due to 1-homogeneity of f the integral does not depend on the choice of σ as long as $\mu\ll\sigma$. Let $\tau\in\mathcal{M}_+(\Omega)$ be another measure such that $\mu\ll\tau$ and let $\nu\in\mathcal{M}_+(\Omega)$ be such that $\sigma,\,\tau\ll\nu$. One finds for any $\phi\in C(\Omega)^n$

$$\int_\Omega \phi\cdot\mathrm{d}\mu=\int_\Omega\phi\cdot\tfrac{\mathrm{d}\mu}{\mathrm{d}\sigma}\mathrm{d}\sigma=\int_\Omega\phi\cdot\tfrac{\mathrm{d}\mu}{\mathrm{d}\sigma}\tfrac{\mathrm{d}\sigma}{\mathrm{d}\nu}\mathrm{d}\nu\,.$$

Therefore,

$$\tfrac{\mathrm{d}\mu}{\mathrm{d}\nu}=\tfrac{\mathrm{d}\mu}{\mathrm{d}\sigma}\tfrac{\mathrm{d}\sigma}{\mathrm{d}\nu}=\tfrac{\mathrm{d}\mu}{\mathrm{d}\tau}\tfrac{\mathrm{d}\tau}{\mathrm{d}\nu}$$

and consequently:

$$\int_\Omega f\left(\tfrac{\mathrm{d}\mu}{\mathrm{d}\sigma}\right)\mathrm{d}\sigma=\int_\Omega f\left(\tfrac{\mathrm{d}\mu}{\mathrm{d}\sigma}\right)\tfrac{\mathrm{d}\sigma}{\mathrm{d}\nu}\mathrm{d}\nu=\int_\Omega f\left(\tfrac{\mathrm{d}\mu}{\mathrm{d}\sigma}\tfrac{\mathrm{d}\sigma}{\mathrm{d}\nu}\right)\mathrm{d}\nu=\int_\Omega f\left(\tfrac{\mathrm{d}\mu}{\mathrm{d}\tau}\tfrac{\mathrm{d}\tau}{\mathrm{d}\nu}\right)\mathrm{d}\nu=\int_\Omega f\left(\tfrac{\mathrm{d}\mu}{\mathrm{d}\tau}\right)\mathrm{d}\tau\,.$$

Definition 5.5 (Action functional)**.** Let

$$\Phi:\mathbb{R}\times\mathbb{R}^d\to\mathbb{R}\cup\{\infty\},\qquad\qquad (\rho,\omega)\mapsto\begin{cases}\frac{\|\omega\|^2}{\rho} & \text{if } \rho>0,\\ 0 & \text{if } (\rho,\omega)=(0,0),\\ +\infty & \text{else.}\end{cases}$$

Comment: If $(\rho,\omega)=(\rho,\rho\,v)$, $\rho\geq 0$ then $\Phi(\rho,\omega)=\|v\|^2\,\rho$. So this coincides with informal definition of action.

Φ is jointly convex and lower semi-continuous since sub-level sets where $\Phi(\rho,\omega)\leq C$ are given by $\rho\geq C^{-1}\,\|\omega\|^2$. Moreover, Φ is jointly 1-homogeneous.

Sketch: Φ and its sub-level sets.

The action is then given by

$$\mathcal{A}(\rho,\omega)=\int_{[0,1]\times\Omega}\Phi\left(\tfrac{\mathrm{d}(\rho,\omega)}{\mathrm{d}\sigma}\right)\mathrm{d}\sigma(t,x)$$

where σ is any measure in $\mathcal{M}_+([0,1]\times\Omega)$ such that $(\rho,\omega)\ll\sigma$. Due to Remark 5.4 the definition does not depend on the choice of σ.

Now we have gathered the ingredients for the rigorous definition of the Benamou–Brenier formulation.

Definition 5.6 (Benamou–Brenier formulation)**.**

$$W_{\mathrm{BB}}(\mu,\nu)^2 := \inf\{\mathcal{A}(\rho,\omega)|(\rho,\omega)\in\mathcal{CE}(\mu,\nu)\} \tag{6}$$

In the following we use the 'time projection' map:

$$T:[0,1]\times\Omega\to[0,1], \qquad\qquad (t,x)\mapsto t\,. \tag{7}$$

So far, ρ and ω were only defined as joint measures in time and space. For potential minimizers of the Benamou–Brenier formulation we show that they can be decomposed into 'time-slices' such that for Lebesgue-almost every time t the spatial arrangement of mass is fixed.

Proposition 5.7 (Time-disintegration of ρ)**.** Let $(\rho,\omega)\in\mathcal{CE}(\mu,\nu)$, $\mathcal{A}(\rho,\omega)<\infty$. Then

- $\rho\in\mathcal{M}_+([0,1]\times\Omega)$, $\|\rho\|_{\mathcal{M}([0,1]\times\Omega)}=1$,
- $T_\sharp\rho=\mathcal{L}_{[0,1]}$ where the latter is the Lebesgue measure on $[0,1]$. There is a Lebesgue-a.e. unique family of measures $(\rho_t)_{t\in[0,1]}$ with $\rho_t\in\mathcal{P}(\Omega)$ such that for any $\phi\in C([0,1]\times\Omega)$
$$\int_{[0,1]\times\Omega}\phi(t,x)\,\mathrm{d}\rho(t,x)=\int_{[0,1]}\left[\int_\Omega\phi(t,x)\,\mathrm{d}\rho_t(x)\right]\mathrm{d}t\,.$$
- $\omega\ll\rho$. (This implies that ω can also be decomposed into time-slices.)

Proof.

- If $\rho\notin\mathcal{M}_+([0,1]\times\Omega)$ then $\frac{\mathrm{d}\rho}{\mathrm{d}\sigma}<0$ for a set that is not σ-negligible, for every $\sigma\in\mathcal{M}_+([0,1]\times\Omega)$ with $\rho\ll\sigma$ and thus $\mathcal{A}(\rho,\omega)=\infty$.
- Use the test function $\phi(t,x)=t$ in the definition of $\mathcal{CE}(\mu,\nu)$ to get
$$\int_{[0,1]\times\Omega}\mathrm{d}\rho(t,x)=\int_\Omega\mathrm{d}\nu(x)=1\,.$$
This is the Radon-norm of ρ since ρ is non-negative.
- Let $f\in C([0,1])$ and let
$$F(t):=\int_0^t f(s)\,\mathrm{d}s\,.$$
By construction $\partial_tF=f$ and $F\circ T\in C^1([0,1]\times\Omega)$. So from the continuity equation we know that
$$\int_{[0,1]}f\,\mathrm{d}(T_\sharp\rho)\int_{[0,1]\times\Omega}(\partial_tF)\circ T\,\mathrm{d}\rho=\int_\Omega(F\circ T)(1,\cdot)\,\mathrm{d}\nu-\int_\Omega(F\circ T)(0,\cdot)\,\mathrm{d}\mu=F(1)-F(0)\,.$$
So the integral against $T_\sharp\rho$ coincides with the Lebesgue measure for all test functions, hence $T_\sharp\rho=\mathcal{L}_{[0,1]}$.
- Let σ be some reference measure with $(\rho,\omega)\ll\sigma$. If $\omega\not\ll\rho$ then there must be a set $A\subset[0,1]\times\Omega$ with $\sigma(A)>0$ where $\frac{\mathrm{d}\rho}{\mathrm{d}\sigma}=0$ but $\frac{\mathrm{d}\omega}{\mathrm{d}\sigma}\neq0$ and consequently $\Phi(\frac{\mathrm{d}\rho}{\mathrm{d}\sigma},\frac{\mathrm{d}\omega}{\mathrm{d}\sigma})=\infty$ and thus $\mathcal{A}(\rho,\omega)=\infty$.

□

Next, we establish that minimizers exist.

Proposition 5.8. If $W_{\mathrm{BB}}(\mu,\nu)<\infty$ then minimizers of the Benamou–Brenier formulation (6) exist.

Finiteness of $W_{\mathrm{BB}}(\mu,\nu)$ will follow from the equivalence results of the next section.

Proof.

- Let $(\rho_n,\omega_n)_n$ be a minimizing sequence. We may assume $(\rho_n,\omega_n)\in\mathcal{CE}(\mu,\nu)$ for all n and that $\mathcal{A}(\rho_n,\omega_n)\le C$ for some $C<\infty$.
- By Proposition 5.7 $\rho_n\ge 0$, $\|\rho_n\|_{\mathcal{M}([0,1]\times\Omega)}=1$ and $\omega_n\ll\rho_n$. Therefore we can pick $\sigma=\rho_n$ as reference measure in the definition of the action $\mathcal{A}(\rho_n,\omega_n)$:
$$\mathcal{A}(\rho_n,\omega_n)=\int_{[0,1]\times\Omega}\Phi\left(\tfrac{\mathrm{d}\rho_n}{\mathrm{d}\rho_n},\tfrac{\mathrm{d}\omega_n}{\mathrm{d}\rho_n}\right)\,\mathrm{d}\rho_n=\int_{[0,1]\times\Omega}\left\|\tfrac{\mathrm{d}\omega_n}{\mathrm{d}\rho_n}\right\|^2\,\mathrm{d}\rho_n$$
With this can bound Radon-norm of ω_n:
$$\|\omega_n\|_{\mathcal{M}([0,1]\times\Omega)^d}=\int_{[0,1]\times\Omega}\left\|\tfrac{\mathrm{d}\omega_n}{\mathrm{d}\rho_n}\right\|\,\mathrm{d}\rho_n\overset{\text{Cauchy–Schwarz}}{\le}$$
$$\left(\int_{[0,1]\times\Omega}\left\|\tfrac{\mathrm{d}\omega_n}{\mathrm{d}\rho_n}\right\|^2\,\mathrm{d}\rho_n\cdot\int_{[0,1]\times\Omega}1^2\mathrm{d}\rho_n\right)^{1/2}=\mathcal{A}(\omega_n,\rho_n)^{1/2}\le C^{1/2}\,.$$
- Therefore, the sequence is uniformly bounded in norm and thus, by Banach–Alaoglu (Theorem 1.22) must have a convergent subsequence with limit (ρ,ω).
- Since $\mathcal{CE}(\mu,\nu)$ is weak$*$-closed, one has $(\rho,\omega)\in\mathcal{CE}(\mu,\nu)$. Moreover, since Φ is convex, lower semi-continuous and 1-homogeneous, the functional $\mathcal{A}$ is lower semi-continuous (see for instance [Ambrosio et al., 2000, Theorem 2.38]). Therefore, (ρ,ω) must be a minimizer of W_{BB}.

□

5.3 Equivalence with Kantorovich formulation

So far, we have given a rigorous definition of the Benamou–Brenier formulation and established that minimizers exist. Now we will sketch how to show that it is actually equivalent to the Kantorovich formulation for W_2. We will proof that $W_{\mathrm{BB}}\le W_2$. The converse inequality requires several tedious smoothing arguments. We only sketch the idea of the proof and refer to the literature for the details (e.g. [Villani, 2003, Theorem 8.1]).

Proposition 5.9.
$$W_{\mathrm{BB}}(\mu,\nu)\le W_2(\mu,\nu)$$

Sketch of proof.

- For fixed $x,\,y\in\Omega$ let
$$\gamma_{x,y}:[0,1]\to\Omega,\qquad\qquad t\mapsto(1-t)\,x+t\,y$$
parametrize the constant-speed geodesic between x and y along the straight line between them. Define measures $\rho_{x,y}\in\mathcal{M}_+([0,1]\times\Omega)$ and $\omega_{x,y}\in\mathcal{M}([0,1]\times\Omega)^d$ by
$$\int_{[0,1]\times\Omega}\phi\,\mathrm{d}\rho_{x,y}=\int_{[0,1]}\phi(t,\gamma_{x,y}(t))\,\mathrm{d}t,\quad\int_{[0,1]\times\Omega}\psi\cdot\mathrm{d}\omega_{x,y}=\int_{[0,1]}\psi(t,\gamma_{x,y}(t))\cdot(y-x)\,\mathrm{d}t$$

for $\phi \in C([0,1]\times\Omega)$, $\psi \in C([0,1]\times\Omega)^d$. Note that $\omega_{x,y} \ll \rho_{x,y}$ and that $\frac{\mathrm{d}\omega_{x,y}}{\mathrm{d}\rho_{x,y}} = y - x$.

Comment: $(\rho_{x,y}, \omega_{x,y})$ describe Dirac particle moving at constant speed from x to y.

- $(\rho_{x,y}, \omega_{x,y}) \in \mathcal{CE}(\delta_x, \delta_y)$: For $\phi \in C^1([0,1]\times\Omega)$ have

$$\begin{aligned}\int_{[0,1]\times\Omega} \partial_t\phi \,\mathrm{d}\rho_{x,y} + \int_{[0,1]\times\Omega} \nabla\phi \,\mathrm{d}\omega_{x,y} \\ = \int_{[0,1]} [(\partial_t\phi)(t,\gamma_{x,y}(t)) + (\nabla\phi)(t,\gamma_{x,y}(t))\cdot(y-x)]\,\mathrm{d}t \\ = \int_{[0,1]} \left[\frac{\mathrm{d}}{\mathrm{d}t}\phi(t,\gamma_{x,y}(t))\right]\mathrm{d}t = \phi(1,y) - \phi(0,x)\,.\end{aligned}$$

- Compute $\mathcal{A}(\rho_{x,y}, \omega_{x,y})$. Use that $\omega_{x,y} \ll \rho_{x,y}$.

$$\mathcal{A}(\rho_{x,y}, \omega_{x,y}) = \int_{[0,1]\times\Omega} \Phi\left(1, \tfrac{\mathrm{d}\omega_{x,y}}{\mathrm{d}\rho_{x,y}}\right)\mathrm{d}\rho_{x,y} = \int_{[0,1]\times\Omega} \|x-y\|^2\,\mathrm{d}\rho_{x,y} = \|x-y\|^2 = d(x,y)^2$$

- For μ, ν let π be optimal coupling. Define ρ, ω by

$$\rho = \int_{\Omega\times\Omega} \rho_{x,y}\,\mathrm{d}\pi(x,y), \qquad \omega = \int_{\Omega\times\Omega} \omega_{x,y}\,\mathrm{d}\pi(x,y).$$

Comment: The above is called a 'Pettis integral'. Think of it as superposition of $\rho_{x,y}$ with weights $\pi(x,y)$.

- One can quickly verify that $(\rho,\omega) \in \mathcal{CE}(\mu,\nu)$. By using that $\mathcal{A}$ is convex and 1-homogeneous and a generalization of Jensen's inequality to topological vector spaces and Pettis integrals, [Perlman, 1974], one finds:

$$\begin{aligned}W_{\mathrm{BB}}(\mu,\nu)^2 \le \mathcal{A}(\rho,\omega) = \mathcal{A}\left(\int_{\Omega\times\Omega} \rho_{x,y}\,\mathrm{d}\pi(x,y), \int_{\Omega\times\Omega} \omega_{x,y}\,\mathrm{d}\pi(x,y)\right) \\ \le \int_{\Omega\times\Omega} \mathcal{A}(\rho_{x,y},\omega_{x,y})\,\mathrm{d}\pi(x,y) = \int_{\Omega\times\Omega} d(x,y)^2\,\mathrm{d}\pi(x,y) = W_2(\mu,\nu)^2\end{aligned}$$

- One can avoid Pettis integrals and the generalization of Jensen's inequality as follows: approximate μ and ν by sequences of superpositions of finite numbers of Dirac measures. For each step, the optimal coupling will be a finite number of Dirac measures for which standard finite sub-additivity can be used. Then go to the limit and use that W_2 is weak$*$ continuous since it metrizes the weak$*$ topology.

$\square$

Now we sketch how to proof the converse inequality.

- Let $v \in C^1([0,1]\times\Omega)^d$ be a differentiable vector field, let $\varphi : [0,1]\times\Omega \to \Omega$ be its flow, i.e. φ is defined by

$$\partial_t\varphi(t,x) = v(t,\varphi(t,x)), \qquad \varphi(0,x) = x.$$

- Set $\rho_t = \varphi_{t\sharp}\mu$. Note that $\rho_0 = \mu$. Assume that v is a flow field that takes the mass μ to ν. That means $\rho_1 = \nu$.
- Define measures ρ and ω by

$$\int_{[0,1]\times\Omega} \phi \,\mathrm{d}\rho = \int_{[0,1]} \left[\int_\Omega \phi(t,x)\,\mathrm{d}\rho_t(x)\right]\mathrm{d}t, \quad \int_{[0,1]\times\Omega} \psi \cdot \mathrm{d}\omega = \int_{[0,1]} \left[\int_\Omega \psi(t,x)\cdot v(t,x)\,\mathrm{d}\rho_t(x)\right]\mathrm{d}t.$$

  for $\phi \in C([0,1]\times\Omega)$ and $\psi \in C([0,1]\times\Omega)^d$.
- Note that $(\rho,\omega) \in \mathcal{CE}(\mu,\nu)$: For $\phi \in C^1([0,1]\times\Omega)$ have

$$\begin{aligned}
\int_{[0,1]\times\Omega} \partial_t\phi\,\mathrm{d}\rho &= \int_{[0,1]\times\Omega} (\partial_t\phi)(t,\cdot)\circ\varphi(t,\cdot)\,\mathrm{d}\mu\,\mathrm{d}t \\
&= \int_{[0,1]\times\Omega} \left[\frac{\mathrm{d}}{\mathrm{d}t}(\phi(t,\cdot)\circ\varphi(t,\cdot)) - ((\nabla\phi)(t,\cdot)\circ\varphi(t,\cdot)\cdot v(t,\cdot)\circ\varphi(t,\cdot)\right]\mathrm{d}\mu\,\mathrm{d}t \\
&= \int_\Omega [\phi(t,\cdot)\circ\varphi(t,\cdot))]_0^1\,\mathrm{d}\mu - \int_{[0,1]\times\Omega} (\nabla\phi(t,\cdot)\cdot v(t,\cdot))\circ\varphi(t,\cdot)\,\mathrm{d}\mu\,\mathrm{d}t \\
&= \int_\Omega \phi(1,\cdot)\,\mathrm{d}\nu - \int_\Omega \phi(0,\cdot)\,\mathrm{d}\mu - \int_{[0,1]\times\Omega} \nabla\varphi\cdot\mathrm{d}\omega\,.
\end{aligned}$$

- For the action get:

$$\begin{aligned}
\mathcal{A}(\rho,\omega) &= \int_{[0,1]\times\Omega} \Phi\left(1, \tfrac{\mathrm{d}\omega}{\mathrm{d}\rho}\right)\,\mathrm{d}\rho = \int_{[0,1]\times\Omega} \|v(t,x)\|^2\,\mathrm{d}\rho(t,x) \\
&= \int_{[0,1]\times\Omega} \|v(t,\varphi(t,x))\|^2\,\mathrm{d}\mu(x)\,\mathrm{d}t = \int_\Omega \left[\int_{[0,1]} \|(\partial_t\varphi(t,x)\|^2\mathrm{d}t\right]\mathrm{d}\mu(x) \\
&\geq \int_\Omega \underbrace{\left\|\int_{[0,1]} \partial_t\varphi(t,x)\,\mathrm{d}t\right\|^2}_{=\|\varphi(1,x)-\varphi(0,x)\|^2} \mathrm{d}\mu(x) = \int_\Omega \|\varphi(1,x)-x\|^2\mathrm{d}\mu(x)
\end{aligned}$$

  This is the Monge-cost associated with the transport map $\varphi(1,\cdot)$. And thus this is larger than $W_2(\mu,\nu)^2$.
- To extend this to a full proof for the inequality $W_{\mathrm{BB}} \geq W_2$ one needs to apply a sequence of smoothing arguments to approximate any (ρ,ω) in such a way that $v = \frac{\mathrm{d}\omega}{\mathrm{d}\rho}$ is in $C^1([0,1]\times\Omega)$.

6 Entropy regularization and Sinkhorn algorithm

Solving the optimal transport problem numerically is a considerable challenge. The Monge formulation entails a non-convex optimization problem with non-linear constraints. While the Kantorovich formulation yields convex primal and dual problems, their naive dimensionality grows quadratically with the discretization resolution of the base domain Ω. There is a diverse zoo of fundamentally very different numerical approaches that

- apply to different types of base spaces and cost functions,

- have different requirements on the structure and regularity of the marginals,
- and scale differently in terms of run-time and memory.

A decent compromise between flexibility, accuracy, run-time and implementation complexity is the entropy regularization approach.

6.1 Discretization

For the sake of simplicity and with numerical optimization in mind we will now turn to discretized problems. The Kantorovich formulation allows elegant discretization via discretization of the marginals. Let

$$\boldsymbol{\Omega}_0 = \{x_0, \dots, x_{M-1}\}, \qquad \boldsymbol{\Omega}_1 = \{y_0, \dots, y_{N-1}\}$$

be discrete, finite subsets of Ω with cardinalities $M = \#(\boldsymbol{\Omega}_0)$, $N = \#(\boldsymbol{\Omega}_1)$. Assume μ and ν are given by superpositions of Dirac measures located on $\boldsymbol{\Omega}_0$ and $\boldsymbol{\Omega}_1$:

$$\mu = \sum_{i=0}^{M-1} \boldsymbol{\mu}_i \, \delta_{x_i}, \qquad \nu = \sum_{i=0}^{N-1} \boldsymbol{\nu}_i \, \delta_{y_i}$$

with coefficient vectors $\boldsymbol{\mu} \in \mathbb{R}^M$, $\boldsymbol{\nu} \in \mathbb{R}^N$. Then any feasible coupling $\pi \in \Pi(\mu, \nu)$ must be concentrated on $\boldsymbol{\Omega}_0 \times \boldsymbol{\Omega}_1$ and can be written as

$$\pi = \sum_{i=0}^{M-1} \sum_{j=0}^{N-1} \boldsymbol{\pi}_{i,j} \, \delta_{(x_i, y_j)}$$

for a nonnegative matrix $\boldsymbol{\pi} \in \mathbb{R}_+^{M \times N}$. The marginal constraints ${\mathrm{proj}_0}_\sharp \pi = \mu$ and ${\mathrm{proj}_1}_\sharp \pi = \nu$ then become

$$\boldsymbol{\mu}_i = \sum_{j=0}^{N-1} \boldsymbol{\pi}_{i,j}, \qquad \boldsymbol{\nu}_j = \sum_{i=0}^{M-1} \boldsymbol{\pi}_{i,j} \, .$$

These can be written compactly as

$$\boldsymbol{\mu} = \boldsymbol{\pi} \, \mathbb{1}_N \, , \qquad \boldsymbol{\nu} = \boldsymbol{\pi}^\top \mathbb{1}_M$$

where $\mathbb{1}_N$ denotes the vector of $\mathbb{R}^N$ with all entries being 1 and $\pi \, \mathbb{1}_N$ denotes standard matrix-vector multiplication.
For a cost function $c \in C(\Omega \times \Omega)$ denote by $\boldsymbol{c} \in \mathbb{R}^{M \times N}$,

$$\boldsymbol{c}_{i,j} = c(x_i, y_j)$$

the discrete cost coefficient matrix. The linear transport cost is then given by

$$\int_{\Omega \times \Omega} c \, \mathrm{d}\pi = \langle \boldsymbol{c}, \boldsymbol{\pi} \rangle \quad \text{with} \quad \langle \boldsymbol{c}, \boldsymbol{\pi} \rangle := \sum_{i=0}^{M-1} \sum_{j=0}^{N-1} \boldsymbol{c}_{i,j} \, \boldsymbol{\pi}_{i,j} \, .$$

Finally, the discretized transport problem reads

$$\min \left\{ \langle \boldsymbol{c}, \boldsymbol{\pi} \rangle \middle| \boldsymbol{\pi} \in \mathbb{R}_+^{M \times N}, \, \boldsymbol{\pi} \, \mathbb{1}_N = \boldsymbol{\mu}, \, \boldsymbol{\pi}^\top \mathbb{1}_M = \boldsymbol{\nu} \right\} \, . \tag{8}$$

The corresponding dual problem is

$$\max\big\{\langle\boldsymbol{\alpha},\boldsymbol{\mu}\rangle+\langle\boldsymbol{\beta},\boldsymbol{\nu}\rangle\big|(\boldsymbol{\alpha},\boldsymbol{\beta})\in\mathbb{R}^M\times\mathbb{R}^N,\\ \boldsymbol{\alpha}_i+\boldsymbol{\beta}_j\leq\boldsymbol{c}_{i,j}\text{ for all }i=0,\ldots,M-1,\,j=0,\ldots,N-1\big\}\,. \tag{9}$$

Remark 6.1. A nice property of the Kantorovich formulation is that after discretization of the marginals, the exact transport problem becomes finite-dimensional and no further discretization is required. The stability result, Proposition 2.12, guarantees that the discrete optimal couplings π converge to a continuous optimal coupling as μ and ν are approximated with increasing resolution and accuracy.

6.2 Entropy regularization of the Kantorovich problem

Definition 6.2 (Kullback–Leibler divergence)**.** Let Ω be a measure space and let $\nu\in\mathcal{M}_+(\Omega)$. The Kullback–Leibler divergence (or relative entropy) with respect to ν is defined as

$$\mathrm{KL}(\cdot|\nu):\mathcal{M}(\Omega)\to\mathbb{R}\cup\{\infty\},$$
$$\mathrm{KL}(\mu|\nu)=\begin{cases}\int_\Omega\log\left(\frac{\mathrm{d}\mu}{\mathrm{d}\nu}\right)\mathrm{d}\mu-\mu(\Omega)+\nu(\Omega) & \text{if }\mu\in\mathcal{M}_+(\Omega),\,\mu\ll\nu,\\ +\infty & \text{else.}\end{cases}$$

The entropy regularized variant of the Kantorovich transport problem is defined as

$$\inf\left\{\int_{\Omega\times\Omega}c\,\mathrm{d}\pi+\varepsilon\,\mathrm{KL}(\pi|\rho)\middle|\pi\in\Pi(\mu,\nu)\right\} \tag{10}$$

where $\varepsilon>0$ is the regularization weight and $\rho\in\mathcal{M}_+(\Omega\times\Omega)$ is a reference measure on $\Omega\times\Omega$. Now let μ and ν be discrete measures, as in the previous section, supported on $\boldsymbol{\Omega}_0$ and $\boldsymbol{\Omega}_1$. For π to be feasible its support must be contained in $\boldsymbol{\Omega}_0\times\boldsymbol{\Omega}_1$. At the same time, for $\mathrm{KL}(\pi|\rho)<\infty$ it is necessary that $\pi\ll\rho$. Thus, for a reasonable discretization of (10) we must pick a suitable discrete reference measure ρ of the form

$$\rho=\sum_{i=0}^{M-1}\sum_{j=0}^{N-1}\boldsymbol{\rho}_{i,j}\,\delta_{(x_i,y_j)}$$

with non-negative coefficients $\boldsymbol{\rho}\in\mathbb{R}^{M\times N}$. For simplicity we assume that $\boldsymbol{\rho}$ is strictly positive, $\boldsymbol{\rho}_{i,j}>0$ for all (i,j).
The Kullback–Leibler divergence $\mathrm{KL}(\pi|\rho)$ can then be expressed directly in terms of $\boldsymbol{\pi}$ and $\boldsymbol{\rho}$:

$$\mathrm{KL}(\pi|\rho)=\mathbf{KL}(\boldsymbol{\pi}|\boldsymbol{\rho}):=\begin{cases}\sum_{i=0}^{M-1}\sum_{j=0}^{N-1}\left(\log\left(\frac{\boldsymbol{\pi}_{i,j}}{\boldsymbol{\rho}_{i,j}}\right)\cdot\boldsymbol{\pi}_{i,j}-\boldsymbol{\pi}_{i,j}+\boldsymbol{\rho}_{i,j}\right) & \text{if }\boldsymbol{\pi}\geq0,\boldsymbol{\pi}\ll\boldsymbol{\rho},\\ +\infty & \text{else.}\end{cases}$$

with the convention that $0\log 0=0$. Here $\boldsymbol{\pi}\ll\boldsymbol{\rho}$ means that $[\boldsymbol{\rho}_{i,j}=0]\Rightarrow[\boldsymbol{\pi}_{i,j}=0]$ which is the natural discrete analogue of absolute continuity. Note that for fixed ρ or $\boldsymbol{\rho}$, $\mathrm{KL}(\cdot|\rho)$ and $\mathbf{KL}(\cdot|\boldsymbol{\rho})$ are convex, lower-semicontinuous functions.
So after discretization of μ, ν and ρ, (10) becomes the regularized equivalent of (8):

$$\inf\left\{\langle\boldsymbol{c},\boldsymbol{\pi}\rangle+\varepsilon\,\mathbf{KL}(\boldsymbol{\pi}|\boldsymbol{\rho})\middle|\boldsymbol{\pi}\in\mathbb{R}^{M\times N},\,\boldsymbol{\pi}\,\mathbb{1}_N=\boldsymbol{\mu},\,\boldsymbol{\pi}^\top\mathbb{1}_M=\boldsymbol{\nu}\right\} \tag{11}$$

Note that the non-negativity constraint for $\boldsymbol{\pi}$ in (8) is already implied by the discrete Kullback–Leibler divergence in (11).
Now we introduce the kernel matrix $\mathbf{k} \in \mathbb{R}_+^{M \times N}$,

$$\mathbf{k}_{i,j} = \exp(-\boldsymbol{c}_{i,j}/\varepsilon) \cdot \boldsymbol{\rho}_{i,j} \, . \tag{12}$$

We find

$$\boldsymbol{c}_{i,j} \cdot \boldsymbol{\pi}_{i,j} + \varepsilon \log\left(\frac{\boldsymbol{\pi}_{i,j}}{\boldsymbol{\rho}_{i,j}}\right) = \varepsilon \log\left(\frac{\boldsymbol{\pi}_{i,j}}{\mathbf{k}_{i,j}}\right)$$

and thus

$$\langle \boldsymbol{c}, \boldsymbol{\pi} \rangle + \varepsilon \, \mathbf{KL}(\boldsymbol{\pi}|\boldsymbol{\rho}) = \varepsilon \, \mathbf{KL}(\boldsymbol{\pi}|\mathbf{k}) - \mathbf{k}(\boldsymbol{\Omega}_0 \times \boldsymbol{\Omega}_1) + \boldsymbol{\rho}(\boldsymbol{\Omega}_0 \times \boldsymbol{\Omega}_1) \, .$$

Therefore, up to a constant, (11) is equivalent to

$$\inf \left\{ \varepsilon \, \mathbf{KL}(\boldsymbol{\pi}|\mathbf{k}) \middle| \boldsymbol{\pi} \in \mathbb{R}^{M \times N}, \, \boldsymbol{\pi} \, \mathbb{1}_N = \boldsymbol{\mu}, \, \boldsymbol{\pi}^\top \mathbb{1}_M = \boldsymbol{\nu} \right\} \tag{13}$$

As usual, we also study a corresponding dual problem. For this we need the Fenchel–Legendre conjugate of $\mathbf{KL}(\cdot|\mathbf{k})$ for fixed $\mathbf{k}$.

Lemma 6.3.

$$\mathbf{KL}^*(\cdot|\mathbf{k}) : \mathbb{R}^{M \times N} \to \mathbb{R} \cup \{\infty\}, \qquad \boldsymbol{\phi} \mapsto \sum_{i=0}^{M-1} \sum_{j=0}^{N-1} [\exp(\boldsymbol{\phi}_{i,j}) - 1] \, \mathbf{k}_{i,j}$$

Proposition 6.4. Minimizers of (13) exist and a dual problem is given by

$$\sup \left\{ \langle \boldsymbol{\alpha}, \boldsymbol{\mu} \rangle + \langle \boldsymbol{\beta}, \boldsymbol{\nu} \rangle - \varepsilon \, \mathbf{KL}^*(\boldsymbol{\alpha} \oplus \boldsymbol{\beta}/\varepsilon|\mathbf{k}) \middle| (\boldsymbol{\alpha}, \boldsymbol{\beta}) \in \mathbb{R}^M \times \mathbb{R}^N \right\} \tag{14}$$

where $\boldsymbol{\alpha} \oplus \boldsymbol{\beta} \in \mathbb{R}^{M \times N}$ is defined by $(\boldsymbol{\alpha} \oplus \boldsymbol{\beta})_{i,j} = \boldsymbol{\alpha}_i + \boldsymbol{\beta}_j$.

Proof.

- We write (14) as

$$-\inf \left\{ f(\boldsymbol{\alpha}, \boldsymbol{\beta}) + g(A\,(\boldsymbol{\alpha}, \boldsymbol{\beta})) \middle| (\boldsymbol{\alpha}, \boldsymbol{\beta}) \in \mathbb{R}^M \times \mathbb{R}^N \right\}$$

  with

$$\begin{aligned} f &: \mathbb{R}^M \times \mathbb{R}^N \to \mathbb{R} \cup \{\infty\}, & (\boldsymbol{\alpha}, \boldsymbol{\beta}) &\mapsto -\langle \boldsymbol{\alpha}, \boldsymbol{\mu} \rangle - \langle \boldsymbol{\beta}, \boldsymbol{\nu} \rangle \, , \\ g &: \mathbb{R}^{M \times N} \to \mathbb{R} \cup \{\infty\}, & \boldsymbol{\phi} &\mapsto \varepsilon \, \mathbf{KL}^*(\boldsymbol{\phi}/\varepsilon|\mathbf{k}), \\ A &: \mathbb{R}^M \times \mathbb{R}^N \to \mathbb{R}^{M \times N}, & (\boldsymbol{\alpha}, \boldsymbol{\beta}) &\mapsto \boldsymbol{\alpha} \oplus \boldsymbol{\beta} \, . \end{aligned}$$

- f and g are convex and lower continuous, A is linear and bounded, hence we can apply the Fenchel–Rockafellar theorem, Theorem 2.4, to find that this is equivalent to solving

$$-\max \left\{ -f^*(-A^* \, \boldsymbol{\pi}) - g^*(\boldsymbol{\pi}) \middle| \boldsymbol{\pi} \in \mathbb{R}^{M \times N} \right\} = \min \left\{ f^*(-A^* \, \boldsymbol{\pi}) + g^*(\boldsymbol{\pi}) \middle| \boldsymbol{\pi} \in \mathbb{R}^{M \times N} \right\} \, .$$

- We obtain:

$$\begin{aligned} f^* &: \mathbb{R}^M \times \mathbb{R}^N \to \mathbb{R} \cup \{\infty\}, & (\boldsymbol{\sigma}, \boldsymbol{\tau}) &\mapsto \iota_{\{-\boldsymbol{\mu}\}}(\boldsymbol{\sigma}) + \iota_{\{-\boldsymbol{\nu}\}}(\boldsymbol{\tau}), \\ g^* &: \mathbb{R}^{M \times N} \to \mathbb{R} \cup \{\infty\}, & \boldsymbol{\pi} &\mapsto \varepsilon \, \mathbf{KL}(\boldsymbol{\pi}|\mathbf{k}), \\ A^* &: \mathbb{R}^{M \times N} \to \mathbb{R}^M \times \mathbb{R}^N, & \boldsymbol{\pi} &\mapsto (\boldsymbol{\pi} \, \mathbb{1}_M, \boldsymbol{\pi}^\top \mathbb{1}_N) \, . \end{aligned}$$

- In particular:

$$f^*(-A^*\boldsymbol{\pi}) = \iota_{\{-\boldsymbol{\mu}\}}(-\boldsymbol{\pi}\,\mathbb{1}_M) + \iota_{\{-\boldsymbol{\nu}\}}(-\boldsymbol{\pi}^\top\mathbb{1}_N) = \iota_{\{\boldsymbol{\mu}\}}(\boldsymbol{\pi}\,\mathbb{1}_M) + \iota_{\{\boldsymbol{\nu}\}}(\boldsymbol{\pi}^\top\mathbb{1}_N)$$

  which enforces the marginal constraints.

- Existence of a minimizer of (13) follows again from the Fenchel–Rockafellar theorem.

□

Note that the minimal value may be $+\infty$ if there is no coupling $\boldsymbol{\pi}$ such that $\boldsymbol{\pi} \ll \mathbf{k}$. However, since we assumed that $\boldsymbol{\rho} > 0$, we excluded this case. Similarly to the standard transport case we now establish existence of dual optimizers.

Proposition 6.5. If $\boldsymbol{\mu}_i > 0$, $\boldsymbol{\nu}_j > 0$ and $\mathbf{k}_{i,j} > 0$ for all (i,j) then optimizers of the dual regularized problem (14) exist.

Proof.
- In this proof we use intrinsically finite-dimensional arguments.
- Let $\boldsymbol{\mu} \otimes \boldsymbol{\nu} \in \mathbb{R}^{M\times N}$ be the product measure of $\boldsymbol{\mu}$ and $\boldsymbol{\nu}$:

$$(\boldsymbol{\mu} \otimes \boldsymbol{\nu})_{i,j} = \boldsymbol{\mu}_i \cdot \boldsymbol{\nu}_j$$

  Then

$$\langle \boldsymbol{\alpha}, \boldsymbol{\mu}\rangle + \langle \boldsymbol{\beta}, \boldsymbol{\nu}\rangle = \langle \boldsymbol{\alpha} \oplus \boldsymbol{\beta}, \boldsymbol{\mu} \otimes \boldsymbol{\nu}\rangle$$

  Thus, (14) can be rewritten as

$$\sum_{i=0}^{M-1}\sum_{j=0}^{N-1} f_{i,j}(\boldsymbol{\alpha}_i + \boldsymbol{\beta}_j) \quad \text{where} \quad f_{i,j}(z) = z \cdot (\boldsymbol{\mu}_i \cdot \boldsymbol{\nu}_j) - \varepsilon[\exp(z/\varepsilon) - 1] \cdot \mathbf{k}_{i,j}$$

- Each $f_{i,j}$ has compact super-level sets (this is where the assumption of strict positivity is used) and since there are only a finite number of terms and each component of $\boldsymbol{\alpha} \oplus \boldsymbol{\beta}$ is used by one term, for any minimizing sequence of (14) we can extract a subsequence such that $\boldsymbol{\alpha} \oplus \boldsymbol{\beta}$ converges.
- Note that shifting $\boldsymbol{\alpha}$ by adding a constant to each component and simultaneously shifting $\boldsymbol{\beta}$ by subtracting the same constant from each component does not change the objective value of (14). Hence, we may renormalize the minimizing sequence such that always $\boldsymbol{\alpha}_0 = 0$. Then convergence of $\boldsymbol{\alpha} \oplus \boldsymbol{\beta}$ implies convergence of $\boldsymbol{\beta}$ and thus eventually also convergence of $\boldsymbol{\alpha}$.

□

The regularized dual problem, (14) is an unconstrained, smooth, convex optimization problem. This will allow simple optimization (see next section) and a simple relation between primal and dual optimizers which we establish now.
The following Lemma for convex, lower semi-continuous functions is often useful.

Lemma 6.6. Let X be a Hilbert space and let f be a proper, convex, lower-semicontinuous function $f : X \to \mathbb{R} \cup \{\infty\}$. Then the Fenchel–Young inequality holds:

$$f(x) + f^*(y) \geq \langle x, y\rangle$$

Moreover, one has

$$[y \in \partial f(x)] \Leftrightarrow [f(x) + f^*(y) = \langle x, y\rangle] \Leftrightarrow [x \in \partial f(y)]$$

Proof. See, for instance, [Bauschke and Combettes, 2011, Proposition 13.13 and Theorem 16.23]. □

This leads to the following result on the relation between primal and dual optimizers.

Proposition 6.7. Let X and Y be Hilbert spaces, let $f : X \to \mathbb{R} \cup \{\infty\}$, $g : Y \to \mathbb{R} \cup \{\infty\}$ be proper, convex, lower semi-continuous functions, let $A : X \to Y$ be linear and bounded. Assume, that there is duality between the primal and dual problem

$$\inf \{f(x) + g(A\,x) | x \in X\} = \sup \{-f^*(-A^*y) - g^*(y) | y \in Y\}$$

and that the optimal value is finite. Then the following are equivalent:

(i) x and y are primal and dual optimizers,

(ii) $-A^*y \in \partial f(x)$ and $y \in \partial g(A\,x)$.

Proof.

- The proof follows [Bauschke and Combettes, 2011, Theorem 19.1].

$$\begin{aligned}
\text{(i)} &\Leftrightarrow f(x) + g(A\,x) = -f^*(-A^*y) - g^*(y) \\
&\Leftrightarrow f(x) + f^*(-A^*y) + g(A\,x) + g^*(y) = 0 = \langle x, -A^*y\rangle_X + \langle A\,x, y\rangle_Y
\end{aligned}$$

with Fenchel–Young inequality:

$$\begin{aligned}
&\Leftrightarrow [f(x) + f^*(-A^*y) = \langle x, -A^*y\rangle_X] \wedge [g(A\,x) + g^*(y) = \langle A\,x, y\rangle_Y] \\
&\Leftrightarrow [-A^*y \in \partial f(x)] \wedge [y \in \partial g(A\,x)] \Leftrightarrow \text{(ii)}
\end{aligned}$$

□

Corollary 6.8. We can apply this to the entropy regularized transport problem. Then

- $X = \mathbb{R}^M \times \mathbb{R}^n$, $Y = \mathbb{R}^{M \times N}$,
- $f : (\boldsymbol{\alpha}, \boldsymbol{\beta}) \mapsto -\langle \boldsymbol{\alpha}, \boldsymbol{\mu}\rangle - \langle \boldsymbol{\beta}, \boldsymbol{\nu}\rangle$, $\partial f(\boldsymbol{\alpha}, \boldsymbol{\beta}) = \{\nabla f(\boldsymbol{\alpha}, \boldsymbol{\beta})\}$, $\nabla f(\boldsymbol{\alpha}, \boldsymbol{\beta}) = (-\boldsymbol{\mu}, -\boldsymbol{\nu})$,
- $g = \varepsilon\, \mathbf{KL}^*(\cdot/\varepsilon | \mathbf{k})$, $\partial g(\boldsymbol{\phi}) = \{\nabla g(\boldsymbol{\phi})\}$ with $[\nabla g(\boldsymbol{\phi})]_{i,j} = \exp(\boldsymbol{\phi}_{i,j}/\varepsilon) \cdot \mathbf{k}_{i,j}$,
- $A(\boldsymbol{\alpha}, \boldsymbol{\beta}) = \boldsymbol{\alpha} \oplus \boldsymbol{\beta}$, $A^*\boldsymbol{\pi} = (\boldsymbol{\pi}\, \mathbb{1}_M, \boldsymbol{\pi}^\top \mathbb{1}_N)$,

and therefore, $\boldsymbol{\pi}$ and $(\boldsymbol{\alpha}, \boldsymbol{\beta})$ are primal and dual optimizers if and only if

- $(\boldsymbol{\pi}\, \mathbb{1}_M, \boldsymbol{\pi}^\top \mathbb{1}_N) = (\boldsymbol{\mu}, \boldsymbol{\nu})$ and
- $\boldsymbol{\pi}_{i,j} = \exp((\boldsymbol{\alpha}_i + \boldsymbol{\beta}_j)/\varepsilon) \cdot \mathbf{k}_{i,j}$.

Comment: Discuss relation / limit to unregularized problem.

6.3 Sinkhorn algorithm

Now we devise an algorithm to solve the regularized dual problem (14). For convenience, we introduce a name for the objective:

$$J : \mathbb{R}^M \times \mathbb{R}^N \to \mathbb{R}, \qquad (\boldsymbol{\alpha}, \boldsymbol{\beta}) \mapsto \langle \boldsymbol{\alpha}, \boldsymbol{\mu} \rangle + \langle \boldsymbol{\beta}, \boldsymbol{\nu} \rangle - \varepsilon\, \mathbf{KL}^*(\boldsymbol{\alpha} \oplus \boldsymbol{\beta}/\varepsilon | \mathbf{k})$$

Since J is smooth, concave and the maximization problem is unconstrained, a necessary and sufficient optimality condition is a vanishing gradient. We obtain for the gradient of J:

$$\frac{\partial J}{\partial \boldsymbol{\alpha}_i} = \boldsymbol{\mu}_i - \exp(\boldsymbol{\alpha}_i/\varepsilon) \sum_{j=0}^{N-1} \exp(\boldsymbol{\beta}_j/\varepsilon)\, \mathbf{k}_{i,j}, \qquad \frac{\partial J}{\partial \boldsymbol{\beta}_j} = \boldsymbol{\nu}_j - \exp(\boldsymbol{\beta}_j/\varepsilon) \sum_{i=0}^{M-1} \exp(\boldsymbol{\alpha}_i/\varepsilon)\, \mathbf{k}_{i,j}.$$

We need to find a pair $(\boldsymbol{\alpha}, \boldsymbol{\beta})$ such that $\nabla J = 0$. We propose to do this in an alternating fashion. Assume some $\boldsymbol{\beta}^{(\ell)} \in \mathbb{R}^N$ is given. For fixed $\boldsymbol{\beta}^{(\ell)}$ the optimal $\boldsymbol{\alpha}^{(\ell+1)}$ can be found by solving $\frac{\partial J}{\partial \boldsymbol{\alpha}_i} = 0$. Then, for fixed $\boldsymbol{\alpha}^{(\ell+1)}$ the optimal dual variable $\boldsymbol{\beta}^{(\ell+1)}$ can be found by solving $\frac{\partial J}{\partial \boldsymbol{\beta}_j} = 0$. One obtains:

$$\exp(\boldsymbol{\alpha}^{(\ell+1)}/\varepsilon) = \frac{\boldsymbol{\mu}}{\mathbf{k}\, \exp(\boldsymbol{\beta}^{(\ell)}/\varepsilon)}, \qquad \exp(\boldsymbol{\beta}^{(\ell+1)}/\varepsilon) = \frac{\boldsymbol{\nu}}{\mathbf{k}^\top \exp(\boldsymbol{\alpha}^{(\ell+1)}/\varepsilon)}. \tag{15}$$

Here the exponential function and the fraction are understood to act component-wise. The algorithm becomes particularly simple when one substitutes the variables

$$\mathbf{u}^{(\ell)} = \exp(\boldsymbol{\alpha}^{(\ell)}/\varepsilon), \qquad \mathbf{v}^{(\ell)} = \exp(\boldsymbol{\beta}^{(\ell)}/\varepsilon). \tag{16}$$

One obtains:

$$\mathbf{u}^{(\ell+1)} = \frac{\boldsymbol{\mu}}{\mathbf{k}\, \mathbf{v}^{(\ell)}}, \qquad \mathbf{v}^{(\ell+1)} = \frac{\boldsymbol{\nu}}{\mathbf{k}^\top \mathbf{u}^{(\ell+1)}}. \tag{17}$$

These are the famous Sinkhorn iterations [Sinkhorn, 1964]. With Corollary 6.8 the primal dual optimality condition then becomes $\boldsymbol{\pi} = \operatorname{diag}(\mathbf{u})\, \mathbf{k}\, \operatorname{diag}(\mathbf{v})$ where $\operatorname{diag}(\mathbf{u})$ and $\operatorname{diag}(\mathbf{v})$ are the diagonal matrices with diagonal entries given by $\mathbf{u}$ and $\mathbf{v}$ respectively.

Proposition 6.9. If $\mathbf{k}$ is strictly positive, iterations (15) converge to a solution of the regularized dual problem (14), up to constant shifts of $\boldsymbol{\alpha}$ and $\boldsymbol{\beta}$.

Sketch of proof.

- As long as the iterates change, the dual objective is strictly increasing. If the iterates do not change, by virtue of the first order optimality conditions an optimal solution is found.
- Since $\mathbf{k}$ is strictly positive, the dual objective is bounded from above, and hence, arguing with compact super-level sets as above, up to constant shifts, one must find a convergent subsequence.
- Since the map that takes a given iterate $(\boldsymbol{\alpha}^{(\ell)}, \boldsymbol{\beta}^{(\ell)})$ to the subsequent iterate $(\boldsymbol{\alpha}^{(\ell+1)}, \boldsymbol{\beta}^{(\ell+1)})$ is continuous, any cluster point of the iterates must be a fixed point of the iteration and thus be an optimal dual solution.

□

Remark 6.10 (Discussion of Sinkhorn algorithm).

- Iterations (17) only consist of matrix-vector multiplications and pointwise division and are thus trivial to implement and parallelize.

- The algorithm works on any cost function $\boldsymbol{c}$ and does not depend on additional theoretical properties. Some caution is required if $\boldsymbol{c}$ may be infinite.
- As $\varepsilon \to 0$ one approaches the original unregularized transport problem but numerically two important issues arise: 1) convergence becomes increasingly slow, 2) numerical values in $\mathbf{k}$, $\mathbf{u}$ and $\mathbf{v}$ become increasingly small / large and will eventually exceed the computational numerical range.
- In summary: if one can tolerate (or even wishes) a moderate regularization strength $\varepsilon > 0$ the Sinkhorn algorithm is a good starting point to get numerical results.

References

L. Ambrosio, N. Fusco, and D. Pallara. *Functions of Bounded Variation and Free Discontinuity Problems*. Oxford mathematical monographs. Oxford University Press, 2000.

L. Ambrosio and P. Tilli. *Topics on Analysis in Metric Spaces*. Number 25 in Oxford Lecture Series in Mathematics and its Applications. Oxford University Press, 2004.

L. Ambrosio, N. Gigli, and G. Savaré. *Gradient Flows in Metric Spaces and in the Space of Probability Measures*. Lectures in Mathematics. Birkhäuser Boston, 2005.

H. H. Bauschke and P. L. Combettes. *Convex Analysis and Monotone Operator Theory in Hilbert Spaces*. CMS Books in Mathematics. Springer, 1st edition, 2011.

D. P. Bertsekas and J. Eckstein. Dual coordinate step methods for linear network flow problems. *Mathematical Programming, Series B*, 42:203–243, 1988.

Y. Brenier. Polar factorization and monotone rearrangement of vector-valued functions. *Comm. Pure Appl. Math.*, 44(4):375–417, 1991.

A. J. Kurdila and M. Zabarankin. *Convex functional analysis*, volume 1 of *Systems and Control: Foundations and Applications*. Birkhäuser, 2005.

M. D. Perlman. Jensen's inequality for a convex vector-valued function on an infinite-dimensional space. *Journal of Multivariate Analysis*, 4(1):52–65, 1974.

R. T. Rockafellar. Duality and stability in extremum problems involving convex functions. *Pacific J. Math*, 21(1):167–187, 1967.

W. Rudin. *Real and complex analysis*. McGraw-Hill Book Company, 3rd edition, 1986.

F. Santambrogio. *Optimal Transport for Applied Mathematicians*, volume 87 of *Progress in Nonlinear Differential Equations and Their Applications*. Birkhäuser Boston, 2015.

B. Schmitzer and B. Wirth. Dynamic models of Wasserstein-1-type unbalanced transport. arXiv:1705.04535, 2017.

R. Sinkhorn. A relationship between arbitrary positive matrices and doubly stochastic matrices. *Ann. Math. Statist.*, 35(2):876–879, 1964.

C. Villani. *Topics in Optimal Transportation*, volume 58 of *Graduate Studies in Mathematics*. American Mathematical Society, Providence, 2003.

C. Villani. *Optimal Transport: Old and New*, volume 338 of *Grundlehren der mathematischen Wissenschaften*. Springer, 2009.

W. P. Ziemer. *Weakly Differentiable Functions*, volume 120 of *Graduate Texts in Mathematics*. Springer New York, 1989.