



HAL
open science

Quantitative Stability in Quadratic Optimal Transport

Alex Delalande

► **To cite this version:**

Alex Delalande. Quantitative Stability in Quadratic Optimal Transport. Optimization and Control [math.OC]. Université Paris-Saclay, 2022. English. NNT : 2022UPASM044 . tel-03935445

HAL Id: tel-03935445

<https://theses.hal.science/tel-03935445v1>

Submitted on 11 Jan 2023

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Quantitative Stability in Quadratic Optimal Transport

Stabilité quantitative en transport optimal quadratique

Thèse de doctorat de l'université Paris-Saclay

Ecole Doctorale de Mathématique Hadamard (EDMH) n° 574

Spécialité de doctorat: Mathématiques appliquées

Graduate School : Mathématiques, Référent : Faculté des sciences d'Orsay

Thèse préparée dans le Laboratoire de Mathématiques d'Orsay (Université Paris-Saclay, CNRS), sous la direction de Quentin MÉRIGOT, Professeur (Université Paris-Saclay), et la co-direction de Frédéric CHAZAL, Directeur de Recherche (Inria Saclay)

Thèse soutenue à Paris-Saclay, le 14 décembre 2022, par

Alex DELALANDE

Composition du jury

Julie Delon Professeure, Université Paris Cité	Présidente
Max Fathi Professeur, Université Paris Cité	Rapporteur & Examineur
Dejan Slepčev Professeur, Carnegie Mellon University	Rapporteur & Examineur
Yann Brenier Directeur de Recherche, CNRS - Université Paris-Saclay	Examineur
Alessio Figalli Professeur, ETH Zürich	Examineur
Quentin Mérigot Professeur, Université Paris-Saclay	Directeur de thèse

Remerciements

Je veux exprimer ici ma sincère reconnaissance à celles et ceux qui ont participé, de près ou de loin, à rendre possible la rédaction et la soutenance de cette thèse.

To Max Fathi and Dejan Slepčev, who agreed to spend a substantial amount of their time proofreading this manuscript carefully, and to Yann Brenier, Julie Delon and Alessio Figalli, who gave me the honor and pleasure to participate in the jury of my defense.

À mes directeurs de thèse, pour leur encadrement sans faille au cours de ces trois années et quelques. À Quentin, qui m'a initié à de nombreux et jolis domaines des mathématiques, avec qui les discussions ont toujours été stimulantes et sympathiques et dont les idées et les conseils ont toujours été plus que judicieux. Son approche des problèmes et son efficacité resteront un modèle pour ma carrière scientifique. À Frédéric également, qui a su me faire confiance dès ma sortie de M2 en me proposant un sujet riche et une grande liberté, et qui m'a toujours fourni de bons conseils.

À Guillaume, avec qui j'ai eu la chance et la joie de collaborer, que j'admire pour son recul et remercie pour sa bienveillance et ses nombreux conseils.

Aux membres permanents et non-permanents du LMO et de l'équipe DataShape, et plus largement à toutes les travailleuses et travailleurs de l'Université Paris-Saclay et de l'Inria Saclay, qui ont contribué à établir l'environnement de travail stimulant et agréable dans lequel s'est déroulée cette thèse. J'ai une pensée particulière pour tous mes camarades (ex-)doctorants, collègues et amis, spécialistes du transport optimal ou autres types de chantiers en tous genres, avec lesquels les conversations scientifiques et non-scientifiques se sont souvent mélangées à mon plus grand bonheur.

À celles et ceux qui m'ont toujours fait me sentir entouré. Aux ami.es de plus ou moins longue date, parfois (ex-)Centraliens ou ex-singapouriens, qui ont toujours répondu présent quand il s'agissait de se changer les idées, et parmi lesquel.les j'ai trouvé de nombreuses sources d'inspiration. À ma famille, véritable constante de toutes ces années, qui n'a jamais manqué à me supporter même dans les périodes les plus rudes. Aux autres familles aussi, Muller, Wolfermann et assimilé.es, pour leur accueil et leur tendresse.

Enfin, et surtout, à Maud, pour les conversations scientifiques et non-scientifiques, pour ta présence, pour la source d'inspiration que tu es, pour ta constance, ta tendresse, et pour bien plus encore.

Orsay, le 28 novembre 2022

Contents

Introduction générale	1
General introduction	15
Notation	28
I Strong convexity of the quadratic optimal transport problem	29
1 Quadratic optimal transport and strong convexity of the dual	31
1.1 Quadratic optimal transport	31
1.2 Strong convexity of the Kantorovich functional	42
1.3 Some geometric and functional inequalities	49
1.4 Extension to source measures with non-convex support	52
2 A semi-discrete approach	59
2.1 Introduction	59
2.2 Semi-discrete optimal transport	60
2.3 Proof in the semi-discrete case	62
2.4 From the semi-discrete case to the general case	66
3 A continuous approach	69
3.1 Introduction	69
3.2 Derivatives of the Kantorovich functional for regular potentials	70
3.3 Proof in the regular case	71
3.4 From the regular case to the general case	74
4 An entropic approach	79
4.1 Introduction	79
4.2 Entropic Kantorovich functional	81
4.3 Semi-discrete entropic Kantorovich functional	82
4.4 Strong convexity of entropic Kantorovich functional	85
4.5 From the entropic case to the classical case	88

II	Consequences for the stability of solutions to optimal transport problems	91
5	Quantitative stability of optimal transport maps with respect to the target measure	93
5.1	Introduction	93
5.2	Stability of potentials	100
5.3	Stability of optimal transport maps	103
5.4	Gagliardo–Nirenberg type inequality	108
6	Quantitative stability of Wasserstein barycenters with respect to the marginals	113
6.1	Introduction	113
6.2	Stability of potentials	122
6.3	Stability of push-forwards	125
6.4	Dual formulation	129
7	Quantitative stability of Schrödinger potentials with respect to the temperature in the semi-discrete setting	133
7.1	Introduction	133
7.2	Bounds for semi-discrete entropic optimal transport	135
7.3	A governing ODE	140
7.4	Bound on the second term of the ODE	142
7.5	Estimates for entropic semi-discrete solutions	148
7.6	Expansion of the regularized cost	151
7.7	Numerical illustrations	155
III	Numerical applications: the Linearized Optimal Transport framework	159
8	Linearized optimal transport and applications	161
8.1	Introduction	161
8.2	Behavior of the LOT embedding	163
8.3	Applications	166
	Conclusion	179

IV	Appendix	183
A	Optimal transport facts	A-1
A.1	Monge and Kantorovich formulations	A-1
A.2	Dual formulation	A-2
A.3	Stability of solutions	A-4
A.4	Wasserstein distances and spaces	A-5

Introduction générale

Le problème du transport optimal constitue le principal objet d'étude de cette thèse. En termes simples, ce problème peut être formulé de la manière suivante :

Étant données deux distributions d'une même quantité de masse sur un certain espace et connaissant le coût de transfert d'une unité de masse depuis n'importe quel endroit vers n'importe quel autre endroit, comment transporter toute la masse de la première distribution à la seconde de la manière la moins coûteuse possible ?

Un tel problème présente une capacité évidente à modéliser des situations pratiques : la masse évoquée peut être matérialisée par une foule de personnes, ou par une collection de particules de gaz, ou par un ensemble de marchandises ; le coût peut correspondre à une distance (pour mesurer un effort de déplacement), ou au carré d'une distance (pour mesurer une énergie cinétique), ou simplement à un coût financier. En présence d'un problème mathématique qui vise à modéliser des phénomènes physiques, il est pertinent de vérifier si ce problème est *bien posé* au sens de Hadamard (Hadamard, 1902). Cette notion informelle regroupe certaines des caractéristiques souhaitables qu'un problème *naturel* doit présenter et qui doivent aider à sa résolution. Un problème donné est dit bien posé si :

- (i) il admet effectivement une solution ;
- (ii) cette solution est unique ;
- (iii) cette solution dépend continûment des données du problème.

En citant Evans dans l'introduction de son livre *Partial Differential Equations* (Evans, 2010), la condition (iii) "est particulièrement importante pour les problèmes issus d'applications physiques : nous préférons que notre solution (unique) ne change qu'un peu lorsque les conditions spécifiant le problème changent un peu. Pour de nombreux problèmes, en revanche, l'unicité n'est pas à espérer".

Dans sa formulation moderne, le problème du transport optimal est globalement bien posé. Une solution à ce problème correspond à une méthode concrète pour effectivement transporter toute la masse positionnée dans une première configuration vers une seconde configuration tout en réalisant le plus faible coût global de transport possible. Nous verrons plus loin que la propriété d'existence de la solution (i) est maintenant parfaitement comprise et qu'il est démontré qu'elle est valable dans la plupart des cas intéressants. La question (ii) de l'unicité de la solution a également été largement étudiée et de nombreux cas où la solution est effectivement unique, ou à l'inverse où elle n'est pas susceptible de l'être, ont été décrits. L'étude de la dernière propriété (iii) est cependant moins avancée. Il existe en général des garanties abstraites de stabilité qui assurent que les solutions de problèmes de transport optimal dépendent de manière continue des distributions de masse qui définissent le problème. Cependant, sauf dans de très rares cas, ces garanties ne sont pas quantitatives : nous ne savons pas en général comment un changement donné dans ces distributions affecte les solutions de transport optimal correspondantes. Comme

les applications modernes obligent presque toujours les praticiens à utiliser des approximations statistiques ou informatiques des données d'intérêt, ce manque de garanties quantitatives est problématique. L'objectif de cette thèse est de travailler à combler cette lacune dans la théorie du transport optimal.

Théorie du transport optimal et applications : un (très) bref aperçu. Le problème du transport optimal a été introduit pour la première fois par Monge en 1781, avec en tête des applications militaires et d'ingénierie (Monge, 1781). Il a formulé le problème général consistant à trouver le moyen le moins cher de transporter une quantité donnée de terre d'un site d'extraction à un site de construction, le coût de transport de chaque *molécule* de terre étant proportionnel à la distance qu'elle parcourt. L'étude de ce problème l'a conduit à la découverte de concepts importants en géométrie des surfaces, mais le problème est resté largement non résolu. Le problème du transport optimal a été relancé par Kantorovich, qui a donné en 1942 (Kantorovich, 1942) sa formulation moderne sous la forme d'un problème de programmation linéaire. En termes mathématiques bruts, le problème de Kantorovich peut être décrit comme suit. Considérons un certain espace Ω , typiquement un espace polonais compact, et deux mesures de probabilité ρ, μ dans $\mathcal{P}(\Omega)$ qui représentent chacune une distribution de masse. Étant donné une fonction $c(x, y)$ qui représente le coût du transfert d'une unité de masse d'un emplacement x dans Ω à un emplacement y dans Ω , Kantorovich propose de résoudre

$$\min_{\gamma \in \Gamma(\rho, \mu)} \int_{\Omega \times \Omega} c(x, y) d\gamma(x, y), \quad (1)$$

où $\Gamma(\rho, \mu)$ est l'ensemble des *plans de transport* entre ρ et μ , c'est-à-dire l'ensemble des mesures de probabilité sur $\Omega \times \Omega$ avec pour première marginale ρ et pour seconde marginale μ . Un candidat γ dans $\Gamma(\rho, \mu)$ propose un plan pour transporter la masse de ρ à μ en envoyant une portion $d\gamma(x, y)$ de la masse $d\rho(x)$ d'un emplacement source x à un emplacement cible y . La linéarité de la formulation de Kantorovich lui a permis d'assurer, sous des hypothèses faibles sur c , l'existence de solutions à son problème dans des espaces métriques compacts ainsi que d'établir une formulation duale et des conditions d'optimalité. Kantorovich et d'autres auteurs ont rapidement compris que la valeur du coût de transport optimal (1) entre deux mesures ρ et μ pouvait donner une idée quantitative de la *similarité* entre ρ et μ . Il a été montré que lorsque $c(x, y) = d_\Omega(x, y)^p$ est la p -ième puissance d'une distance d_Ω sur Ω pour un certain $p \geq 1$, la valeur de (1) correspond elle-même à la p -ième puissance d'une distance entre ρ et μ . Cette distance, généralement appelée p -ième distance de Wasserstein¹ et noté W_p , confère à l'ensemble des mesures de probabilité $\mathcal{P}(\Omega)$ une riche structure géométrique *hissée* de l'espace de base Ω . Par exemple, un espace métrique compact (Ω, d_Ω) est plongé isométriquement dans l'espace de Wasserstein- p $(\mathcal{P}(\Omega), W_p)$ par l'application $x \mapsto \delta_x$ (où δ_x désigne la masse de Dirac en x). La géométrie fournie par les métriques de Wasserstein sur les espaces de mesures de probabilité s'est avérée très pratique, tant pour des considérations théoriques qu'appliquées. L'unique cas quadratique $p = 2$ sur $\Omega = \mathbb{R}^d$ a produit à lui seul une théorie très substantielle. Il a permis de définir des notions d'interpolations (McCann, 1997) ou de barycentres (Agueh and Carlier, 2011) sur des

¹L'attribution de ce nom à cette distance est souvent remise en question. Dans les travaux de Kantorovich, la première apparition d'une telle notion est dans son travail conjoint avec Rubinstein (Kantorovich and Rubinstein, 1958) pour le cas $p = 1$. Nous renvoyons aux notes bibliographiques du Chapitre 6 de (Villani, 2008) pour plus de détails.




Figure 1: (Haut) Interpolation linéaire $(1-t)\rho_0 + t\rho_1$ et (Bas) Interpolation *par déplacement* $\rho_t := \arg \min_{\mu \in \mathcal{P}(\mathbb{R})} (1-t)W_2^2(\rho_0, \mu) + tW_2^2(\rho_1, \mu)$ de deux distributions Gaussiennes ρ_0, ρ_1 sur $(\mathbb{R}, |\cdot|)$, avec pour moyenne respective 0 et 15 et chacune de variance unitaire. Dans beaucoup d'applications, le mouvement *horizontal* de l'interpolant *par déplacement* est préférable au mouvement *vertical* de l'interpolant linéaire.

familles de mesures de probabilité présentant de fortes caractéristiques géométriques (voir Figure 1). Plus généralement, la géométrie de l'espace Wasserstein-2 s'est révélée donner une structure riemannienne physiquement pertinente aux espaces de mesures de probabilité, dans laquelle certaines EDP d'évolution classiques (telles que les équations de Fokker-Planck ou des milieux poreux) ont pu être exprimées comme les flots de gradient de fonctionnelles d'énergie bien choisies sur l'espace des distributions de probabilité. D'autres EDP d'évolution importantes ont trouvé une formulation variationnelle dans les espaces de Wasserstein, comme les équations d'Euler en dynamique des fluides (Brenier, 1989, 1999).

Parallèlement au développement de sa théorie, le problème du transport optimal a fait de nombreuses incursions réussies dans les applications. Kantorovich a introduit son programme linéaire (1) pour modéliser des problèmes courants d'allocation de ressources survenant en économie, un domaine où le transport optimal suscite encore un vif intérêt (Galichon, 2016). Dans ces applications, les ressources considérées sont souvent de nature discrète et les mesures de probabilité ρ et μ peuvent être supposées à support fini. Dans un tel contexte, le problème (1) correspond à un programme linéaire classique de dimension finie, un problème qui a rapidement été résolu numériquement avec l'algorithme du simplexe de Dantzig (Dantzig, 1949, 1951) et de manière plus efficace avec les algorithmes pour les problèmes de flux à coût minimal (Ford and Fulkerson, 1962; Goldberg and Tarjan, 1989). La version discrète du problème (1) est également étroitement liée au problème d'assignement, qui a été résolu efficacement avec l'algorithme d'enchères de Bertsekas (Bertsekas, 1981; Bertsekas and Eckstein, 1988). Depuis les années 2000, le problème du transport optimal a également été de plus en plus utilisé pour résoudre diverses tâches de traitement de formes, d'images et de vidéos telles que le recalage (Haker et al., 2004), la réduction de scintillement (Delon, 2006), le transfert de couleurs (Pitié et al., 2007; Bonneel et al., 2016), le débruitage (Lellmann et al., 2014) ou la segmentation (Rabin and Papadakis, 2015). Dans le domaine de l'apprentissage automatique, le transport optimal a été utilisé pour la recherche d'images par le contenu (Rubner et al., 2000), l'apprentissage semi-supervisé (Solomon et al., 2014), la modélisation générative (Arjovsky et al., 2017), l'adaptation de domaine (Courty et al., 2017) ou l'optimisation robuste au sens des distributions (Kuhn et al., 2019). Le nombre de ces applications a considérablement augmenté suite aux avancées algorithmiques dues à (Cuturi, 2013), voir (Peyré and Cuturi, 2019) pour plus de références. Enfin, d'autres applications nota-

bles du problème du transport optimal se trouvent en chimie quantique (Buttazzo et al., 2012; Cotar et al., 2013), en conception optique (Oliker, 2003; Caffarelli and Oliker, 2008) et en statistique, où il a permis d'étendre les notions de quantiles aux variables aléatoires multivariées (Carlier et al., 2016; Chernozhukov et al., 2017), de construire des estimateurs de densité efficaces en inférence géométrique (Weed and Berthet, 2019; Divol, 2022) et d'analyser la convergence d'algorithmes d'échantillonnage tels que l'algorithme de Langevin Monte Carlo (Dalalyan, 2017; Bernton, 2018) ou l'algorithme de descente de gradient de Stein variationnel (Korba et al., 2020).

Caractère bien-posé du problème du transport optimal. La forte capacité du problème du transport optimal à modéliser des phénomènes physiques soulève de manière pressante la question de son caractère bien-posé : peut-on espérer une solution, la solution est-elle unique et répond-elle de manière continue aux modifications des données du problème ?

Nous avons déjà mentionné que Kantorovich a prouvé l'existence de solutions à (1) dans des cas généraux. Ce résultat d'existence a été généralisé plus encore dans (Kellerer, 1984), et on peut s'attendre en général à ce que (1) admette des solutions par exemple quand Ω est un espace polonais et c est semi-continue inférieurement et minorée.

L'unicité d'une solution à (1) n'est pas à attendre en général (voir la Figure 2 pour un exemple). Il existe cependant des cas particuliers intéressants où l'unicité de la solution est garantie. Le plus célèbre de ces cas est sans doute dû à Brenier (Brenier, 1987), qui a montré pour Ω un sous-ensemble compact de \mathbb{R}^d et $c(x, y) = \|x - y\|^2$ que lorsque la mesure source ρ est absolument continue par rapport à la mesure de Lebesgue, la solution de (1) est unique et, plus important encore, elle est supportée sur le graphe du gradient d'une fonction convexe (voir également (Knott and Smith, 1984; Smith and Knott, 1987; Rüschendorf and Rachev, 1990)). Incidemment, cette caractérisation a permis d'adopter le point de vue EDP suivant sur le problème de transport optimal : lorsque ρ et μ admettent des densités (notées avec les mêmes lettres), une fonction régulière et strictement convexe ϕ dont le graphe du gradient supporte la solution de transport optimal entre ρ et μ doit vérifier pour tout x de Ω la formule de changement de variable suivante :

$$\det(D^2\phi(x))\mu(\nabla\phi(x)) = \rho(x). \quad (2)$$

Ceci correspond à une équation de Monge-Ampère en ϕ , dont la solution fournit la solution au problème de transport optimal entre ρ et μ sous des conditions aux limites appropriées. Le résultat de Brenier a ensuite été généralisé à des coûts et domaines plus généraux, voir par exemple (Gangbo and McCann, 1996; Trudinger and Wang, 2001; McCann, 2001; Caffarelli et al., 2002; Bernard and Buffoni, 2007; Fathi and Figalli, 2010).

La stabilité des solutions de problèmes de transport optimal par rapport aux données qui les définissent est établie dans des cas généraux. Par exemple, le Théorème 5.19 de (Villani, 2008) assure que pour Ω un espace polonais et c une fonction de coût continue et bornée, la convergence faible des mesures source et cible ρ_n, μ_n dans $\mathcal{P}(\Omega)$ vers des limites respectives ρ, μ dans $\mathcal{P}(\Omega)$ entraîne, à sous-suite près, la convergence faible de solutions de transport optimal γ_n entre ρ_n et μ_n vers une solution de transport optimal γ entre ρ et μ . D'autres résultats assurent également en général la stabilité d'autres quantités de transport optimal telles que les interpolants et les barycentres dans les espaces de Wasserstein mentionnés plus haut. Ces garanties ne sont pas anecdotiques : elles assurent




Figure 2: Dans \mathbb{R}^2 , soit $\rho = \frac{1}{2}(\delta_{(-1,0)} + \delta_{(1,0)})$ et $\mu = \frac{1}{2}(\delta_{(0,-1)} + \delta_{(0,1)})$. Le transport optimal entre ρ et μ avec la distance euclidienne comme coût ($c(x, y) = \|x - y\|$) est réalisé par $\gamma^0 = \frac{1}{2}(\delta_{(-1,0) \times (0,1)} + \delta_{(1,0) \times (0,-1)})$ (Gauche) ou par $\gamma^1 = \frac{1}{2}(\delta_{(-1,0) \times (0,-1)} + \delta_{(1,0) \times (0,1)})$ (Droite), ou par n'importe laquelle des combinaisons convexes $(1-t)\gamma^0 + t\gamma^1$ pour $t \in [0, 1]$.




Figure 3: Représentation par histogrammes de l'interpolation *par déplacement* $\hat{\rho}_t^n := \arg \min_{\mu \in \mathcal{P}(\mathbb{R})} (1-t)W_2^2(\hat{\rho}_0^n, \mu) + tW_2^2(\hat{\rho}_1^n, \mu)$, où pour $k \in \{0, 1\}$, $\hat{\rho}_k^n = \frac{1}{n} \sum_{i=1}^n \delta_{x_k^i}$ avec $(x_k^i)_{1 \leq i \leq n} \sim \rho_k$ où $n = 2000$ et ρ_k pris comme en Figure 1.

par exemple qu'un plan de transport optimal approximatif entre deux mesures ρ et μ est donné par un plan de transport optimal $\hat{\gamma}$ entre deux approximations $\hat{\rho}, \hat{\mu}$ de ρ, μ . Ceci est particulièrement utile dans les applications, où l'utilisation d'approximations $\hat{\rho}, \hat{\mu}$ au lieu de ρ, μ peut être nécessaire soit en raison de limitations computationnelles, soit dans un contexte statistique où seuls des échantillons des mesures d'intérêt sont disponibles. Dans de telles applications cependant, les utilisateurs peuvent avoir besoin de garanties quantitatives sur les approximations commises. Prenons l'exemple de l'interpolation *par déplacement* présentée dans la Figure 3. Cette interpolation est une *approximation statistique* de celle affichée dans la partie inférieure de la Figure 1, obtenue en approxinant les mesures des extrémités ρ_0 et ρ_1 avec des mesures empiriques $\hat{\rho}_0^n$ et $\hat{\rho}_1^n$ construites à partir d'échantillons avant de calculer leur interpolation. Dans des contextes typiques, des garanties statistiques sur la qualité des approximations de ρ_0 et ρ_1 en distance de Wasserstein sont disponibles, c'est-à-dire que des bornes (en espérance ou avec grande probabilité) sur $W_2(\rho_0, \hat{\rho}_0^n)$ et $W_2(\rho_1, \hat{\rho}_1^n)$ sont connues. Pour les applications en aval, il peut alors être important de savoir si ces garanties de qualité sont transmises aux interpolants, c'est-à-dire si la distance $W_2(\rho_t, \hat{\rho}_t^n)$ pour un certain $t \in (0, 1)$ peut être majorée en termes de $W_2(\rho_0, \hat{\rho}_0^n)$ et $W_2(\rho_1, \hat{\rho}_1^n)$. Dans le cadre unidimensionnel des Figures 1 et 3, la réponse est positive et on peut toujours assurer le comportement Lipschitz suivant :

$$W_p(\rho_t, \hat{\rho}_t^n) \leq (1-t)W_p(\rho_0, \hat{\rho}_0^n) + tW_p(\rho_1, \hat{\rho}_1^n).$$

Cependant, cette borne est spécifique à $\Omega = \mathbb{R}^d$ avec $d = 1$ et il n'existe pas de garantie quantitative similaire dès lors que $d \geq 2$. Cela soulève la question de la stabilité quantitative générale du problème de transport optimal.

Pour certaines EDP elliptiques, la stabilité quantitative des solutions peut être dé-

duite de garanties de forte ellipticité. Considérons par exemple sur un domaine régulier borné Ω de \mathbb{R}^d l'équation de Poisson

$$\Delta\phi = f, \tag{3}$$

avec une condition aux limites de Dirichlet nulle ($\phi = 0$ sur $\partial\Omega$). Dans cette équation, Δ désigne l'opérateur laplacien, dont l'ellipticité donne l'inégalité de Poincaré sur Ω . Cette inégalité assure en particulier qu'il existe une constante C telle que pour toute solution faible $\phi, \hat{\phi} \in H^1(\Omega)$ de (3) avec des seconds membres respectifs f, \hat{f} , on a

$$\|\phi - \hat{\phi}\|_{L^2(\Omega)}^2 \leq C \|\nabla\phi - \nabla\hat{\phi}\|_{L^2(\Omega)}^2 = -C \int_{\Omega} (\phi - \hat{\phi})(f - \hat{f}) dx \leq C \|\phi - \hat{\phi}\|_{L^2(\Omega)} \|f - \hat{f}\|_{L^2(\Omega)},$$

de sorte que $\|\phi - \hat{\phi}\|_{L^2(\Omega)} \leq C \|f - \hat{f}\|_{L^2(\Omega)}$. Une telle inégalité quantifie précisément l'effet d'une perturbation de la donnée d'entrée f sur la solution correspondante ϕ dans (3). Malheureusement, cette approche elliptique ne peut pas être facilement appliquée au problème du transport optimal. Par exemple, nous avons mentionné que dans le contexte quadratique et euclidien (qui correspond au cadre le plus étudié et sans doute le plus simple), le problème de transport optimal pouvait être reformulé dans certains cas en termes de l'équation de Monge-Ampère (2). En général, cette équation est seulement dégénérée elliptique, et des garanties de forte ellipticité ne sont disponibles que lorsque l'inconnue ϕ est régulière et fortement convexe, ce qui est rarement le cas en pratique. Cela rend la question de la stabilité quantitative du problème de transport optimal particulièrement difficile.

Contributions principales. Dans cette thèse, nous suivons une approche classique dans la théorie du transport optimal qui consiste à étudier le problème dual de (1) pour obtenir des informations qualitatives et quantitatives sur les solutions de transport optimal. Nous nous concentrons exclusivement sur le cadre euclidien et quadratique (c'est-à-dire $\Omega = \mathbb{R}^d$ et $c(x, y) = \|x - y\|^2$), laissant les généralisations de nos résultats à des travaux futurs. Dans ce cadre, le problème dual de (1) correspond essentiellement au problème de minimisation

$$\min_{\psi: \mathbb{R}^d \rightarrow \mathbb{R}} \int \psi^* d\rho + \int \psi d\mu, \tag{4}$$

où ψ^* désigne la conjuguée convexe du *potentiel* ψ .

Au lieu de chercher directement des estimations quantitatives de stabilité pour (1), nous cherchons d'abord des estimations quantitatives de stabilité pour (4). La fonctionnelle $\psi \mapsto \int \psi^* d\rho$ apparaissant dans ce problème dual, que nous appelons *fonctionnelle de Kantorovich* associée à une source ρ , est convexe. En tant que telle, des garanties quantitatives sur la stabilité des minimiseurs de (4) peuvent être déduites d'estimations de la forte convexité de cette fonctionnelle. Dans les Chapitres 2–4, nous obtenons des estimations explicites de la forte convexité de la fonctionnelle de Kantorovich, en nous appuyant principalement sur les inégalités de Brunn-Minkowski, Brascamp-Lieb et Prékopa-Leindler. Il était déjà compris, depuis l'article précurseur de (McCann, 1997), que ces inégalités géométriques et fonctionnelles sont liées au problème du transport optimal puisqu'elles peuvent être déduites de la convexité géodésique de certaines fonctionnelles d'énergie sur l'espace de Wasserstein-2. Cette thèse renforce ce lien, quelque

peu dans une direction opposée, en utilisant ces inégalités pour quantifier la forte convexité du dual du problème de transport optimal quadratique.

Ensuite, nous rassemblons dans les Chapitres 5–7 des conséquences des estimations de forte convexité des Chapitres 2–4 concernant la stabilité quantitative de solutions de problèmes de transport optimal par rapport aux données qui les définissent. En particulier, nous obtenons des estimations quantitatives de stabilité pour les applications de transport optimal par rapport à leurs mesures cibles et pour les barycentres de Wasserstein par rapport à leurs marginales. Au-delà des garanties qu’elles offrent pour les applications numériques et statistiques, ces estimations donnent également de nouvelles perspectives sur la géométrie de l’espace de Wasserstein-2 et sa plongeabilité dans un espace de Hilbert, que nous exploitons dans des applications en apprentissage automatique au Chapitre 8. Enfin, nous nous concentrons sur la variante du problème de transport optimal obtenue en ajoutant un terme de régularisation entropique dans (1) pondéré par un paramètre de *température*. Cette variante est connue pour être liée au problème de Schrödinger en physique statistique et nous donnons, dans un cadre spécifique, des estimations quantitatives de stabilité pour ses solutions par rapport au paramètre de température.

Plan détaillé et résumé des contributions

Partie I : Forte convexité du problème de transport optimal quadratique

Dans la première partie de cette thèse, nous obtenons des estimations de forte convexité pour le dual du problème de transport optimal quadratique sous différentes conditions et en utilisant différentes techniques, et nous établissons les relations entre ces estimations.

Chapitre 1 : Transport optimal quadratique et forte convexité du dual

Ce chapitre donne une introduction à la première partie de cette thèse. Dans la Section 1.1, nous rappelons les formulations de Monge et de Kantorovich du problème de transport optimal quadratique dans \mathbb{R}^d . Nous prouvons la formulation duale de Kantorovich (4), qui motive la définition de la *fonctionnelle de Kantorovich* $\mathcal{K}_\rho : \psi \mapsto \int \psi^* d\rho$ associée à une mesure source ρ . Cette fonctionnelle est convexe, et elle caractérise formellement un *potentiel de Kantorovich* ψ_μ entre ρ et μ solution au problème (4) par la condition de premier ordre

$$\partial\mathcal{K}_\rho(\psi_\mu) + \mu \ni 0 \quad \iff \quad \psi_\mu \in (\partial\mathcal{K}_\rho)^{-1}(-\mu). \quad (5)$$

Cette caractérisation nous conduit à étudier les propriétés (sous-)différentielles de \mathcal{K}_ρ . Nous accordons une attention particulière au cas où la mesure source ρ est absolument continue par rapport à la mesure de Lebesgue et nous remarquons dans ce cadre que la mesure signée $-(\nabla\psi^*)_{\#}\rho$ est (formellement) dans le sous-différentiel de \mathcal{K}_ρ à $\psi : \mathbb{R}^d \rightarrow \mathbb{R}$, c’est-à-dire pour tout $\psi : \mathbb{R}^d \rightarrow \mathbb{R}$,

$$-\langle (\nabla\psi^*)|\tilde{\psi} - \psi \rangle \leq \mathcal{K}_\rho(\tilde{\psi}) - \mathcal{K}_\rho(\psi). \quad (6)$$

Sous des hypothèses de compacité, nous montrons que $-(\nabla\psi^*)_{\#}\rho$ correspond en fait au *gradient* de \mathcal{K}_ρ à ψ (dans un sens à préciser). Ce fait, associé à la condition de

premier ordre (5), assure que dans un tel cadre, ψ_μ est solution de (4) si et seulement si $(\nabla\psi_\mu^*)_{\#\rho} = \mu$. Ceci nous permet de retrouver dans ce cadre le théorème de Brenier (Brenier, 1987) que nous avons mentionné plus haut. Ces idées illustrent comment l'étude de la fonctionnelle de Kantorovich peut aider à obtenir des informations qualitatives sur les solutions des problèmes de transport optimal de Monge et de Kantorovich.

Dans la Section 1.2, nous proposons de pousser plus loin l'étude de la fonctionnelle de Kantorovich afin d'obtenir également des informations quantitatives sur les solutions aux problèmes de transport optimal de Monge et de Kantorovich. En particulier, nous nous demandons dans cette section sous quelles conditions nous pouvons espérer des estimations de convexité forte pour \mathcal{K}_ρ , c'est-à-dire des estimations qui quantifient l'écart dans l'inégalité sous-différentielle (6). En raison de la condition d'optimalité (5), ces estimations de convexité forte pourraient être directement traduites en estimations de stabilité pour ψ_μ par rapport à μ . Cette question naturelle a déjà été abordée dans (Gigli, 2011) et (Hütter and Rigollet, 2021) dans le contexte de l'étude de la stabilité des applications de transport optimal, mais les estimations obtenues dans ces travaux n'étaient valables qu'au voisinage de potentiels très réguliers (fortement convexes). Comme nous le verrons juste après, ceci n'est pas optimal. Après avoir mentionné les conditions nécessaires sur la source ρ pour assurer la forte convexité de \mathcal{K}_ρ , nous annonçons la forme des estimations que nous calculons dans les Chapitres 2, 3 et 4. Essentiellement, ces estimations garantissent que pour une mesure source ρ absolument continue à support compact et convexe, on a pour tout potentiel de Kantorovich $\psi_\mu, \psi_\nu : \mathbb{R}^d \rightarrow \mathbb{R}$ entre ρ et $\mu, \nu \in \mathcal{P}_2(\mathbb{R}^d)$ la borne

$$\mathrm{Var}_\rho(\psi_\nu^* - \psi_\mu^*) \lesssim \mathcal{K}_\rho(\psi_\nu) - \mathcal{K}_\rho(\psi_\mu) + \langle (\nabla\psi_\mu^*) | \psi_\nu - \psi_\mu \rangle, \quad (7)$$

à condition que μ et ν satisfassent certaines hypothèses de moment. Le calcul d'estimées de la forme de (7) effectué dans les Chapitres 2, 3 et 4 repose principalement sur des inégalités géométriques et fonctionnelles bien connues (les inégalités de Brunn-Minkowsky, Brascamp-Lieb et Prékopa-Leindler). Nous rappelons dans la Section 1.3 les énoncés de ces inégalités et nous discutons succinctement de la manière dont elles ont déjà interféré avec le transport optimal dans des travaux antérieurs. Nous donnons enfin dans la Section 1.4 une extension de l'estimation (7) à des mesures sources ρ qui peuvent avoir un support non-convexe mais qui satisfont une inégalité de Poincaré-Wirtinger et de faibles hypothèses géométriques supplémentaires sur leur support.

Chapitre 2 : Une approche semi-discrète

Ce chapitre contient la preuve d'une première estimation de forte convexité de la forme (7). Cette estimation n'est valable que pour des mesures cibles μ, ν à support compact dans (7).

La preuve détaillée dans ce chapitre fonctionne par approximation. Nous supposons d'abord dans la Section 2.2 que les mesures de probabilité cibles μ, ν sont discrètement supportées sur un ensemble fini commun de N points. Ceci nous place dans le contexte du transport optimal semi-discrète, avec une source absolument continue ρ et des cibles discrètes μ, ν . Dans ce contexte, la fonctionnelle de Kantorovich \mathcal{K}_ρ peut être vue comme une fonction \mathcal{K}_ρ convexe de classe \mathcal{C}^2 sur \mathbb{R}^N pour laquelle le gradient et la hessienne sont connus. Dans la Section 2.3, nous tirons parti de la structure laplacienne de la matrice hessienne de \mathcal{K}_ρ pour donner une borne inférieure explicite à sa plus petite valeur propre non nulle. Nous déduisons ensuite de cette borne inférieure combinée à l'inégalité de

Brunn-Minkowski une estimation de forte convexité du type de (7), fonctionnant pour des cibles discrètes. En utilisant un argument d'approximation, nous généralisons finalement dans la Section 2.4 l'estimation de forte convexité de la Section 2.3 à toutes mesures cibles supportées de manière compacte.

Chapitre 3 : Une approche continue

Ce chapitre donne la preuve d'une deuxième estimation de convexité forte de la forme de (7). Cette seconde estimation est valable pour des mesures de probabilité cibles μ, ν dans (7) qui sont telles que les conjuguées convexes ψ_μ^*, ψ_ν^* de leurs potentiels de Kantorovich dans le transport optimal entre ρ et μ, ν sont bornées sur le support compact de ρ . D'après l'inégalité de Morrey et le plongement de Sobolev qui en résulte, c'est le cas par exemple chaque fois que μ et ν admettent un moment fini d'ordre $p > d$. L'estimation de ce chapitre couvre donc le cas des mesures cibles supportées de manière compacte et non-compacte et peut être vue comme une extension de l'estimation du Chapitre 2.

La preuve de cette deuxième estimation est également effectuée par approximation. Dans la Section 3.2, nous supposons que les cibles μ, ν dans (7) sont absolument continues et suffisamment régulières pour que les potentiels de Kantorovich ψ_μ, ψ_ν soient \mathcal{C}^2 et fortement convexes. Cette hypothèse nous permet de calculer la dérivée seconde de la fonctionnelle de Kantorovich en ψ_μ dans la direction $\psi_\nu - \psi_\mu$. Nous utilisons ensuite dans la Section 3.3 l'inégalité de Brascamp-Lieb (dite *de concentration*) pour obtenir une borne inférieure explicite sur la valeur de cette dérivée seconde, de laquelle nous déduisons une estimation de forte convexité du type de (7) fonctionnant pour des cibles μ et ν suffisamment régulières. Enfin, dans la Section 3.4, nous généralisons par densité l'estimation de forte convexité de la Section 3.3 à des cibles μ, ν qui sont seulement telles que leurs potentiels de Kantorovich admettent des conjuguées bornées.

Chapitre 4 : Une approche entropique

Nous considérons dans ce chapitre la variante *entropique* du problème de transport optimal de Kantorovich obtenue en ajoutant un terme de régularisation entropique à (1) :

$$\min_{\gamma \in \Gamma(\rho, \mu)} \int_{\mathbb{R}^d \times \mathbb{R}^d} \|x - y\|^2 d\gamma(x, y) + \varepsilon \text{KL}(\gamma | \rho \otimes \mu), \quad (8)$$

où $\varepsilon \geq 0$ est un paramètre de régularisation et KL désigne la divergence de Kullback-Leibler ou l'entropie relative. Cette variante est liée au problème de Schrödinger en physique statistique et elle a été popularisée ces dernières années dans les applications pour ses avantages computationnels et statistiques. Nous montrons dans ce chapitre que le dual de ce problème présente une fonctionnelle de Kantorovich *entropique* qui satisfait des estimations de forte convexité rappelant celles présentées dans les Chapitres 2 et 3.

Dans la Section 4.2, nous définissons la fonctionnelle de Kantorovich entropique et discutons le rôle joué par la mesure cible dans cette fonctionnelle. Cette discussion nous amène à faire, à nouveau, une hypothèse semi-discrète et à ne considérer que des mesures cibles discrètes. Sous cette hypothèse semi-discrète, nous calculons dans la Section 4.3 les dérivées première et seconde de la fonctionnelle de Kantorovich entropique. L'inégalité de Prékopa-Leindler est ensuite utilisée dans la Section 4.4 pour obtenir une borne inférieure sur la plus petite valeur propre de la hessienne de la fonctionnelle de Kantorovich

entropique semi-discrète, à partir de laquelle une estimation de forte convexité de type (7) est déduite. Nous exposons enfin dans la Section 4.5 comment cette estimation entropique peut être utilisée pour retrouver l'estimation de forte convexité du Chapitre 3.

Partie II : Conséquences pour la stabilité de solutions à des problèmes de transport optimal

Dans la deuxième partie de cette thèse, nous rassemblons des conséquences des estimations de forte convexité de la Partie I pour la stabilité quantitative de solutions à des problèmes de transport optimal par rapport à certaines des données qui les définissent.

Chapitre 5 : Stabilité quantitative des applications de transport optimal par rapport à la mesure cible

Dans ce chapitre, nous donnons des estimations de stabilité quantitative pour l'application de transport optimal quadratique entre une densité de probabilité fixe ρ et une mesure de probabilité μ sur \mathbb{R}^d , que nous désignons par T_μ , définie comme étant le minimiseur de

$$\min_{T_{\#\rho}=\mu} \int_{\mathbb{R}^d} \|T(x) - x\|^2 d\rho(x),$$

où $T_{\#\rho}$ est la mesure image de ρ par T . En supposant que la densité de la source ρ est bornée inférieurement et supérieurement sur un ensemble convexe compact, nous prouvons que l'application $\mu \mapsto T_\mu$ est bi-Hölder continue par rapport à la métrique de Wasserstein-2 sur de grandes familles de mesures de probabilité, comme l'ensemble des mesures de probabilité dont le moment d'ordre $p > d$ est borné par une certaine constante. Un peu plus précisément, pour un certain $p > d$ et tout $\mu, \nu \in \mathcal{P}_p(\mathbb{R}^d)$ qui admettent un p -ième moment borné par une constante commune, nous montrons que T_μ et T_ν satisfont des bornes du type

$$W_2(\mu, \nu) \leq \|T_\mu - T_\nu\|_{L^2(\rho, \mathbb{R}^d)} \lesssim W_2(\mu, \nu)^{\frac{p}{6p+16d}}. \quad (9)$$

où W_2 désigne la métrique Wasserstein-2. Ces estimations de stabilité montrent que la métrique $W_{2,\rho}(\mu, \nu) = \|T_\mu - T_\nu\|_{L^2(\rho, \mathbb{R}^d)}$ de *transport optimal linéarisé* (Wang et al., 2013) (pour laquelle les géodésiques correspondent à des *géodésiques généralisées* dans (Ambrosio et al., 2008)) est bi-Hölder équivalente à la distance de Wasserstein-2 sur de grands sous-ensembles de $\mathcal{P}_2(\mathbb{R}^d)$, ce qui justifie son utilisation dans les applications (voir le Chapitre 8 pour des exemples d'applications). Ce résultat répond aussi partiellement à la question de géométrie métrique de la plongeabilité de l'espace de Wasserstein $(\mathcal{P}_2(\mathbb{R}^d), W_2)$ dans un espace de Hilbert. Alors qu'un fort résultat négatif prouvé dans (Andoni et al., 2018) assure que l'espace de Wasserstein-2 ne peut pas être entièrement plongé dans un espace de Hilbert quelconque d'une manière bi-Hölder, nos estimations assurent qu'au moins de grands sous-ensembles de $(\mathcal{P}_2(\mathbb{R}^d), W_2)$ peuvent être explicitement plongés dans un espace L^2 avec un contrôle bi-Hölder de la distorsion induite sur la métrique. Une dernière interprétation possible de l'estimation (9) est en termes de l'interprétation riemannienne de l'espace (de dimension infinie) $(\mathcal{P}_2(\mathbb{R}^d), W_2)$: cette estimation peut être vue comme une estimation quantitative de la continuité de l'*application exponentielle inverse* $\mu \in (\mathcal{P}_2(\mathbb{R}^d), W_2) \mapsto T_\mu - \text{id} \in L^2(\rho, \mathbb{R}^d)$.

Afin de prouver les estimations du type de (9), nous tirons parti du fait que l'application de transport optimal T_μ entre une source fixe ρ et une cible μ s'écrit

$T_\mu = (\nabla\psi_\mu^*)_{\#}\rho$ où ψ_μ est un potentiel de Kantorovich pour le problème de transport entre ρ et μ , c'est-à-dire un minimiseur de (4). Dans la Section 5.2, nous déduisons d'abord des estimations de stabilité pour $\mu \mapsto \psi_\mu$ et $\mu \mapsto \psi_\mu^*$ comme conséquences directes des estimations de forte convexité calculées dans la Partie I. La stabilité de $\mu \mapsto T_\mu$ est ensuite obtenue dans la Section 5.3, en s'appuyant notamment sur une nouvelle inégalité de type Gagliardo-Nirenberg pour la différence de fonctions convexes prouvée dans la Section 5.4 qui peut présenter un intérêt indépendant.

Chapitre 6 : Stabilité quantitative des barycentres de Wasserstein par rapport à leurs marginales

Dans ce chapitre, nous donnons des estimations quantitatives de la stabilité des barycentres de Wasserstein par rapport à leurs marginales. Les barycentres de Wasserstein sont des moyennes de Fréchet dans les espaces de Wasserstein-2 : pour Ω un sous-ensemble compact de \mathbb{R}^d et \mathbb{P} une mesure de probabilité sur l'ensemble des mesures de probabilité sur Ω , c'est-à-dire $\mathbb{P} \in \mathcal{P}(\mathcal{P}(\Omega))$, un barycentre de Wasserstein de \mathbb{P} est défini comme un minimiseur $\mu_{\mathbb{P}}$ de

$$\min_{\mu \in \mathcal{P}(\Omega)} \frac{1}{2} \int_{\mathcal{P}(\Omega)} W_2^2(\rho, \mu) d\mathbb{P}(\rho). \quad (10)$$

De tels barycentres donnent des notions de *moyennes* de mesures de probabilité avec de fortes caractéristiques géométriques. Leur utilisation est de plus en plus populaire dans les domaines appliqués, tels que le traitement des images ou du langage ou en géométrie computationnelle. Cependant, dans ces domaines, la mesure de probabilité d'intérêt $\mathbb{P} \in \mathcal{P}(\mathcal{P}(\Omega))$ n'est souvent pas accessible dans sa totalité et le praticien peut avoir à traiter une approximation statistique ou informatique $\mathbb{Q} \in \mathcal{P}(\mathcal{P}(\Omega))$ à la place. Dans ce chapitre, nous quantifions l'effet de telles approximations sur les barycentres correspondants. Nous montrons que les barycentres de Wasserstein dépendent d'une manière Hölder-continue de leurs marginales sous des hypothèses de régularité relativement faibles. Formellement, notre résultat est le suivant: soit $\mathbb{P} \in \mathcal{P}(\mathcal{P}(\Omega))$ tel que \mathbb{P} charge un ensemble de mesures dont les fonctionnelles de Kantorovich associées (étudiées dans la Partie I) satisfont des estimations de forte convexité de type (7). Alors pour tout $\mathbb{Q} \in \mathcal{P}(\mathcal{P}(\Omega))$, le barycentre (unique) $\mu_{\mathbb{P}}$ de \mathbb{P} et tout barycentre $\mu_{\mathbb{Q}}$ de \mathbb{Q} satisfont

$$W_2(\mu_{\mathbb{P}}, \mu_{\mathbb{Q}}) \lesssim \mathcal{W}_1(\mathbb{P}, \mathbb{Q})^{1/6}, \quad (11)$$

où \mathcal{W}_1 désigne la métrique de Wasserstein-1 sur l'espace métrique $(\mathcal{P}(\Omega), W_2)$.

Avant de prouver cette estimation de stabilité, nous étudions dans la Section 6.1 sous quelles conditions nous pouvons espérer un quelconque résultat de stabilité. Lorsque la dimension ambiante d est supérieure à 2, nous montrons que des hypothèses de régularité (telles que l'absolue continuité et la connexité du support) doivent être faites sur certaines des mesures marginales de \mathbb{P} ou de \mathbb{Q} dans (11). Nous présentons ensuite la formulation duale de (10) ainsi que nos principales hypothèses et estimations, et nous donnons quelques conséquences immédiates mais utiles dans les applications. Nous montrons ensuite que la preuve des estimations du type (11) peut être déduite de deux estimations de stabilité : une première estimation de stabilité pour les solutions duales du problème (10) par rapport aux mesures marginales, donnée dans la Section 6.2, et une seconde estimation de stabilité pour l'opération poussé-en-avant par une application de transport optimal (pas nécessairement régulière), donnée dans la Section 6.3. La Section 6.4 donne enfin la preuve de la formulation duale de (10).

Chapitre 7 : Stabilité quantitative des potentiels de Schrödinger par rapport à la température dans le cadre semi-discret

Il est maintenant bien connu que pour un paramètre de régularisation suffisamment grand, le problème de transport optimal entropique est plus facile à résoudre que son équivalent non-régularisé et que les quantités associées ont un meilleur comportement statistique. Le praticien peut donc être intéressé par l'approximation des quantités de transport optimal classiques à l'aide de leurs versions entropiques. Ces pratiques nécessitent l'étude de l'erreur d'approximation qu'elles induisent. Dans le cas du transport optimal discret (où la source et la cible sont discrètes), la littérature assure des vitesses de convergence rapides et non-asymptotiques des quantités entropiques vers leurs analogues classiques lorsque le paramètre de régularisation tend vers zéro. Le cas du transport optimal semi-discret est cependant moins avancé : dans ce cadre, rien n'était connu quantitativement de l'effet de la régularisation entropique jusqu'au travail récent de (Altschuler et al., 2022) où des bornes asymptotiques ont été décrites. Dans ce chapitre, nous améliorons ces limites pour obtenir des limites non asymptotiques et presque optimales. Nous nous donnons une densité de probabilité source fixe ρ supportée sur un ensemble compact et convexe \mathcal{X} et une mesure cible discrète fixe $\mu = \sum_{i=1}^N \mu_i \delta_{y_i}$ supportée sur un ensemble fini $\mathcal{Y} = \{y_i\}_{1 \leq i \leq N}$. Pour ces source et cible, nous considérons pour tout $\varepsilon \geq 0$ le problème de transport optimal régularisé par l'entropie (8) déjà étudié dans le Chapitre 4 et notons γ^ε sa solution (unique) ainsi que $\psi^\varepsilon \in \mathbb{R}^N$ la solution de son problème dual vérifiant $\sum_{i=1}^N \psi_i^\varepsilon = 0$ (une telle solution duale est souvent appelée potentiel de Schrödinger ou de Sinkhorn). Nous montrons que dès que la densité source ρ est α -Hölder continue pour un certain $\alpha > 0$, l'application $\varepsilon \mapsto \psi^\varepsilon$ est *mieux que Lipschitz* : pour toute $0 < \varepsilon' \leq \varepsilon \leq 1$ et $\alpha' \in (0, 1)$, nous assurons

$$\|\psi^\varepsilon - \psi^{\varepsilon'}\|_\infty \lesssim \varepsilon^{\alpha'} (\varepsilon - \varepsilon').$$

Ce fait peut constituer un premier pas vers une justification mathématique de l'heuristique d' ε -*scaling* utilisée pour la résolution numérique du transport optimal semi-discret régularisé, où ε est progressivement diminué au cours des itérations d'un algorithme qui vise à résoudre le problème dual de (8). En laissant ε' tendre vers zéro, cette borne assure un taux de convergence superlinéaire en ε de ψ^ε vers la solution non régularisée ψ^0 . Ce résultat assure également qu'il existe une fonction $c : \mathcal{X} \rightarrow \mathbb{R}_+$ qui est strictement positive ρ -presque partout et qui vérifie pour tout $x \in \mathcal{X}$ et $y \in \mathcal{Y}$

$$|\gamma^\varepsilon(x, y) - \gamma^0(x, y)| \lesssim e^{-c(x)/\varepsilon}.$$

Ce résultat peut être considéré, dans le cadre semi-discret, comme une version non-asymptotique du principe de grandes déviations montré récemment par (Bernton et al., 2022). Enfin, les bornes présentées impliquent également un développement non asymptotique et optimal de la différence entre les coûts entropiques et non régularisés. A savoir, en notant

$$W_{2,\varepsilon}(\rho, \mu) = \left(\int_{\mathcal{X} \times \mathcal{Y}} \|x - y\|^2 d\gamma^\varepsilon(x, y) \right)^{1/2}$$

la distance de Wasserstein *approximée* entre ρ et μ , nous prouvons qu'il existe une constante explicite $C(\rho, \mu)$ qui ne dépend que de ρ et μ telle que pour tout $\alpha' \in (0, 1)$,

$$|W_{2,\varepsilon}^2(\rho, \mu) - W_2^2(\rho, \mu) - \varepsilon^2 C(\rho, \mu)| \lesssim \varepsilon^{2+\alpha'}.$$

Afin de prouver ces estimations, nous rappelons d’abord dans la Section 7.2 des éléments de transport optimal semi-discret (entropique) des Chapitres 2 et 4 et énonçons nos principaux résultats. La Section 7.3 donne l’EDO à partir de laquelle commence la preuve de notre principale estimation. Cette EDO présente deux termes qui impliquent tous deux la fonctionnelle de Kantorovich entropique semi-discrète introduite dans le Chapitre 4. L’estimation de forte convexité de ce chapitre ainsi qu’une autre estimation obtenue dans la Section 7.4 permettent alors de prouver notre principale estimation. Les deux corollaires de cette estimation donnant les taux de convergence de ψ^ε vers ψ^0 et de $W_{2,\varepsilon}$ vers W_2 sont respectivement prouvés dans la Section 7.5 et la Section 7.6. La Section 7.7 illustre enfin nos résultats théoriques avec des exemples numériques simples uni-dimensionnels.

Partie III : Applications numériques : le cadre du transport optimal linéarisé

Chapitre 8 : Transport optimal linéarisé et applications

Ce dernier chapitre rassemble des illustrations numériques et des expériences dans le cadre *transport optimal linéarisé* (LOT) de (Wang et al., 2013), un cadre de transport optimal approximatif qui est analysé dans une certaine mesure dans le Chapitre 5. Dans la Section 8.2, nous illustrons les résultats théoriques du Chapitre 8.2 et observons la distorsion de la métrique induite par le plongement LOT sur certains exemples bi-dimensionnels. Nous mentionnons également comment le plongement LOT peut être utilisé pour approcher les barycentres dans l’espace de Wasserstein-2. Ensuite, dans la Section 8.3, nous donnons deux exemples d’extensions des méthodes classiques d’analyse de données hilbertiennes à des mesures de probabilité dans le cadre du LOT. Ces extensions concernent des problèmes de K -moyennes et d’apprentissage de dictionnaires dans l’espace de Wasserstein-2.

Appendice

Chapitre A : Faits relatifs au transport optimal

Cette courte annexe rassemble certains faits relatifs au transport optimal qui sont utiles dans cette thèse mais qui ne sont pas traités dans le Chapitre 1.

Publications

- *Quantitative stability of optimal transport maps and linearization of the 2-Wasserstein space.* Q. Mérigot, A. Delalande, F. Chazal. Proceedings of the 23rd International Conference on Artificial Intelligence and Statistics (AISTATS), PMLR 108:3186-3196, 2020.
- *Quantitative stability of optimal transport maps under variations of the target measure.* A. Delalande, Q. Mérigot. Under revision at the Duke Mathematical Journal.
- *Nearly tight convergence bounds for semi-discrete entropic optimal transport.* A. Delalande. Proceedings of the 25th International Conference on Artificial Intelligence and Statistics (AISTATS), PMLR 151:1619-1642, 2022.

- *Quantitative stability of barycenters in the Wasserstein space.* G. Carlier, A. Delalande, Q. Mérigot. Preprint.

Ressources pour les expériences numériques

- https://github.com/alex-delalande/stability_ot_maps_and_linearization_wasserstein_space
- <https://github.com/alex-delalande/potentials-entropic-sd-ot>
- https://github.com/alex-delalande/linearized_wasserstein_dictionary_learning

General introduction

The optimal transport problem corresponds to the main object of study of the present thesis. In plain language, this problem can be formulated as follows:

Given two distributions of a same amount of mass over some space and given the knowledge of the cost of transferring a unit of mass from any location to any other location, how can we transport all the mass from the first distribution to the second in the cheapest possible way?

Such a problem has the ability to model practical issues: the mass aforementioned may be instantiated by a crowd of people, or by a collection of gas particles, or by a set of commodities; and the cost may correspond to a distance (to measure a traveling effort), or to the square of a distance (to measure a kinetic energy), or simply to a financial cost. In the presence of a mathematical problem that aims at modeling physical phenomena, it is relevant to verify whether this problem is *well-posed* in the sense of Hadamard (Hadamard, 1902). This informal notion gathers some of the desirable features that a *natural* problem should present and that should help its resolution. A given problem is said to be well-posed if:

- (i) it admits a solution;
- (ii) this solution is unique;
- (iii) this solution depends continuously on the data of the problem.

Quoting Evans in the introduction to his book *Partial Differential Equations* (Evans, 2010), condition (iii) "is particularly important for problems arising from physical applications: we would prefer that our (unique) solution changes only a little when the conditions specifying the problem change a little. For many problems, on the other hand, uniqueness is not to be expected."

In its modern formulation, the optimal transport problem is overall well-posed. A solution to this problem corresponds to a concrete plan to indeed morph all the mass positioned in a first configuration into a second configuration while achieving the minimal overall transportation cost. We shall see below that the existence of solution property (i) is by now perfectly understood and shown to hold in most interesting cases. The question (ii) of the uniqueness of the solution has also been extensively studied and many of the instances where the solution is indeed unique, or conversely where it is not likely to be unique, have been described. The study of the last property (iii) is however less advanced. There are in general abstract stability guarantees that ensure that optimal transport solutions do change continuously with the mass distributions that define the problem. However, except in very rare cases, these guarantees are not quantitative: we do not know in general how a given change in these distributions impacts the corresponding optimal transport solutions. Because modern days applications almost always require the practitioners to deal with statistical or computational approximations of the data of interest, this lack of quantitative guarantees is problematic. The aim of this thesis is to work towards closing this gap.

Optimal transport theory and applications: a (very) brief overview. The optimal transport problem was first introduced by Monge in 1781 with military and engineering applications in mind (Monge, 1781). He formulated the general problem of finding the cheapest way to transport a given amount of dirt from an extraction site to a building site, the cost of transport of each *molecule* of dirt being proportional to the distance it travels. The study of this problem led him to the discovery of important concepts in the geometry of surfaces, but the problem was left largely unresolved. The optimal transport problem was revived with Kantorovich, who gave in 1942 (Kantorovich, 1942) its modern formulation as a linear program. In raw mathematical terms, Kantorovich's problem may be described as follows. Consider some space Ω , typically a compact Polish space, and two probability measures ρ, μ in $\mathcal{P}(\Omega)$ that each represent a distribution of mass. Then, given a function $c(x, y)$ that represents the cost of transferring a unit of mass from a location x in Ω to a location y in Ω , solve

$$\min_{\gamma \in \Gamma(\rho, \mu)} \int_{\Omega \times \Omega} c(x, y) d\gamma(x, y), \quad (12)$$

where $\Gamma(\rho, \mu)$ is the set of *transport plans* between ρ and μ , that is the set of probability measures over $\Omega \times \Omega$ with first marginal ρ and second marginal μ . A candidate γ in $\Gamma(\rho, \mu)$ proposes a plan to transport the mass from ρ to μ by sending a portion $d\gamma(x, y)$ of the mass $d\rho(x)$ from a source location x to a target location y . The linearity of Kantorovich's formulation allowed him to ensure, under mild assumptions on c , the existence of solutions to his problem in compact metric spaces as well as to establish a dual formulation and optimality conditions. He and other authors soon understood that the value of the optimal transport cost (12) between two measures ρ and μ could give a quantitative idea of how *similar* ρ and μ are. It was shown that when $c(x, y) = d_\Omega(x, y)^p$ is the p -th power of a distance d_Ω on Ω for some $p \geq 1$, the value of (12) corresponds itself to the p -th power of a distance between ρ and μ . This distance, generally called the p -Wasserstein distance² and denoted W_p , endows the set of probability measures $\mathcal{P}(\Omega)$ with a rich geometric structure *lifted* from the base space Ω . For instance, a compact metric space (Ω, d_Ω) embeds isometrically into the p -Wasserstein space $(\mathcal{P}(\Omega), W_p)$ through the mapping $x \mapsto \delta_x$ (where δ_x denotes the Dirac mass at x). The geometry provided by Wasserstein metrics on spaces of probability measures has proven to be very convenient, both for theoretical and applied considerations. The unique quadratic case $p = 2$ on $\Omega = \mathbb{R}^d$ has by itself produced a very substantial theory. It has allowed to define notions of interpolations (McCann, 1997) or barycenters (Agueh and Carlier, 2011) on families of probability measures coming with strong geometrical flavors (see Figure 4). More generally, the 2-Wasserstein geometry revealed itself to give a physically relevant Riemannian structure to spaces of probability measures, in which some well known evolution PDEs (such as the Fokker-Planck or porous medium equations) could be expressed as gradient flows of well-chosen energy functionals on the space of probability distributions (Otto, 1998; Jordan et al., 1998; Otto, 2001). Other important evolution PDEs found a variational formulation in Wasserstein spaces, such as Euler equations in fluid dynamics (Brenier, 1989, 1999).

In parallel to the development of its theory, the optimal transport problem has made many successful incursions in applications. Kantorovich introduced his linear program (12) to model common problems of resource allocation arising in economics, a field

²the attribution of this name to this distance is often questioned. In Kantorovich's works, the first appearance of such notion is in his joint work with Rubinstein (Kantorovich and Rubinstein, 1958) for the case $p = 1$. We refer to the bibliographical notes in Chapter 6 of (Villani, 2008) for more details.




Figure 4: (Top) Linear interpolation $(1 - t)\rho_0 + t\rho_1$ and (Bottom) Displacement interpolation $\rho_t := \arg \min_{\mu \in \mathcal{P}(\mathbb{R})} (1 - t)W_2^2(\rho_0, \mu) + tW_2^2(\rho_1, \mu)$ of two Gaussian distributions ρ_0, ρ_1 on $(\mathbb{R}, |\cdot|)$, with respective mean 0 and 15 and unit variance each. In many applications, the *horizontal* movement of the displacement interpolant is preferable to the *vertical* movement of its linear counterpart.

where optimal transport still draws vivid interest (Galichon, 2016). In these applications, the considered resources are often of discrete nature and the probability measures ρ and μ can be taken finitely and discretely supported. In such a setting, problem (12) corresponds to a classical finite-dimensional linear program, a problem that was soon solved numerically with Dantzig’s simplex algorithm (Dantzig, 1949, 1951) and in more efficient manners with algorithms for min-cost flow problems (Ford and Fulkerson, 1962; Goldberg and Tarjan, 1989). The discrete version of problem (12) is also closely linked to the assignment problem, which was solved efficiently with Bertsekas’ auction algorithm (Bertsekas, 1981; Bertsekas and Eckstein, 1988). Since the years 2000, the optimal transport problem has also been increasingly used to solve various tasks of shape, image and video processing such as registration (Haker et al., 2004), flicker reduction (Delon, 2006), color transfer (Pitié et al., 2007; Bonneel et al., 2016), denoising (Lellmann et al., 2014) or segmentation (Rabin and Papadakis, 2015). In machine learning, optimal transport was used for image retrieval (Rubner et al., 2000), semi-supervised learning (Solomon et al., 2014), generative modeling (Arjovsky et al., 2017), domain adaptation (Courty et al., 2017) or distributionally robust optimization (Kuhn et al., 2019). The number of these applications vastly increased following the computational advances due to (Cuturi, 2013), see (Peyré and Cuturi, 2019) for more references. Finally, other notable applications of optimal transport can be found in quantum chemistry (Buttazzo et al., 2012; Cotar et al., 2013), optics design (Oliker, 2003; Caffarelli and Oliker, 2008) and in statistics, where it has helped to extend notions of quantiles to multivariate random variables (Carlier et al., 2016; Chernozhukov et al., 2017), to build efficient density estimators in geometric inference (Weed and Berthet, 2019; Divol, 2022) and to analyze the convergence of sampling algorithms such as the Langevin Monte Carlo algorithm (Dalalyan, 2017; Bernton, 2018) or the Stein variational gradient descent algorithm (Korba et al., 2020).

Well-posedness of optimal transport problems. The strong ability of optimal transport to model physical phenomena raises urgently the question of its well-posedness: can we hope for any solution, is the solution unique and does it respond continuously to modifications of the problem data?

We have already mentioned that Kantorovich proved the existence of solutions to (12) in general cases. This existence result was even more generalized in (Kellerer, 1984),




Figure 5: In \mathbb{R}^2 , consider $\rho = \frac{1}{2}(\delta_{(-1,0)} + \delta_{(1,0)})$ and $\mu = \frac{1}{2}(\delta_{(0,-1)} + \delta_{(0,1)})$. The optimal transport between ρ and μ with the Euclidean distance as ground cost ($c(x, y) = \|x - y\|$) is achieved by $\gamma^0 = \frac{1}{2}(\delta_{(-1,0) \times (0,1)} + \delta_{(1,0) \times (0,-1)})$ (left) as well as by $\gamma^1 = \frac{1}{2}(\delta_{(-1,0) \times (0,-1)} + \delta_{(1,0) \times (0,1)})$ (right), or by any of the convex combination $(1 - t)\gamma^0 + t\gamma^1$ for $t \in [0, 1]$.

and one can expect that (12) admits solutions for instance whenever Ω is Polish and c is lower semicontinuous and lower-bounded.

The uniqueness of a solution to (12) is not to be expected in general (see Figure 5 for an example). There are however interesting particular cases where uniqueness of the solution holds. The most famous of these cases is undoubtedly due to Brenier (Brenier, 1987), who showed for Ω a compact subset of \mathbb{R}^d and $c(x, y) = \|x - y\|^2$ that whenever the source measure ρ is absolutely continuous with respect to the Lebesgue measure, the solution to (12) is unique and, more importantly, it is supported on the graph of the gradient of a convex function (see also (Knott and Smith, 1984; Smith and Knott, 1987; Rüschemdorf and Rachev, 1990)). Incidentally, this characterization allowed to adopt the following PDE point of view on the optimal transport problem: whenever ρ and μ admit densities (denoted with the same letters), a smooth and strictly convex function ϕ whose gradient's graph supports the optimal transport solution between ρ and μ must verify for all x in Ω the change of variable formula

$$\det(D^2\phi(x))\mu(\nabla\phi(x)) = \rho(x). \quad (13)$$

This corresponds to a Monge-Ampère equation in ϕ , whose solution actually provides the solution to the optimal transport problem between ρ and μ under suitable boundary conditions. Brenier's result was then generalized to more general costs and domains, see for instance (Gangbo and McCann, 1996; Trudinger and Wang, 2001; McCann, 2001; Caffarelli et al., 2002; Bernard and Buffoni, 2007; Fathi and Figalli, 2010).

The stability of optimal transport solutions with respect to the data that defines them is established in general cases. For instance, Theorem 5.19 of (Villani, 2008) ensures that for Ω a Polish space and c a continuous and bounded cost function, the weak convergence of source and target measures ρ_n, μ_n in $\mathcal{P}(\Omega)$ to respective limits ρ, μ in $\mathcal{P}(\Omega)$ entails, up to a subsequence, the weak convergence of optimal transport solutions γ_n between ρ_n and μ_n to an optimal transport solution γ between ρ and μ . Other results also ensure in general the stability of other optimal transport quantities such as the above mentioned interpolants and barycenters in Wasserstein spaces. These guarantees are not anecdotal: they ensure for instance that an approximate optimal transport plan between two measures ρ and μ is given by an optimal transport plan $\hat{\gamma}$ between two approximations $\hat{\rho}, \hat{\mu}$ of ρ, μ . This is particularly useful in applications, where opting for approximations $\hat{\rho}, \hat{\mu}$ of ρ, μ can be necessary either because of computational limitations, or in a statistical




Figure 6: Histogram representation of the displacement interpolation $\hat{\rho}_t^n := \arg \min_{\mu \in \mathcal{P}(\mathbb{R})} (1-t)W_2^2(\hat{\rho}_0^n, \mu) + tW_2^2(\hat{\rho}_1^n, \mu)$, where for $k \in \{0, 1\}$, $\hat{\rho}_k^n = \frac{1}{n} \sum_{i=1}^n \delta_{x_k^i}$ with $(x_k^i)_{1 \leq i \leq n} \sim \rho_k$ for $n = 2000$ and ρ_k as in Figure 4.

context where only samples of the measures of interest are available. In such applications however, the practitioner may need to have quantitative guarantees on the committed approximations. Consider for instance the displacement interpolation displayed in Figure 6. This interpolation is a *statistical approximation* of the one displayed in the bottom row of Figure 4, obtained by approximating the end measures ρ_0 and ρ_1 with empirical measures $\hat{\rho}_0^n$ and $\hat{\rho}_1^n$ built from samples before computing their interpolation. In typical settings, the statistician is able to say something on the quality of the approximations of ρ_0 and ρ_1 in Wasserstein distance, i.e. bounds (in expectation or with high probability) on $W_2(\rho_0, \hat{\rho}_0^n)$ and $W_2(\rho_1, \hat{\rho}_1^n)$ are available. For downstream applications, it can then be important to know whether these quality guarantees are transmitted to the interpolants, i.e. whether the distance $W_2(\rho_t, \hat{\rho}_t^n)$ for some $t \in (0, 1)$ can be bounded in terms of $W_2(\rho_0, \hat{\rho}_0^n)$ and $W_2(\rho_1, \hat{\rho}_1^n)$. In the one-dimensional setting of Figures 4 and 6, the answer is positive and one can always ensure the following Lipschitz behavior:

$$W_p(\rho_t, \hat{\rho}_t^n) \leq (1-t)W_p(\rho_0, \hat{\rho}_0^n) + tW_p(\rho_1, \hat{\rho}_1^n).$$

However, this bound is specific to $\Omega = \mathbb{R}^d$ with $d = 1$ and there is no similar quantitative guarantee whenever $d \geq 2$. This raises the question of the general quantitative stability of the optimal transport problem.

For some elliptic PDEs, quantitative stability of solutions can be deduced from strong ellipticity. Consider for instance on a smooth bounded domain Ω of \mathbb{R}^d the Poisson equation

$$\Delta \phi = f, \tag{14}$$

with zero Dirichlet boundary condition ($\phi = 0$ on $\partial\Omega$). In this equation Δ denotes the Laplace operator, whose ellipticity gives the Poincaré inequality on Ω . This inequality ensures in particular that there exists a constant C such that for any weak solutions $\phi, \hat{\phi} \in H^1(\Omega)$ of (14) with respective second members f, \hat{f} , one has

$$\|\phi - \hat{\phi}\|_{L^2(\Omega)}^2 \leq C \|\nabla \phi - \nabla \hat{\phi}\|_{L^2(\Omega)}^2 = -C \int_{\Omega} (\phi - \hat{\phi})(f - \hat{f}) dx \leq C \|\phi - \hat{\phi}\|_{L^2(\Omega)} \|f - \hat{f}\|_{L^2(\Omega)},$$

so that $\|\phi - \hat{\phi}\|_{L^2(\Omega)} \leq C \|f - \hat{f}\|_{L^2(\Omega)}$. Such an inequality quantifies precisely the effect of a perturbation of the input data f on the corresponding solution ϕ in (14). Unfortunately, this ellipticity approach cannot be readily applied to the optimal transport problem. For instance, we have mentioned that in the quadratic and Euclidean setting (which corresponds to the most studied and arguably the simplest setting), the optimal transport problem could be reformulated in some cases in terms of the Monge-Ampère equation (13). In general, this equation is merely degenerate elliptic, and strong ellipticity only holds when the unknown ϕ is smooth and strongly convex, which is rarely the case in practice. This makes the question of the quantitative stability of the optimal transport problem a particularly difficult one.

Main contributions. In this thesis, we follow a classical approach in optimal transportation theory that consists in studying the dual problem of (12) to infer qualitative and quantitative information about optimal transport solutions. We only focus on the simpler Euclidean and quadratic setting (i.e. $\Omega = \mathbb{R}^d$ and $c(x, y) = \|x - y\|^2$), leaving the generalizations of our results to future work. In this setting, the dual problem of (12) essentially corresponds to the minimization problem

$$\min_{\psi: \mathbb{R}^d \rightarrow \bar{\mathbb{R}}} \int \psi^* d\rho + \int \psi d\mu, \quad (15)$$

where ψ^* denotes the convex conjugate of the *potential* ψ .

Instead of directly looking for quantitative stability estimates for (12), we first look for quantitative stability estimates for (15). The functional $\psi \mapsto \int \psi^* d\rho$ appearing in this dual problem, which we call the *Kantorovich functional* associated to a source ρ , is convex. As such, the derivation of stability estimates for minimizers of (15) can be carried out through the estimation of the strong convexity of this functional. In chapters 2–4, we derive explicit estimates of the strong convexity of the Kantorovich functional, mainly relying on the Brunn-Minkowski, Brascamp-Lieb and Prékopa-Leindler inequalities. It was already understood, since the seminal work of (McCann, 1997), that these well-known geometric and functional inequalities are linked to the optimal transport problem since they can be deduced from the geodesic convexity of some energy functionals on the 2-Wasserstein space. This thesis reinforces this link, somewhat in an opposite direction, by using these inequalities to quantify the strong convexity of the dual quadratic optimal transport problem.

Then, we gather in chapters 5–7 consequences of the strong convexity estimates of chapters 2–4 regarding the quantitative stability of optimal transport solutions with respect to the data that defines them. In particular, we derive quantitative stability estimates for optimal transport maps with respect to their target measures and for Wasserstein barycenters with respect to their marginals. Beyond the guarantees they offer for numerical and statistical applications, these estimates also give new insights about the geometry of the 2-Wasserstein space and its embeddability in Hilbert spaces that we leverage in machine learning applications in Chapter 8. Finally, we focus on the variant of the optimal transport problem obtained by adding an entropic regularization term in (12) weighted by a *temperature* parameter. This variant is known to be related to the Schrödinger problem in statistical physics and we derive, in specific settings, quantitative stability estimates for its solutions with respect to the temperature parameter.

Detailed outline and summary of contributions

Part I: Strong convexity of the quadratic optimal transport problem

In the first part of this thesis, we derive strong-convexity estimates for the dual quadratic optimal transport problem under different conditions and using different techniques, and we establish the relationships between these estimates.

Chapter 1: Quadratic optimal transport and strong convexity of the dual

This chapter gives an introduction to the first part of this thesis. In Section 1.1, we recall the Monge and Kantorovich formulations of the quadratic optimal transport problem in \mathbb{R}^d . We prove the Kantorovich duality formula (15), that motivates the definition of the *Kantorovich functional* $\mathcal{K}_\rho : \psi \mapsto \int \psi^* d\rho$ associated to a source measure ρ . This functional is convex, and it formally characterizes a *Kantorovich potential* ψ_μ between ρ and μ solution to problem (15) through the first order condition

$$\partial\mathcal{K}_\rho(\psi_\mu) + \mu \ni 0 \quad \iff \quad \psi_\mu \in (\partial\mathcal{K}_\rho)^{-1}(-\mu). \quad (16)$$

This characterization leads us to study the (sub)differential properties of \mathcal{K}_ρ . We pay particular attention to the case where the source measure ρ is absolutely continuous with respect to the Lebesgue measure and we notice in such setting that the signed measure $-(\nabla\psi^*)_{\#}\rho$ is (formally) in the subdifferential of \mathcal{K}_ρ at $\psi : \mathbb{R}^d \rightarrow \bar{\mathbb{R}}$, that is for any $\tilde{\psi} : \mathbb{R}^d \rightarrow \bar{\mathbb{R}}$,

$$-\langle(\nabla\psi^*)|\tilde{\psi} - \psi\rangle \leq \mathcal{K}_\rho(\tilde{\psi}) - \mathcal{K}_\rho(\psi). \quad (17)$$

Under compactness assumptions, we show that $-(\nabla\psi^*)_{\#}\rho$ actually corresponds to the *gradient* of \mathcal{K}_ρ at ψ (in a sense to be made precise). This fact together with the first order condition (16) ensure that in such setting, ψ_μ is solution to (15) if and only if $(\nabla\psi_\mu^*)_{\#}\rho = \mu$. This allows us to recover in this setting Brenier's theorem (Brenier, 1987) that we mentioned above. These ideas illustrate how the study of the Kantorovich functional can help get qualitative information about the solutions to Monge and Kantorovich optimal transport problems.

In Section 1.2, we propose to push further the study of the Kantorovich functional in order to also get quantitative information about the solutions to Monge and Kantorovich optimal transport problems. In particular, we wonder in this section under which conditions we can hope for strong convexity estimates for \mathcal{K}_ρ , i.e. estimates that quantify the gap in the subdifferential inequality (17). Because of the optimality condition (16), such strong convexity estimates could be directly translated into stability estimates for ψ_μ with respect to μ . This natural question has already been tackled in (Gigli, 2011) and (Hütter and Rigollet, 2021) in the context of the study of the stability of optimal transport maps, but the estimates derived in these works were only valid near very regular (i.e. strongly convex) potentials. As we shall see right after, this is not optimal. After mentioning necessary conditions on the source ρ to ensure strong convexity of \mathcal{K}_ρ , we announce the form of the estimates that we derive in the subsequent chapters 2, 3 and 4. In broad terms, these estimates ensure that for an absolutely continuous, compactly and convexly supported source measure ρ , one has for any Kantorovich potentials $\psi_\mu, \psi_\nu : \mathbb{R}^d \rightarrow \bar{\mathbb{R}}$ between ρ and $\mu, \nu \in \mathcal{P}_2(\mathbb{R}^d)$ the bound

$$\text{Var}_\rho(\psi_\nu^* - \psi_\mu^*) \lesssim \mathcal{K}_\rho(\psi_\nu) - \mathcal{K}_\rho(\psi_\mu) + \langle(\nabla\psi_\mu^*)|\psi_\nu - \psi_\mu\rangle, \quad (18)$$

provided that μ and ν satisfy some moment assumptions. The derivations of estimates of the form of (18) carried out in chapters 2, 3 and 4 mainly rely on well-known geometric and functional inequalities (the Brunn-Minkowsky, Brascamp-Lieb and Prékopa-Leindler inequalities). We recall in Section 1.3 the statements of these inequalities and we discuss succinctly how they have already interfered with optimal transport in previous works. We finally give in Section 1.4 an extension of estimate (18) to source measures ρ that may not be convexly supported but that satisfy a Poincaré-Wirtinger inequality and additional mild geometric assumptions on their support.

Chapter 2: A semi-discrete approach

This chapter contains the proof of a first strong convexity estimate of the form of (18). This estimate is only valid for compactly supported target measures μ, ν in (18).

The proof derived in this chapter works by approximation arguments. We first assume in Section 2.2 that the target probability measures μ, ν are discretely supported on a common finite set of N points. This places ourselves in the context of semi-discrete optimal transport, with an absolutely continuous source ρ and discrete targets μ, ν . In this context, the Kantorovich functional \mathcal{K}_ρ can be seen as a \mathcal{C}^2 convex function \mathcal{K}_ρ on \mathbb{R}^N for which the gradient and Hessian are known. In Section 2.3, we leverage the Laplacian structure of the Hessian matrix of \mathcal{K}_ρ to give an explicit lower-bound on its smallest non-zero eigenvalue. We then deduce from this lower-bound combined with the Brunn-Minkowski inequality a strong convexity estimate of type (18) working for discrete targets. Using an approximation argument, we finally generalize in Section 2.4 the strong-convexity estimate of Section 2.3 to any compactly supported target measures.

Chapter 3: A continuous approach

This chapter gives the proof of a second strong convexity estimate of the form of (18). This second estimate is valid for target probability measures μ, ν in (18) that are such that the convex conjugates ψ_μ^*, ψ_ν^* of their Kantorovich potentials in the optimal transport between ρ and μ, ν are bounded on the compact support of ρ . From Morrey's inequality and the resulting Sobolev embedding, this is the case for instance whenever μ and ν admit finite moments of order $p > d$. The estimate of this chapter thus covers the case of compactly and non-compactly supported target measures and can be seen as an extension of the estimate of Chapter 2.

The proof of this second estimate is also carried out by approximation. In Section 3.2, we assume that the targets μ, ν in (18) are absolutely continuous and regular enough so that the Kantorovich potentials ψ_μ, ψ_ν are smooth and strongly-convex. This assumption allows us to compute the second order derivative of the Kantorovich functional at ψ_μ in the direction $\psi_\nu - \psi_\mu$. We then use in Section 3.3 the Brascamp-Lieb concentration inequality to get an explicit lower-bound on the value of this second derivative, from which we deduce a strong-convexity estimate of type (18) working for regular enough targets μ and ν . Finally in Section 3.4, we generalize with density arguments the strong convexity estimate of Section 3.3 to targets μ, ν that are only such that their Kantorovich potentials satisfy some boundedness assumptions.

Chapter 4: An entropic approach

We consider in this chapter the *entropic* variant of the Kantorovich optimal transport problem obtained by adding an entropic regularization term to (12):

$$\min_{\gamma \in \Gamma(\rho, \mu)} \int_{\mathbb{R}^d \times \mathbb{R}^d} \|x - y\|^2 d\gamma(x, y) + \varepsilon \text{KL}(\gamma | \rho \otimes \mu), \quad (19)$$

where $\varepsilon \geq 0$ is a regularization parameter and KL denotes the Kullback-Leibler divergence or relative entropy. This variant is linked to the Schrödinger problem in statistical physics and it has been popularized in the recent years in applied fields for its computational and statistical advantages. We show in this chapter that the dual of this problem

features an *entropic* Kantorovich functional that enjoys strong-convexity estimates reminiscent of the ones presented in chapters 2 and 3.

In Section 4.2, we define the entropic Kantorovich functional and discuss the role played by the target measure in this functional. This discussion leads us to make, again, a semi-discrete assumption and consider only discrete target measures. Under this semi-discrete assumption, we compute in Section 4.3 the first and second derivatives of the entropic Kantorovich functional. The Prékopa-Leindler inequality is then used in Section 4.4 to derive a lower-bound on the smallest eigenvalue of the Hessian of the semi-discrete entropic Kantorovich functional, from which a strong convexity estimate of type (18) is derived. We finally expose in Section 4.5 how this entropic estimate can be used to recover the strong convexity estimate of Chapter 3.

Part II: Consequences for the stability of solutions to optimal transport problems

In the second part of this thesis, we collect consequences of the strong-convexity estimates of Part I regarding the quantitative stability of solutions to optimal transport problems with respect to some of the data that defines them.

Chapter 5: Quantitative stability of optimal transport maps with respect to the target measure

We derive in this chapter quantitative stability estimates for the quadratic optimal transport map between a fixed probability density ρ and a probability measure μ on \mathbb{R}^d , which we denote T_μ , defined as being the minimizer of

$$\min_{T_{\#\rho}=\mu} \int_{\mathbb{R}^d} \|T(x) - x\|^2 d\rho(x),$$

where $T_{\#\rho}$ is the image measure of ρ by T . Assuming that the source density ρ is bounded from above and below on a compact convex set, we prove that the map $\mu \mapsto T_\mu$ is bi-Hölder continuous with respect to the 2-Wasserstein metric on large families of probability measures, such as the set of probability measures whose moment of order $p > d$ is bounded by some constant. A bit more precisely, for some $p > d$ and any $\mu, \nu \in \mathcal{P}_p(\mathbb{R}^d)$ that admit a p -th moment upper bounded by a common constant, we show that T_μ and T_ν satisfy bounds of the type

$$W_2(\mu, \nu) \leq \|T_\mu - T_\nu\|_{L^2(\rho, \mathbb{R}^d)} \lesssim W_2(\mu, \nu)^{\frac{p}{6p+16d}}. \quad (20)$$

These stability estimates show that the *linearized optimal transport* (Wang et al., 2013) metric $W_{2,\rho}(\mu, \nu) = \|T_\mu - T_\nu\|_{L^2(\rho, \mathbb{R}^d)}$ (with respect to which geodesic curves correspond to *generalized geodesics* in (Ambrosio et al., 2008)) is bi-Hölder equivalent to the 2-Wasserstein distance on large subsets of $\mathcal{P}_2(\mathbb{R}^d)$, justifying its use in applications (see Chapter 8 for examples of applications). This result also partially answers the metric geometry question of the embeddability of the Wasserstein space $(\mathcal{P}_2(\mathbb{R}^d), W_2)$ into a Hilbert space. While a strong negative result found in (Andoni et al., 2018) ensures that the whole Wasserstein space cannot be embedded into any Hilbert space in a bi-Hölder way, our bounds ensure that at least large subsets of $(\mathcal{P}_2(\mathbb{R}^d), W_2)$ can be explicitly

embedded into a L^2 space with a controlled bi-Hölder distortion. A final possible interpretation of estimate (20) is in terms of the Riemannian interpretation of the infinite-dimensional space $(\mathcal{P}_2(\mathbb{R}^d), W_2)$: this estimate can be seen as a quantitative continuity estimate for the *inverse exponential map* $\mu \in (\mathcal{P}_2(\mathbb{R}^d), W_2) \mapsto T_\mu - \text{id} \in L^2(\rho, \mathbb{R}^d)$.

In order to prove estimates of type (20), we leverage the fact the optimal transport map T_μ between a fixed source ρ and a target μ reads $T_\mu = (\nabla\psi_\mu^*)_{\#}\rho$ where ψ_μ is a Kantorovich potential for the transport problem between ρ and μ , i.e. a minimizer of (15). In Section 5.2, we first derive stability estimates for $\mu \mapsto \psi_\mu$ and $\mu \mapsto \psi_\mu^*$ as direct consequences of the strong convexity estimates derived in Part I. The stability of $\mu \mapsto T_\mu$ is then obtained in Section 5.3, relying in particular on a new Gagliardo-Nirenberg type inequality for the difference of convex functions proven in Section 5.4 which might be of independent interest.

Chapter 6: Quantitative stability of Wasserstein barycenters with respect to the marginals

In this chapter we derive quantitative stability estimates for Wasserstein barycenters with respect to their marginals. Wasserstein barycenters are Fréchet means in Wasserstein spaces: for Ω a compact subset of \mathbb{R}^d and \mathbb{P} a probability measure on the set of probability measures over Ω , i.e. $\mathbb{P} \in \mathcal{P}(\mathcal{P}(\Omega))$, a Wasserstein barycenter of \mathbb{P} is defined as a minimizer $\mu_{\mathbb{P}}$ of

$$\min_{\mu \in \mathcal{P}(\Omega)} \frac{1}{2} \int_{\mathcal{P}(\Omega)} W_2^2(\rho, \mu) d\mathbb{P}(\rho). \quad (21)$$

Such barycenters give geometrically meaningful notions of *averages* of probability measures. Their use is increasingly popular in applied fields, such as image, geometry or language processing. In these fields however, the probability measure of interest $\mathbb{P} \in \mathcal{P}(\mathcal{P}(\Omega))$ is often not accessible in its entirety and the practitioner may have to deal with a statistical or computational approximation $\mathbb{Q} \in \mathcal{P}(\mathcal{P}(\Omega))$ instead. In this chapter, we quantify the effect of such approximations on the corresponding barycenters. We show that Wasserstein barycenters depend in a Hölder-continuous way on their marginals under relatively mild regularity assumptions. In rough terms, our result is the following. Let $\mathbb{P} \in \mathcal{P}(\mathcal{P}(\Omega))$ be such that \mathbb{P} gives mass to a set of measures whose associated Kantorovich functional (studied in Part I) satisfy strong convexity estimates of type (18). Then for any $\mathbb{Q} \in \mathcal{P}(\mathcal{P}(\Omega))$, the (unique) barycenter $\mu_{\mathbb{P}}$ of \mathbb{P} and any barycenter $\mu_{\mathbb{Q}}$ of \mathbb{Q} satisfy

$$W_2(\mu_{\mathbb{P}}, \mu_{\mathbb{Q}}) \lesssim \mathcal{W}_1(\mathbb{P}, \mathbb{Q})^{1/6}, \quad (22)$$

where \mathcal{W}_1 denotes the 1-Wasserstein metric on the metric space $(\mathcal{P}(\Omega), W_2)$.

Before we prove this stability estimate, we study in Section 6.1 under what conditions we can hope for any stability result. Whenever the ambient dimension d is greater than 2, we show that regularity assumptions (such as absolute continuity and connectedness of the support) must be made on some of the marginal measures of either \mathbb{P} or \mathbb{Q} in (22). We then present the dual formulation of (21) as well as our main assumptions and estimate, and we give some immediate but useful consequences of this result in applications. We next show that the proof of estimates of type (22) can be deduced from two stability estimates: a first stability estimate for the dual solutions to the Wasserstein barycenter problem with respect to the marginal measures, derived in Section 6.2, and a second stability estimate for the push-forward operation under a (not necessarily smooth)

optimal transport map, derived in Section 6.3. Section 6.4 finally gives the proof of the dual formulation of (21).

Chapter 7: Quantitative stability of Schrödinger potentials with respect to the temperature in the semi-discrete setting

It is now well-known that for a large enough regularization parameter, the entropy-regularized optimal transport problem is easier to solve than its non-regularized counterpart and the associated quantities have a better statistical behavior. The practitioner may thus be interested in approximating optimal transport quantities using their entropy-regularized versions. These practices call for the study of the approximation error that they induce. In the case of discrete optimal transport (where both the source and target are discrete), the literature ensures fast and non-asymptotic convergence rates of entropic quantities toward their classical analogues as the regularization parameter goes to zero. The case of semi-discrete optimal transport is however less ahead: in this setting, nothing was known quantitatively of the effect of entropic regularization until the recent work (Altschuler et al., 2022) where asymptotic bounds were derived. In this chapter, we improve these bounds to non-asymptotic and nearly tight ones. We give ourselves a fixed probability source density ρ supported over a compact and convex set \mathcal{X} and a fixed discrete target measure $\mu = \sum_{i=1}^N \mu_i \delta_{y_i}$ supported over a finite set $\mathcal{Y} = \{y_i\}_{1 \leq i \leq N}$. For these source and target, we consider for any $\varepsilon \geq 0$ the entropy-regularized optimal transport problem (19) already studied in Chapter 4 and denote γ^ε its (unique) solution as well as $\psi^\varepsilon \in \mathbb{R}^N$ the solution to the dual of this problem verifying $\sum_{i=1}^N \psi_i^\varepsilon = 0$ (such a dual solution is often called a Schrödinger or Sinkhorn potential). We show that whenever the source density ρ is α -Hölder continuous for some $\alpha > 0$, the mapping $\varepsilon \mapsto \psi^\varepsilon$ is *better than Lipschitz*: for any $0 < \varepsilon' \leq \varepsilon \leq 1$ and $\alpha' \in (0, 1)$, we ensure

$$\|\psi^\varepsilon - \psi^{\varepsilon'}\|_\infty \lesssim \varepsilon^{\alpha'} (\varepsilon - \varepsilon').$$

Such fact may be a first step towards a mathematical justification of ε -scaling heuristics used for the numerical resolution of regularized semi-discrete optimal transport, where ε is gradually decreased over the course of the iterations of an algorithm that aims at solving the dual of (19). Letting ε' go to zero, this bound ensures a super-linear rate of convergence in ε of ψ^ε to the non-regularized solution ψ^0 . This result also ensures that there exists a function $c : \mathcal{X} \rightarrow \mathbb{R}_+$ that is positive ρ -almost everywhere and that verifies for all $x \in \mathcal{X}$ and $y \in \mathcal{Y}$

$$|\gamma^\varepsilon(x, y) - \gamma^0(x, y)| \lesssim e^{-c(x)/\varepsilon}.$$

This result may be seen, in the semi-discrete setting, as a non-asymptotic version of the large deviations result shown recently in (Bernton et al., 2022). Finally, the presented bounds also entail a non-asymptotic and tight expansion of the difference between the entropic and the unregularized costs. Namely, denoting

$$W_{2,\varepsilon}(\rho, \mu) = \left(\int_{\mathcal{X} \times \mathcal{Y}} \|x - y\|^2 d\gamma^\varepsilon(x, y) \right)^{1/2}$$

the *approximated* Wasserstein distance between ρ and μ , we prove that there is an explicit constant $C(\rho, \mu)$ that depends only on ρ and μ such that for any $\alpha' \in (0, 1)$,

$$|W_{2,\varepsilon}^2(\rho, \mu) - W_2^2(\rho, \mu) - \varepsilon^2 C(\rho, \mu)| \lesssim \varepsilon^{2+\alpha'}.$$

In order to prove these estimates, we first recall in Section 7.2 elements of semi-discrete (entropic) optimal transport from Chapters 2 and 4 and state our main results. Section 7.3 derives the ODE from which starts the proof of our main bound. This ODE presents two terms that both involve the entropic semi-discrete Kantorovich functional introduced in Chapter 4. The strong-convexity estimate of this chapter together with another estimate derived in Section 7.4 then allow to prove our main bound. The two corollaries to this main bound giving rates of convergences of ψ^ε to ψ^0 and of $W_{2,\varepsilon}$ to W_2 are respectively proven in Section 7.5 and Section 7.6. Section 7.7 finally illustrates our theoretical results on simple one-dimensional numerical examples.

Part III: Numerical applications: the Linearized Optimal Transport framework

Chapter 8: Linearized optimal transport and applications

This final chapter gathers numerical illustrations and experiments revolving around the *linearized optimal transport* (LOT) framework of (Wang et al., 2013), an approximated optimal transport framework which is analyzed to some extent in Chapter 5. In Section 8.2, we illustrate the theoretical results of Chapter 8.2 and observe the metric distortion induced by the LOT embedding on some two-dimensional examples. We also mention how the LOT embedding may be used to perform barycenter approximation in the 2-Wasserstein space. Then in Section 8.3, we give two example extensions of classical Hilbertian data analysis methods to probability measures within the LOT framework. These extensions concern K -means and dictionary learning problems in the 2-Wasserstein space.

Appendix

Chapter A: Optimal transport facts

This short appendix collects some optimal transport facts that are useful in this thesis but not treated in Chapter 1.

Publications

- *Quantitative stability of optimal transport maps and linearization of the 2-Wasserstein space.* Q. Mérigot, A. Delalande, F. Chazal. Proceedings of the 23rd International Conference on Artificial Intelligence and Statistics (AISTATS), PMLR 108:3186-3196, 2020.
- *Quantitative stability of optimal transport maps under variations of the target measure.* A. Delalande, Q. Mérigot. Under revision at the Duke Mathematical Journal.
- *Nearly tight convergence bounds for semi-discrete entropic optimal transport.* A. Delalande. Proceedings of the 25th International Conference on Artificial Intelligence and Statistics (AISTATS), PMLR 151:1619-1642, 2022.
- *Quantitative stability of barycenters in the Wasserstein space.* G. Carlier, A. Delalande, Q. Mérigot. Preprint.

Resources for the numerical experiments

- https://github.com/alex-delalande/stability_ot_maps_and_linearization_wasserstein_space
- <https://github.com/alex-delalande/potentials-entropic-sd-ot>
- https://github.com/alex-delalande/linearized_wasserstein_dictionary_learning

Notation

Spaces.

$\bar{\mathbb{R}}$: extended real number line, also denoted $\mathbb{R} \cup \{-\infty\} \cup \{+\infty\}$.
$\mathcal{M}(\Omega)$: set of Radon bounded measures supported over $\Omega \subseteq \mathbb{R}^d$.
$\mathcal{M}_+(\Omega)$: set of positive Radon bounded measures supported over $\Omega \subseteq \mathbb{R}^d$.
$\mathcal{M}_p(\Omega)$: set of Radon bounded measures supported over $\Omega \subseteq \mathbb{R}^d$ that admit a finite p -th moment.
$\mathcal{P}(\Omega)$: set of Borel probability measures supported over $\Omega \subseteq \mathbb{R}^d$.
$\mathcal{P}_{a.c.}(\Omega)$: set of Borel probability measures supported over $\Omega \subseteq \mathbb{R}^d$ that are absolutely continuous w.r.t. the Lebesgue measure.
$\mathcal{P}_p(\Omega)$: set of Borel probability measures supported over $\Omega \subseteq \mathbb{R}^d$ that admit a finite p -th moment.
$\mathcal{P}_{p,a.c.}(\Omega)$: set of Borel probability measures supported over $\Omega \subseteq \mathbb{R}^d$ that admit a finite p -th moment and are absolutely continuous w.r.t. the Lebesgue measure.
$\mathcal{C}_b(\Omega)$: set of continuous and bounded functions from $\Omega \subseteq \mathbb{R}^d$ to \mathbb{R} .
$(1 + \ \cdot\ ^2)\mathcal{C}_b(\Omega)$: set of continuous functions from $\Omega \subseteq \mathbb{R}^d$ to \mathbb{R} with at most quadratic growth.
$\mathcal{C}^k(\Omega)$: set of functions from $\Omega \subseteq \mathbb{R}^d$ to \mathbb{R} admitting $k \geq 0$ continuous derivatives.
$\mathcal{C}^{k,\alpha}(\Omega)$: subset of functions of $\mathcal{C}^k(\Omega)$ with α -Hölder continuous k -th derivative.
$L^p(\rho)$: set of L^p -integrable ρ -measurable functions from \mathbb{R}^d to $\mathbb{R} \cup \{+\infty\}$.
$L^p(\rho, \mathbb{R}^d)$: set of L^p -integrable ρ -measurable functions from \mathbb{R}^d to \mathbb{R}^d .
$\text{Lip}_k(\Omega)$: set of k -Lipschitz functions from $\Omega \subseteq \mathbb{R}^d$ to \mathbb{R} .

Measures.

$M_p(\mu)$: p -th moment of a probability measure $\mu \in \mathcal{P}(\mathbb{R}^d)$: $M_p(\mu) = \int_{\Omega} \ x\ ^p d\mu(x)$.
$\mathbb{E}_{\mu}(f)$: expectation of f against μ (equal to $\int_{\mathbb{R}^d} f d\mu$, possibly denoted $\langle f \mu \rangle$).
$\text{Var}_{\mu}(f)$: variance of f against μ (equal to $\int_{\mathbb{R}^d} f^2 d\mu - (\int_{\mathbb{R}^d} f d\mu)^2$).
$\text{spt}(\rho)$: topological support of a probability measure $\rho \in \mathcal{P}(\mathbb{R}^d)$.

Others.

$\text{diam}(\Omega)$: diameter of $\Omega \subset \mathbb{R}^d$.
χ_{Ω}	: for a subset $\Omega \subset \mathbb{R}^d$, denotes the indicator function of S , i.e. $\chi_{\Omega}(x) = 1$ if $x \in \Omega$, and 0 else.
ι_C	: for a convex set $C \subset \mathbb{R}^d$, denotes the convex characteristic function of C , valued 0 on C and $+\infty$ everywhere else.
$\text{diag}(v)$: for $v \in \mathbb{R}^N$, denotes the diagonal matrix of $\mathbb{R}^{N \times N}$ with v on the diagonal.
$B(x, R)$: Euclidean ball of \mathbb{R}^d centered at $x \in \mathbb{R}^d$ and of radius $R > 0$.
$\phi \oplus \psi$: for two functions $\phi, \psi : \Omega \subseteq \mathbb{R}^d \rightarrow \mathbb{R}$, denotes the function from $\Omega \times \Omega$ to \mathbb{R} defined by $\phi \oplus \psi : (x, y) \mapsto \phi(x) + \psi(y)$.
$\langle \cdot \cdot \rangle$: canonical pairing between dual spaces (in \mathbb{R}^d , corresponds to the inner product).
S^{d-1}	: $(d-1)$ -sphere in \mathbb{R}^d .

Part I

Strong convexity of the quadratic optimal transport problem

Quadratic optimal transport and strong convexity of the dual

Abstract

This chapter stands as an introduction to the first part of this thesis. In a first section (§1.1), we recall some well-known facts about quadratic optimal transport and introduce what we call the *Kantorovich functional*, that appears in the dual formulation of the quadratic optimal transport problem. It is shown that the study of this functional can be useful to derive qualitative and quantitative properties of the solutions to quadratic optimal transport problems. These ideas motivate in particular the study of the strong convexity of this functional in a second section (§1.2), where a strong convexity estimate is announced and the outline of the rest of the first part of this thesis is given. The new strong convexity estimates derived in this thesis all follow from well-known geometric and functional inequalities that are recalled in a third section (§1.3). Finally, we give in a fourth section (§1.4) an extension of the main result of the first part of this thesis.

1.1 Quadratic optimal transport

This section gives a brief introduction to the *quadratic* optimal transport problem, that corresponds arguably to the most studied of the optimal transport problems and is the focus of this thesis. This section does not present new results but introduces the main objects and concepts at work in this dissertation. Complementary facts about optimal transport (with possibly other costs than the quadratic one) are gathered in Chapter A of the appendix. We also refer to the following monographs and chapters (from which this section and Chapter A are inspired) for more general presentations of theoretical and computational aspects related to the field: (Villani, 2003, 2008), (Santambrogio, 2015), (Peyré and Cuturi, 2019) and (Mérigot and Thibert, 2021).

1.1.1 Monge and Kantorovich formulations

Monge formulation. The optimal transport problem, as introduced by Monge in 1781 (Monge, 1781), may be stated as follows. Let $\Omega \subseteq \mathbb{R}^d$. Given two probability measures

$\rho, \mu \in \mathcal{P}(\Omega)$ that represent two different distributions of mass and given a cost function $c : \Omega \times \Omega \rightarrow \mathbb{R}_+$ that encodes with $c(x, y)$ the *cost* of transporting a unit of mass from a location $x \in \Omega$ to a location $y \in \Omega$, look for a transport map $T : \Omega \rightarrow \Omega$ that solves the following non-convex optimization problem:

$$\inf_{T_{\#}\rho = \mu} \int_{\Omega} c(x, T(x)) d\rho(x), \quad (\text{MP})$$

where $T_{\#}\rho$ corresponds to the push-forward measure of ρ by T , which satisfies $T_{\#}\rho(A) = \rho(T^{-1}(A))$ for every ρ -measurable set A . The quadratic setting, that we consider from now on, corresponds to choosing the squared distance as a ground cost:

$$c : (x, y) \mapsto \|x - y\|^2,$$

where $\|\cdot\|$ denotes the Euclidean norm in \mathbb{R}^d . When a minimizing map T exists for problem (MP) with the quadratic cost as ground cost, such a map may be called either an *optimal transport map* or a *Monge map* between ρ and μ . The weakness of formulation (MP) is that such a map may not always exist. Indeed, one can first notice that there does not always exist a transport map between any measures ρ and μ : for instance, think of the case where ρ is a Dirac mass, which implies that $T_{\#}\rho$ is also a Dirac mass and prevents from satisfying the constraint $T_{\#}\rho = \mu$ unless μ is a Dirac mass itself. Moreover, even when there exist transport maps between ρ and μ , the existence of a minimizing one remains a challenging question in general: the push-forward constraint is highly non-linear in T , which prevents from easily employ direct methods in the calculus of variations to show existence of solutions (see Example 4.9 of (Villani, 2008) for a concrete example where transport maps exist but no transport map is optimal).

Kantorovich formulation. In 1942, Kantorovich introduced a linear optimization problem that latter on revealed itself to be a *relaxed* version of Monge's problem (MP), and that allowed to overcome its above-mentioned limitations. In Kantorovich's version, instead of looking for a transport map $T : \Omega \rightarrow \Omega$ that sends all the mass located at $x \in \Omega$ to a unique target location $T(x) \in \Omega$, one look for a *transport plan* $\gamma \in \mathcal{P}(\Omega \times \Omega)$ that sends a fraction $d\gamma(x, y)$ of the mass $d\rho(x)$ from a location $x \in \Omega$ to a location $y \in \Omega$, thus allowing to split the mass coming from the source measure during the transport process. Less formally, introduce $\Gamma(\rho, \mu)$ the set of transport plans between ρ and μ , i.e. the set of probability measures over $\Omega \times \Omega$ with first marginal ρ and second marginal μ :

$$\Gamma(\rho, \mu) = \{\gamma \in \mathcal{P}(\Omega \times \Omega) \mid \forall \text{ Borel set } B \subset \Omega, \gamma(B \times \Omega) = \rho(B), \gamma(\Omega \times B) = \mu(B)\}.$$

Kantorovich's formulation of the quadratic optimal transport problem then corresponds to the following convex optimization problem:

$$\inf_{\gamma \in \Gamma(\rho, \mu)} \int_{\Omega \times \Omega} \|x - y\|^2 d\gamma(x, y). \quad (\text{KP})$$

Problem (KP) is a relaxation of (MP) in the sense that if there exists a transport map T between ρ and μ , then one can build an admissible transport plan γ_T from it: $\gamma_T := (\text{id}, T)_{\#}\rho \in \Gamma(\rho, \mu)$. Such transport plan verifies in particular

$$\int_{\Omega \times \Omega} \|x - y\|^2 d\gamma_T(x, y) = \int_{\mathbb{R}^d} \|T(x) - x\|^2 d\rho(x),$$

so that one always has $(\text{KP}) \leq (\text{MP})$. In some cases, one can ensure that $(\text{KP}) = (\text{MP})$. We will study in Section 1.1.3 (in particular with Theorem 1.12) the most famous of these cases, that corresponds to the instance where the source measure ρ is absolutely continuous. When a minimizing transport plan γ exists in (KP), such a plan is called an *optimal transport plan*. A clear advantage of (KP) over (MP) is that whenever the probability measures ρ, μ admit a finite second-order moment (which is automatically the case when Ω is compact), an optimal transport plan always exists:

Theorem 1.1. *Let $\Omega \subseteq \mathbb{R}^d$ and $\rho, \mu \in \mathcal{P}_2(\Omega)$. Then (KP) admits a solution.*

This result can be shown using the direct method in the calculus of variations, i.e. by looking at a minimizing sequence, extracting a converging subsequence by compactness of the set of transport plans (deduced from Prokhorov's theorem) and showing that the limit of the subsequence is a minimizer thanks to the continuity of the quadratic cost, see e.g. Proposition 2.1 of (Villani, 2003). Another advantage of formulation (KP) is that it is a convex optimization problem: it admits as such a dual formulation, that we present in the next subsection.

1.1.2 Dual formulation

For $\rho, \mu \in \mathcal{P}_2(\Omega)$ having a finite second-order moment, we have observed that (KP) admits a minimizer. Developing the square in the integral term of (KP), one can notice that solving this problem with $\rho, \mu \in \mathcal{P}_2(\Omega)$ is equivalent to solving the maximum correlation problem

$$\max_{\gamma \in \Gamma(\rho, \mu)} \int_{\Omega \times \Omega} \langle x|y \rangle d\gamma(x, y), \quad (\text{KP}')$$

with the relation $(\text{KP}) = M_2(\rho) + M_2(\mu) - 2 \times (\text{KP}')$. We can thus now focus exclusively on problem (KP'). This problem presents the constraint that a candidate minimizer γ must be a transport plan between ρ and μ and thus belong to $\Gamma(\rho, \mu)$. Such constraint may be expressed using the method of Lagrange multipliers: for $\gamma \in \mathcal{M}_+(\Omega \times \Omega)$,

$$\inf_{\phi, \psi \in \mathcal{C}^0(\Omega)} \int_{\Omega} \phi d\rho + \int_{\Omega} \psi d\mu - \int_{\Omega \times \Omega} \phi \oplus \psi d\gamma = \begin{cases} 0 & \text{if } \gamma \in \Gamma(\rho, \mu), \\ -\infty & \text{else,} \end{cases}$$

where $\phi \oplus \psi : (x, y) \mapsto \phi(x) + \psi(y)$. With this representation of the constraint, (KP') can be rewritten

$$(\text{KP}') = \sup_{\gamma \in \mathcal{M}_+(\Omega \times \Omega)} \int_{\Omega \times \Omega} \langle \cdot | \cdot \rangle d\gamma + \inf_{\phi, \psi \in \mathcal{C}^0(\Omega)} \int_{\Omega} \phi d\rho + \int_{\Omega} \psi d\mu - \int_{\Omega \times \Omega} \phi \oplus \psi d\gamma.$$

The Lagrangian dual problem of (KP') is then obtained by exchanging the supremum and infimum in this last formulation:

$$\inf_{\phi, \psi \in \mathcal{C}^0(\Omega)} \int_{\Omega} \phi d\rho + \int_{\Omega} \psi d\mu + \sup_{\gamma \in \mathcal{M}_+(\Omega \times \Omega)} \int_{\Omega \times \Omega} (\langle \cdot | \cdot \rangle - \phi \oplus \psi) d\gamma. \quad (1.1)$$

In this formulation, the supremum also represents a constraint:

$$\sup_{\gamma \in \mathcal{M}_+(\Omega \times \Omega)} \int_{\Omega \times \Omega} (\langle \cdot | \cdot \rangle - \phi \oplus \psi) d\gamma = \begin{cases} 0 & \text{if } \forall x, y \in \Omega, \phi(x) + \psi(y) \geq \langle x|y \rangle, \\ +\infty & \text{else.} \end{cases}$$

The unconstrained dual problem (1.1) thus admits the following constrained formulation:

$$\inf_{\phi, \psi \in \mathcal{C}^0(\Omega)} \left\{ \int_{\Omega} \phi d\rho + \int_{\Omega} \psi d\mu \mid \forall x, y \in \Omega, \quad \phi(x) + \psi(y) \geq \langle x|y \rangle \right\}. \quad (\text{DP})$$

From the constraints in the last dual problem, it is easy to check that (DP) \geq (KP') so that naturally weak duality holds. The following result ensures that strong-duality actually holds, i.e. (KP') = (DP).

Theorem 1.2. *Let $\Omega \subseteq \mathbb{R}^d$ and $\rho, \mu \in \mathcal{P}_2(\Omega)$. Then*

$$\max_{\gamma \in \Gamma(\rho, \mu)} \int_{\Omega \times \Omega} \langle \cdot | \cdot \rangle d\gamma = \inf_{\phi, \psi \in \mathcal{C}^0(\Omega), \phi \oplus \psi \geq \langle \cdot | \cdot \rangle} \int_{\Omega} \phi d\rho + \int_{\Omega} \psi d\mu.$$

This strong duality result, often referred to as *Kantorovich duality*, is well-known and holds in much more general settings than our Euclidean setting with the squared distance as a ground cost (see e.g. Theorem 5.9 of (Villani, 2008), reported partially in Theorem A.5 of the appendix). We detail however a possible proof of this result, for completeness and because we will use a similar approach to show another strong duality result in Chapter 6. This approach is inspired from Section 2 of (Agueh and Carlier, 2011) and from Section 1.6.3 of (Santambrogio, 2015), itself suggested by C. Jimenez and adapted from Section 4 of (Bouchitté and Buttazzo, 2001). We will require the notion of convex conjugation, or Legendre transformation:

Definition 1.3 (Convex conjugate). Let V be a real topological vector space and let V^* be its dual space, with canonical pairing denoted $\langle \cdot | \cdot \rangle : V \times V^* \rightarrow \mathbb{R}$. The convex conjugate, or Legendre transform, of a function $f : V \rightarrow \overline{\mathbb{R}}$ is the function defined by

$$f^* : \begin{cases} V^* \rightarrow \overline{\mathbb{R}}, \\ v^* \mapsto \sup_{v \in V} \langle v | v^* \rangle - f(v). \end{cases}$$

As a supremum of linear functions, the convex conjugate is always convex and lower semi-continuous. Moreover, from its definition, the biconjugate f^{**} always satisfies $f^{**} \leq f$. In the proof of Theorem 1.2, we will use the equality case of this inequality given by the Fenchel–Moreau theorem, that ensures that a proper function (a function is said proper if it does not take the value $-\infty$ and it is not equal to $+\infty$ everywhere) $f : V \rightarrow \overline{\mathbb{R}}$ satisfies $f^{**} = f$ if and only if f is convex and lower semi-continuous.

Proof of Theorem 1.2. We will prove a slightly stronger result by restricting the infimum in (DP) and taking ϕ, ψ in the set $(1 + \|\cdot\|^2)\mathcal{C}_b(\Omega)$ of continuous functions over Ω with at most quadratic growth. Define the functional

$$H : \begin{cases} (1 + \|\cdot\|^2)\mathcal{C}_b(\Omega \times \Omega) \rightarrow \mathbb{R}, \\ p \mapsto \inf_{\phi, \psi \in (1 + \|\cdot\|^2)\mathcal{C}_b(\Omega)} \{ \langle \phi | \rho \rangle + \langle \psi | \mu \rangle \mid \phi \oplus \psi - \langle \cdot | \cdot \rangle \geq p \}, \end{cases}$$

and notice that by definition H is convex. Also notice that $H(0) \geq$ (DP). Let's now compute the Legendre transform of H . For any $\gamma \in \mathcal{M}_2(\Omega \times \Omega)$, one has

$$\begin{aligned} H^*(\gamma) &= \sup_p \langle p | \gamma \rangle - H(p) \\ &= \sup_{p, \phi, \psi} \{ \langle p | \gamma \rangle - \langle \phi | \rho \rangle - \langle \psi | \mu \rangle \mid \phi \oplus \psi - \langle \cdot | \cdot \rangle \geq p \}. \end{aligned}$$

If $\gamma \notin \mathcal{M}_+(\Omega \times \Omega)$, then there exists $p_0 \in (1 + \|\cdot\|^2)\mathcal{C}_b(\Omega \times \Omega)$ verifying $p_0 \leq 0$ and such that $\langle p_0 | \gamma \rangle > 0$. Setting $\phi = \psi = 0$ in the last supremum and $p = p_n = -\langle \cdot | \cdot \rangle + np_0$ for $n \geq 0$, this ensures that

$$H^*(\gamma) \geq - \int_{\Omega \times \Omega} \langle x | y \rangle d\gamma(x, y) + n \langle p_0 | \gamma \rangle \geq -M_2(\gamma) + n \langle p_0 | \gamma \rangle.$$

Letting $n \rightarrow +\infty$ in the last inequality shows that $H^*(\gamma) = +\infty$ if $\gamma \notin \mathcal{M}_+(\Omega \times \Omega)$. Assume now that $\gamma \in \mathcal{M}_+(\Omega \times \Omega)$. Then in the computation of $H^*(\gamma)$, p may be chosen as large as possible, i.e.

$$p = \phi \oplus \psi - \langle \cdot | \cdot \rangle.$$

This ensures that for $\gamma \in \mathcal{M}_+(\Omega \times \Omega)$,

$$\begin{aligned} H^*(\gamma) &= - \int_{\Omega \times \Omega} \langle x | y \rangle d\gamma(x, y) + \sup_{\phi, \psi} \langle \phi \oplus \psi | \gamma \rangle - \langle \phi | \rho \rangle - \langle \psi | \mu \rangle \\ &= - \int_{\Omega \times \Omega} \langle x | y \rangle d\gamma(x, y) + \begin{cases} 0 & \text{if } \gamma \in \Gamma(\rho, \mu), \\ +\infty & \text{else.} \end{cases} \end{aligned}$$

Thus for $\gamma \in \Gamma(\rho, \mu)$, $H^*(\gamma)$ is the negative correlation induced by the coupling γ between ρ and μ and we have

$$(\text{KP}') = \sup_{\gamma \in \mathcal{M}_2(\Omega \times \Omega)} -H^*(\gamma) = H^{**}(0).$$

We thus have shown so far

$$H(0) \geq (\text{DP}) \geq (\text{KP}') = H^{**}(0),$$

so that equality between these terms will hold if we can show $H^{**}(0) = H(0)$. Since H is convex, this will follow from the continuity of H at 0 for the following supremum norm for $p \in (1 + \|\cdot\|^2)\mathcal{C}_b(\Omega \times \Omega)$:

$$\|p\| := \sup_{(x, y) \in \Omega \times \Omega} \frac{|p(x, y)|}{1 + \|x\|^2 + \|y\|^2} < +\infty.$$

We can first notice that H never takes the value $-\infty$: for $p \in (1 + \|\cdot\|^2)\mathcal{C}_b(\Omega \times \Omega)$ and any $\phi, \psi \in (1 + \|\cdot\|^2)\mathcal{C}_b(\Omega)$ such that $\phi \oplus \psi - \langle \cdot | \cdot \rangle \geq p$, one has

$$\begin{aligned} \langle \phi | \rho \rangle + \langle \psi | \mu \rangle &\geq \int_{\Omega \times \Omega} (p(x, y) + \langle x | y \rangle) d\rho \otimes \mu(x, y) \\ &\geq \int_{\Omega \times \Omega} \left(-(1 + \|x\|^2 + \|y\|^2) \|p\| - \frac{1}{2}(\|x\|^2 + \|y\|^2) \right) d\rho \otimes \mu(x, y) \\ &\geq -(1 + M_2(\rho) + M_2(\mu)) \|p\| - \frac{1}{2}(M_2(\rho) + M_2(\mu)) > -\infty. \end{aligned}$$

Thus $H(p) > -\infty$. On the other hand, one can notice that H is bounded above in a neighborhood of 0 in $(1 + \|\cdot\|^2)\mathcal{C}_b(\Omega \times \Omega)$: let $p \in (1 + \|\cdot\|^2)\mathcal{C}_b(\Omega \times \Omega)$ be such that $\|p\| \leq 1$. Then for all $x, y \in \Omega$, $p(x, y) \leq 1 + \|x\|^2 + \|y\|^2$. Thus choosing $\phi : x \mapsto 1 + \frac{3}{2}\|x\|^2$, $\psi : y \mapsto \frac{3}{2}\|y\|^2$ yields $\phi \oplus \psi - \langle \cdot | \cdot \rangle \geq p$, so that

$$H(p) \leq \langle \phi | \rho \rangle + \langle \psi | \mu \rangle \leq 1 + \frac{3}{2}M_2(\rho) + \frac{3}{2}M_2(\mu) < +\infty.$$

A standard convex analysis result (Proposition 2.5 of (Ekeland and Témam, 1999)) then ensures that H is continuous at 0, which ensures $H(0) = H^{**}(0)$. \square

Semi-dual formulation. One can notice that (DP) may be turned easily into an unconstrained minimization problem. Indeed, for a given $\psi \in \mathcal{C}^0(\Omega)$ (resp. $\phi \in \mathcal{C}^0(\Omega)$), it is interesting to choose ϕ (resp. ψ) as small as possible while satisfying the constraint, that is:

$$\forall x \in \Omega, \quad \phi(x) = \sup_{y \in \Omega} \langle x|y \rangle - \psi(y) \quad (\text{resp. } \forall y \in \Omega, \quad \psi(y) = \sup_{x \in \Omega} \langle x|y \rangle - \phi(x)).$$

This corresponds to choosing $\phi = \psi^*$ (resp. $\psi = \phi^*$) (see Definition 1.3). Problem (DP) thus admits the following unconstrained formulations, sometimes referred to as the *semi-dual* formulations:

$$(\text{DP}) = \inf_{\psi \in \mathcal{C}^0(\Omega)} \langle \psi^*|\rho \rangle + \langle \psi|\mu \rangle = \inf_{\phi \in \mathcal{C}^0(\Omega)} \langle \phi|\rho \rangle + \langle \phi^*|\mu \rangle.$$

From the last remarks, a natural procedure to improve the value given by a candidate $\psi_0 \in \mathcal{C}^0(\Omega)$ in (DP) could be the following: first set $\psi \leftarrow \psi_0$ and $\phi \leftarrow \psi^*$. Then for all $k \geq 1$, alternate between $\psi \leftarrow \phi^*$ and $\phi \leftarrow \psi^*$. However such procedure *converges* at iteration $k = 1$. Indeed, at this iteration, one sets $\psi \leftarrow \psi_0^{**}$ and $\phi \leftarrow \psi_0^{**}$. But by the Fenchel-Moreau theorem, $\psi_0^{**} = \psi_0^*$, so that we are back to right before iteration $k = 1$. This fact thus prevents from easily solving (DP) from successive convex conjugations, but it at least ensures the following: in the semi-dual formulation, one can impose that ψ (resp. ϕ) is a proper convex function, so that $\psi^{**} = \psi$ (resp. $\phi^{**} = \phi$). This idea allows to prove the existence of solutions to the dual problem:

Theorem 1.4. *Let $\Omega \subseteq \mathbb{R}^d$ and $\rho, \mu \in \mathcal{P}_2(\Omega)$. Then there exists a pair (ψ^*, ψ) of proper lower semi-continuous (l.s.c.) conjugate convex functions on \mathbb{R}^d such that*

$$(\text{DP}) = \langle \psi^*|\rho \rangle + \langle \psi|\mu \rangle.$$

Before we mention how to prove this result, let us recall the definition of the subdifferential of a function and state some convex analysis facts that allow to describe qualitatively the minimizers of (DP).

Definition 1.5 (Subdifferential). Let V be a real topological vector space and denote V^* its dual space. The subdifferential of a function $f : V \rightarrow \mathbb{R} \cup \{+\infty\}$ at $x \in V$ such that $f(x) < +\infty$ is defined by

$$\partial f(x) = \{g \in V^* | \forall y \in V, f(x) + \langle g|y - x \rangle \leq f(y)\}.$$

Remark 1.6 (Convex analysis facts). A proper convex function $\phi : \Omega \rightarrow \bar{\mathbb{R}}$ is continuous and locally Lipschitz on the interior of $\{x \in \Omega | \phi(x) < +\infty\}$. Hence by Rademacher's theorem it is differentiable almost everywhere on this set. By definition of the convex conjugate, it verifies the Fenchel-Young inequality:

$$\forall x, y \in \Omega, \quad \phi(x) + \phi^*(y) \geq \langle x|y \rangle.$$

Moreover, the Fenchel-Young equality case can easily be characterized: for all $x, y \in \Omega$,

$$\phi(x) + \phi^*(y) = \langle x|y \rangle \iff y \in \partial\phi(x) \iff x \in \partial\phi^*(y).$$

Finally, we see from this equivalence (or from the definition of the convex conjugation) that the domain taken in the definition of the convex conjugate is important: if $\Omega \subset \mathbb{R}^d$ is compact and is included in a ball $B(0, R_\Omega)$ centered at 0 and of radius $R_\Omega > 0$, the convex conjugate of a function $\phi : \Omega \rightarrow \mathbb{R}$ – which is defined on the whole \mathbb{R}^d – is R_Ω -Lipschitz continuous.

Theorem 1.4 is a consequence of the above remark. Indeed, we know that we can assume that a candidate minimizer ψ in (DP) is a continuous proper convex function and set $\phi = \psi^*$. For $\Omega \subset B(0, R_\Omega)$ compact, Remark 1.6 ensures that any such candidate ψ is R_Ω -Lipschitz continuous, and it can be assumed to be bounded since for any constant $c \in \mathbb{R}$, a pair (ψ^*, ψ) performs as well in (DP) as the pair $((\psi+c)^*, \psi+c) = (\psi^*-c, \psi+c)$. From a minimizing sequence in (DP), the Arzelà-Ascoli theorem then allows to extract a uniformly converging subsequence, that is easily shown to be maximizing. The general (non-compact) case is then obtained by approximation, see e.g. Theorem 2.9 of (Villani, 2003). Later on, the minimizers ψ or ψ^* that appear in Theorem 1.4 may be called *Kantorovich potentials*.

Remark 1.7 (Primal-dual relations). Let $\Omega \subseteq \mathbb{R}^d$ and $\rho, \mu \in \mathcal{P}_2(\Omega)$. For a primal solution $\gamma \in \Gamma(\rho, \mu)$ to (KP') as given in Theorem 1.1 and a dual solution (ψ^*, ψ) to (DP) as given in Theorem 1.4, the strong duality result (KP') = (DP) of Theorem 1.2 ensures the equality

$$\int_{\Omega} \psi^*(x) d\rho(x) + \int_{\Omega} \psi(y) d\mu(y) = \int_{\Omega \times \Omega} \langle x|y \rangle d\gamma(x, y).$$

Thus by the Fenchel-Young inequality and equality case (Remark 1.6), for γ -almost-every $(x, y) \in \Omega \times \Omega$, $y \in \partial\psi^*(x)$ or, equivalently, $x \in \partial\psi(y)$. This ensures that for μ -almost-every $y \in \Omega$, there exists x in the support $\text{spt}(\rho)$ of ρ such that $\psi(y) = \psi^{**}(y) = \langle x|y \rangle - \psi^*(x)$, so that μ -almost-every $y \in \Omega$ satisfies

$$\psi(y) = \sup_{x \in \text{spt}(\rho)} \langle x|y \rangle - \psi^*(x).$$

Similarly, for ρ -almost every $x \in \Omega$, $\psi^*(x) = \sup_{y \in \text{spt}(\mu)} \langle x|y \rangle - \psi(y)$. These facts show in particular that whenever ρ is supported on a compact set \mathcal{X} included in a ball $B(0, R_{\mathcal{X}})$ of radius $R_{\mathcal{X}} > 0$, ψ can be assumed to be $R_{\mathcal{X}}$ -Lipschitz continuous without any loss of generality and without assuming that μ is compactly supported (to see this, consider replacing ψ^* with $\psi^* + \iota_{B(0, R_{\mathcal{X}})}$ and observe that this does not change the value in (DP)).

We thus have shown that (KP') = (DP) and that both of these problems admit optimizers. In the next subsection, we derive immediate properties of the functional being minimized in (DP) and present the questions we will try to answer in the rest of the first part of this thesis.

1.1.3 The Kantorovich functional

In the preceding paragraphs, for $\Omega \subseteq \mathbb{R}^d$ and $\rho, \mu \in \mathcal{P}_2(\Omega)$, we have established that solving the quadratic optimal transport problem (KP) between ρ and μ is equivalent to solving the following (semi-)dual minimization problem

$$\min_{\psi: \Omega \rightarrow \bar{\mathbb{R}}} \int_{\Omega} \psi^* d\rho + \int_{\Omega} \psi d\mu, \quad (\text{DP}')$$

where ψ can be assumed to be proper, convex and lower semi-continuous. This duality result, often called Kantorovich duality, motivates the definition of the following *Kantorovich functional* defined for a given source measure $\rho \in \mathcal{P}_2(\Omega)$ by

$$\mathcal{K}_{\rho} : \begin{cases} \bar{\mathbb{R}}^{\Omega} & \rightarrow \bar{\mathbb{R}}, \\ \psi & \mapsto \int_{\Omega} \psi^* d\rho. \end{cases}$$

This functional has already been considered many times in the literature. It has been leveraged in particular for the numerical resolution of (semi-discrete) optimal transport problems (Aurenhammer et al., 1998; Mérigot, 2011; De Goes et al., 2012; Genevay et al., 2016; Kitagawa et al., 2019).

Convexity of the Kantorovich functional. It is immediate to notice, from the convexity of the convex conjugation $\psi \mapsto \psi^*$, that the Kantorovich functional \mathcal{K}_ρ is convex for any $\rho \in \mathcal{P}(\mathbb{R}^d)$. We thus recover the fact that the functional being minimized in (DP'), which is the addition of $\psi \mapsto \mathcal{K}_\rho(\psi)$ and the linear term $\psi \mapsto \langle \psi | \mu \rangle$, is convex. Therefore, a bit formally, a minimizer $\psi : \Omega \rightarrow \bar{\mathbb{R}}$ of (DP') may be characterized by the following first order condition:

$$0 \in \partial(\mathcal{K}_\rho(\cdot) + \langle \cdot | \mu \rangle)(\psi) = \partial\mathcal{K}_\rho(\psi) + \mu \iff \psi \in (\partial\mathcal{K}_\rho)^{-1}(-\mu).$$

Such fact, though formal, shows that the study of the Kantorovich functional may help derive qualitative or quantitative information about the minimizers of (DP') (and a posteriori of (KP)). In the following, we observe settings where the subdifferential of the Kantorovich functional can be properly derived.

Subdifferential and differential of the Kantorovich functional. In order to study the (sub)differential properties of the Kantorovich functional, we need to restrict its domain of definition to a relevant topological (or more conveniently normed) space. In general, we can only assume that a minimizer of (DP') is proper, l.s.c. and convex (Theorem 1.4), but such fact does not indicate a natural normed space to work on. Before we worry about this matter, let us state the following lemma that allows to find, for a given convex potential, measures that satisfy the subdifferential inequality in very general cases.

Lemma 1.8. *Let $\Omega \subseteq \mathbb{R}^d$ and let $\rho \in \mathcal{P}_2(\Omega)$. Let $\psi^0 : \Omega \rightarrow \bar{\mathbb{R}}$ be a proper l.s.c. convex function and denote $\phi^0 = (\psi^0)^*$. Assume that the subdifferential of ϕ^0 is not empty ρ -almost-everywhere. Let $g^0 : \Omega \rightarrow \Omega$ be a measurable selection of the subdifferential of ϕ^0 , i.e. for any $x \in \Omega$,*

$$g^0(x) \in \partial\phi^0(x),$$

with the convention $g^0(x) = x_0$ for some $x_0 \in \Omega$ if $\partial\phi^0(x) = \emptyset$. Then, the measure $-g_{\#\rho}^0$ satisfies the subdifferential inequality of \mathcal{K}_ρ at ψ^0 : for any proper $\psi^1 : \Omega \rightarrow \bar{\mathbb{R}}$,

$$\mathcal{K}_\rho(\psi^0) + \langle \psi^1 - \psi^0 | -g_{\#\rho}^0 \rangle \leq \mathcal{K}_\rho(\psi^1).$$

Proof. We have to check that for any proper $\psi^1 : \Omega \rightarrow \bar{\mathbb{R}}$, denoting $\phi^1 = (\psi^1)^*$, we have

$$\mathcal{K}_\rho(\psi^1) - \mathcal{K}_\rho(\psi^0) - \langle \psi^0 - \psi^1 | g_{\#\rho}^0 \rangle = \int_{\Omega} (\phi^1 + \psi^1(g^0) - \phi^0 - \psi^0(g^0)) d\rho \geq 0.$$

But Fenchel-Young inequality and equality (Remark 1.6) give respectively for ρ -almost every $x \in \Omega$

$$\phi^1(x) + \psi^1(g^0(x)) \geq \langle x | g^0(x) \rangle \quad \text{and} \quad \phi^0(x) + \psi^0(g^0(x)) = \langle x | g^0(x) \rangle.$$

This proves the wanted inequality. □

We may now describe properly elements of the subdifferential of \mathcal{K}_ρ in the specific compact setting. Whenever the domain Ω is compact and included in a ball $B(0, R_\Omega)$ centered at 0 and of radius $R_\Omega > 0$, we have observed in Remarks 1.6 and 1.7 that a minimizer ψ of (DP') can be assumed to be R_Ω -Lipschitz continuous. This fact indicates that under such compactness assumptions on Ω and for any $\rho \in \mathcal{P}(\Omega)$, we may restrict the domain of definition of \mathcal{K}_ρ to $\text{Lip}_{R_\Omega}(\Omega)$ without any loss of generality in the resolution of (DP'). We equip $\text{Lip}_{R_\Omega}(\Omega)$ with the norm $\|\cdot\|_\infty + \|\cdot\|_{\text{Lip}}$, where each of these norms is defined for $f \in \text{Lip}_{R_\Omega}(\Omega)$ by

$$\|f\|_\infty = \sup_{x \in \Omega} |f(x)| \quad \text{and} \quad \|f\|_{\text{Lip}} = \sup_{x \neq y \in \Omega} \frac{|f(x) - f(y)|}{\|x - y\|}.$$

Note that the topological dual of $\text{Lip}_{R_\Omega}(\Omega)$ is included in the set $\mathcal{M}(\Omega)$ of Radon bounded measures supported over $\Omega \subseteq \mathbb{R}^d$. Now for any compact set Ω and $\rho \in \mathcal{P}(\Omega)$, it is easy to find from Lemma 1.8 elements of the subdifferential of

$$\mathcal{K}_\rho : (\text{Lip}_{R_\Omega}(\Omega), \|\cdot\|_\infty + \|\cdot\|_{\text{Lip}}) \rightarrow \mathbb{R}.$$

Lemma 1.9. *With the notation and assumptions of Lemma 1.8, assume additionally that $\Omega \subset B(0, R_\Omega)$ is a compact subset of \mathbb{R}^d and that $\psi \in \text{Lip}_{R_\Omega}(\Omega)$. Then, \mathcal{K}_ρ is subdifferentiable at ψ^0 in $(\text{Lip}_{R_\Omega}(\Omega), \|\cdot\|_\infty + \|\cdot\|_{\text{Lip}})$ and*

$$-g_{\#\rho}^0 \in \partial \mathcal{K}_\rho(\psi^0).$$

Proof. Lemma 1.8 already ensures that $-g_{\#\rho}^0$ satisfies the subdifferential inequality for \mathcal{K}_ρ at ψ^0 . The statement thus follows from the fact that $-g_{\#\rho}^0$ is a negative Borel measure with total mass -1 , so that $-g_{\#\rho}^0 \in \mathcal{M}(\Omega)$. \square

Though simple, this lemma is informative: it ensures that in the compact setting, the Kantorovich functional is subdifferentiable in many cases. Moreover, its subdifferential may admit more than one element when ρ has atoms:

Example 1.10. In dimension $d = 1$ and with $\Omega = [-1, 1]$, set $\rho = \delta_0$ and $\psi = 0$. Then for all $x \in \Omega$, the conjugate $\phi = \psi^*$ verifies $\phi(x) = |x|$ so that $\partial\phi(0) = [-1, 1]$. Thus by Lemma 1.8, for any $g \in [-1, 1]$,

$$-\delta_g \in \partial \mathcal{K}_{\delta_0}(\psi).$$

On the other hand, whenever Ω is compact and $\rho \in \mathcal{P}(\Omega)$ is absolutely continuous, Lemma 1.8 describes way fewer elements of the subdifferential of \mathcal{K}_ρ . Indeed, for $\psi \in \text{Lip}_{R_\Omega}(\Omega)$ convex, we know from Remark 1.6 that the conjugate $\phi = \psi^*$ (which is also convex) is differentiable almost-everywhere on $\{\phi < +\infty\}$, that is almost-everywhere on Ω . Therefore, for an absolutely continuous source measure ρ , Lemma 1.8 describes a single element of the subdifferential of \mathcal{K}_ρ at ψ , which may be denoted (slightly abusively)

$$-(\nabla\psi^*)_{\#\rho}.$$

We show in the following lemma that this signed measure is actually the unique element of the subdifferential of \mathcal{K}_ρ at ψ , so that \mathcal{K}_ρ is differentiable at ψ and $-(\nabla\psi^*)_{\#\rho}$ is its Fréchet derivative.

Lemma 1.11. *Let $\Omega \subset B(0, R_\Omega)$ be a compact subset of \mathbb{R}^d . Let $\rho \in \mathcal{P}_{a.c.}(\Omega)$. Then \mathcal{K}_ρ is weakly \mathcal{C}^1 on $(\text{Lip}_{R_\Omega}(\Omega), \|\cdot\|_\infty + \|\cdot\|_{\text{Lip}})$ and its Fréchet derivative reads:*

$$\forall \psi \in \text{Lip}_{R_\Omega}(\Omega), \quad \nabla \mathcal{K}_\rho(\psi) = -(\nabla \psi^*)_{\#} \rho.$$

Proof of Lemma 1.11. Let $\psi \in \text{Lip}_{R_\Omega}(\Omega)$ and let a sequence $(\psi_n)_{n \geq 0} \in \text{Lip}_{R_\Omega}(\Omega)$ that converges to ψ in $(\text{Lip}_{R_\Omega}(\Omega), \|\cdot\|_\infty + \|\cdot\|_{\text{Lip}})$, i.e. such that

$$\|\psi_n - \psi\|_\infty + \|\psi_n - \psi\|_{\text{Lip}} \rightarrow 0.$$

The uniform convergence $\|\psi_n - \psi\|_\infty \rightarrow 0$ entails $\|\psi_n^* - \psi^*\|_\infty \rightarrow 0$, so that (using the dominated convergence theorem), $\mathcal{K}_\rho(\psi_n) \rightarrow \mathcal{K}_\rho(\psi)$. The functional \mathcal{K}_ρ is thus continuous on $(\text{Lip}_{R_\Omega}(\Omega), \|\cdot\|_\infty + \|\cdot\|_{\text{Lip}})$. Now recall that from Lemma 1.8, we have

$$-(\nabla \psi^*)_{\#} \rho \in \partial \mathcal{K}_\rho(\psi).$$

If we can show that the mapping $\psi \mapsto -(\nabla \psi^*)_{\#} \rho$ induces a continuous selection of the subdifferential of \mathcal{K}_ρ (with the dual of $(\text{Lip}_{R_\Omega}(\Omega), \|\cdot\|_\infty + \|\cdot\|_{\text{Lip}})$ equipped with the topology associated to the dual norm¹), then this would show that \mathcal{K}_ρ is \mathcal{C}^1 on $(\text{Lip}_{R_\Omega}(\Omega), \|\cdot\|_\infty + \|\cdot\|_{\text{Lip}})$ and the selection maps to the Fréchet derivative of \mathcal{K}_ρ (see e.g. Corollary 4.5 of (Jourani et al., 2012)). We thus have to show that

$$\sup_{f \in \text{Lip}_{R_\Omega}(\Omega)} \left\{ \langle f(\nabla \psi_n^*) - f(\nabla \psi^*) | \rho \rangle \mid \|f\|_\infty + \|f\|_{\text{Lip}} \leq 1 \right\} \rightarrow 0.$$

But for any $f \in \text{Lip}_{R_\Omega}(\Omega)$ such that $\|f\|_{\text{Lip}} \leq 1$, one has

$$\int_{\Omega} (f(\nabla \psi_n^*) - f(\nabla \psi^*)) d\rho \leq \int_{\Omega} \|\nabla \psi_n^* - \nabla \psi^*\| d\rho.$$

The right-hand side of this inequality does not depend on f and it can be shown to converge to zero using the dominated convergence theorem. Indeed, $(\psi_n^*)_{n \geq 0}, \psi^*$ are R_Ω -Lipschitz so that the integrand is bounded by $2R_\Omega$. Moreover, the pointwise ρ -a.e. convergence of the integrand to 0 may be deduced from the limit $\|\psi_n^* - \psi^*\|_\infty \rightarrow 0$ together with the convexity of ψ^* and ψ_n^* for all $n \geq 0$. Indeed, let $x \in \Omega$ where ψ^* is differentiable (note that ρ -a.e. $x \in \Omega$ works), and let a sequence $(g_n^x)_{n \geq 0}$ be such that $\forall n \geq 0, g_n^x \in \partial \psi_n^*(x)$. Let's show that

$$g_n^x \rightarrow \nabla \psi^*(x).$$

For all $n \geq 0$, $\|g_n^x\| \leq R_\Omega$, so that $(g_n^x)_{n \geq 0}$ is a bounded sequence from which we can extract a converging subsequence. Let's extract a converging subsequence (without relabelling) and denote g_∞^x its limit. By definition of $(g_n^x)_{n \geq 0}$, for any $y \in \Omega$ and for any $n \geq 0$,

$$\langle g_n^x | y - x \rangle \leq \psi_n^*(y) - \psi_n^*(x).$$

If we let n go to ∞ in the last inequality we get for any $y \in \Omega$,

$$\langle g_\infty^x | y - x \rangle \leq \psi^*(y) - \psi^*(x),$$

so that $g_\infty^x \in \partial \psi^*(x) = \{\nabla \psi^*(x)\}$. Thus any converging subsequence of the bounded sequence $(g_n^x)_{n \geq 0}$ converges to $\nabla \psi^*(x)$. This implies that the whole sequence $(g_n^x)_{n \geq 0}$ converges to $\nabla \psi^*(x)$. □

¹Note that this dual norm is equal to the Dudley metric, which is equivalent to the flat norm.

Interestingly, the computation of Lemma 1.11 together with the existence result of Theorem 1.4 allow to recover the following well-known fact (only recovered in the compact setting): for $\Omega \subset \mathbb{R}^d$ compact, $\rho, \mu \in \mathcal{P}(\Omega)$ with ρ absolutely continuous, there exists a proper l.s.c. convex function $\psi : \Omega \rightarrow \mathbb{R}$ such that (ψ^*, ψ) is solution to (DP'), and such ψ is characterized by the first order condition

$$(\nabla\psi^*)_{\#}\rho = \mu.$$

Thus for any ψ satisfying this condition, the mapping $T = \nabla\psi^* : \Omega \rightarrow \Omega$ is a transport map between ρ and μ , and it is an optimal one since we have for $\gamma_T = (\text{id}, T)_{\#}\rho$,

$$(\text{KP}') \geq \int_{\Omega \times \Omega} \langle x|y \rangle d\gamma_T(x, y) = \int_{\Omega} \psi^* d\rho + \int_{\Omega} \psi d\mu \geq (\text{DP}') = (\text{KP}'),$$

where we used the Fenchel-Young equality (Remark 1.6). Therefore when ρ is absolutely continuous, (MP) = (KP) and Monge's formulation admits a (unique) solution that is characterized as being the gradient of a convex function. This corresponds to Brenier's theorem, which holds more generally in the non-compact case:

Theorem 1.12 (Brenier (1991)). *Let $\Omega \subseteq \mathbb{R}^d$. Let $\rho \in \mathcal{P}_{2,a.c.}(\Omega)$ and $\mu \in \mathcal{P}_2(\Omega)$. Then there exists between ρ and μ a unique optimal transport map T and a unique optimal transport plan γ , and these solutions are related by $\gamma = (\text{id}, T)_{\#}\rho$. Further, the map T is the unique transport map that reads as the gradient of a convex function: $T = \nabla\phi$ for all convex functions $\phi : \Omega \rightarrow \bar{\mathbb{R}}$ that satisfy $(\nabla\phi)_{\#}\rho = \mu$.*

Remark 1.13 (Brenier potentials and uniqueness). In this result, the convex functions $\phi : \Omega \rightarrow \bar{\mathbb{R}}$ that satisfy $(\nabla\phi)_{\#}\rho = \mu$ may be called later on *Brenier potentials*. Note that from the proof of Theorem 1.12 below, ϕ is a Brenier potential if and only if $\psi = \phi^*$ is a minimizer of (DP'). Note also that by Theorem 1.12, two Brenier potentials must have that their gradient agree a.e. on $\text{spt}(\rho)$: this entails that whenever $\text{spt}(\rho)$ is the closure of a bounded connected open set, there is a unique Brenier potential up to additive constants (see e.g. Proposition 7.18 of (Santambrogio, 2015)).

Proof of Theorem 1.12. Let $\gamma \in \Gamma(\rho, \mu)$ be optimal for (KP') and let $\psi : \Omega \rightarrow \bar{\mathbb{R}}$ be a proper l.s.c. convex function that is optimal for (DP'). From the equality $\langle \langle \cdot | \cdot \rangle | \gamma \rangle = \langle \psi^* | \rho \rangle + \langle \psi | \mu \rangle$, γ is concentrated on the set $\{(x, y) \in \Omega \times \Omega | \psi^*(x) + \psi(y) = \langle x|y \rangle\}$ (this corresponds to the primal-dual relation described in Remark 1.7). Therefore for any $(x_0, y_0) \in \text{spt}(\gamma)$ such that ψ^* is differentiable in x_0 , one has $y_0 = \nabla\psi^*(x_0)$ by the Fenchel-Young equality case (Remark 1.6). Since ψ^* is proper l.s.c. convex on Ω , it is differentiable ρ -a.e. (Remark 1.6). This entails that γ is induced by the map $\nabla\psi^*$ in the sense that $\gamma = (\text{id}, \nabla\psi^*)_{\#}\rho$. Since γ is an optimal transport plan, $\nabla\psi^*$ is an optimal transport map. Since ψ was not chosen depending on γ , there is a unique optimal transport plan γ and a unique optimal transport map $\nabla\psi^*$ defined ρ -a.e. Finally, any proper l.s.c. convex function $\tilde{\phi} : \Omega \rightarrow \bar{\mathbb{R}}$ that satisfies $(\nabla\tilde{\phi})_{\#}\rho = \mu$ also induces a transport plan $\tilde{\gamma} := (\text{id}, \nabla\tilde{\phi})_{\#}\rho \in \Gamma(\rho, \mu)$ that satisfies

$$\int_{\Omega \times \Omega} \langle x|y \rangle d\tilde{\gamma}(x, y) \leq \int_{\Omega} \tilde{\phi}(x) d\rho(x) + \int_{\Omega} \tilde{\phi}^*(y) d\mu(y) = \int_{\Omega \times \Omega} \langle x|y \rangle d\tilde{\gamma}(x, y),$$

so that by optimality and uniqueness of γ , $\tilde{\gamma} = \gamma$ and $\nabla\phi = \nabla\tilde{\phi}$ ρ -almost-everywhere. \square

Because of the (formal) first order condition

$$\psi \in (\partial\mathcal{K}_\rho)^{-1}(-\mu)$$

that characterizes a minimizer ψ of the dual optimal transport problem (DP') between $\rho, \mu \in \mathcal{P}_2(\mathbb{R}^d)$, we have seen that it is interesting to study the Kantorovich functional \mathcal{K}_ρ to get information about such minimizer, and to get information a posteriori about the solutions to the Monge (MP) or Kantorovich (KP) problems. In particular, in the compact setting, we were able to recover Brenier's theorem through the derivation of the gradient of \mathcal{K}_ρ . This result allows to describe *qualitatively* the solutions to (MP) and (KP) under regularity assumptions on the involved measures. In the following, we wonder if we can push further the study of \mathcal{K}_ρ in order to also get *quantitative* information about the solutions to (DP), (KP) and (MP).

1.2 Strong convexity of the Kantorovich functional

Let $\Omega \subseteq \mathbb{R}^d$. Let's recall here the definition of the Kantorovich functional associated to a given source measure $\rho \in \mathcal{P}_2(\Omega)$:

$$\mathcal{K}_\rho : \begin{cases} \bar{\mathbb{R}}^\Omega & \rightarrow \mathbb{R}, \\ \psi & \mapsto \int_\Omega \psi^* d\rho. \end{cases}$$

As mentioned on several occasions in the preceding section, \mathcal{K}_ρ is a convex functional and a minimizer ψ of the dual quadratic optimal transport problem (DP') between $\rho, \mu \in \mathcal{P}_2(\mathbb{R}^d)$ is characterized formally by the first order condition

$$\psi \in (\partial\mathcal{K}_\rho)^{-1}(-\mu).$$

Still formally, the (multivalued) *mapping* $\mu \mapsto (\partial\mathcal{K}_\rho)^{-1}(-\mu)$ may thus associate to a target measure μ a set of solutions to the dual optimal transport problem (DP') between ρ and μ . Hence, the study of this *mapping* may help get information about solutions to (DP'). In particular, getting (quantitative) continuity estimates for this *mapping* could translate into (quantitative) stability estimates for the solutions to (DP') with respect to μ . As we will see in the rest of this thesis (in particular in Part II), such estimates might be useful, be it theoretically, for the numerical resolution of (DP) or in statistical contexts. The strongest notion of continuity we may hope is Lipschitz continuity. In general, for a convex function F , the Lipschitz continuity of $x \mapsto (\partial F)^{-1}(x)$ is equivalent to a notion of strong convexity of $x \mapsto F(x)$. Let's recall these ideas for a function F defined on \mathbb{R}^d :

Definition 1.14 (Strongly convex function on \mathbb{R}^d). A convex function $F : \mathbb{R}^d \rightarrow \mathbb{R}$ is said (α) -strongly convex if there exists $\alpha > 0$ such that for any $x, y \in \mathbb{R}^d$, for any $g^x \in \partial F(x)$, one has

$$\frac{\alpha}{2} \|x - y\|^2 \leq F(y) - F(x) - \langle g^x | y - x \rangle, \quad (1.2)$$

or, equivalently, for any $t \in [0, 1]$,

$$\frac{t(1-t)\alpha}{2} \|x - y\|^2 \leq (1-t)F(x) + tF(y) - F((1-t)x + ty).$$

Lemma 1.15. *Let $\alpha > 0$ and let $F : \mathbb{R}^d \rightarrow \mathbb{R}$ be an α -strongly convex function. Then for any $x^*, y^* \in \mathbb{R}^d$, one has for $x \in (\partial F)^{-1}(x^*)$ and $y \in (\partial F)^{-1}(y^*)$ that*

$$\|y - x\| \leq \frac{1}{\alpha} \|y^* - x^*\|.$$

In particular, $(\partial F)^{-1}(x^) = \{x\}$ and $(\partial F)^{-1}(y^*) = \{y\}$.*

Proof. By definition, $x^* \in \partial F(x)$ and $y^* \in \partial F(y)$. The strong convexity assumption on F thus ensures

$$\frac{\alpha}{2} \|y - x\|^2 \leq F(y) - F(x) - \langle x^* | y - x \rangle \quad \text{and} \quad \frac{\alpha}{2} \|y - x\|^2 \leq F(x) - F(y) - \langle y^* | x - y \rangle.$$

Summing these inequalities and using Cauchy-Schwartz inequality yields

$$\alpha \|y - x\|^2 \leq \langle y^* - x^* | y - x \rangle \leq \|y^* - x^*\| \|y - x\|. \quad \square$$

Remark 1.16. For any $x, x^* \in \mathbb{R}^d$, we know from Remark 1.6 that $x^* \in \partial F(x)$ if and only if $x \in \partial F^*(x^*)$. Thus from Lemma 1.15, if F is α -strong convex, then F^* is differentiable and $\nabla F^* = (\partial F)^{-1}$ is $\frac{1}{\alpha}$ -Lipschitz continuous. One can easily show the converse, so that these properties are actually equivalent.

Beyond ensuring Lipschitz stability estimates for the inverse gradient map, the strong convexity of a convex functional can be very useful for its (numerical) minimization. Indeed, it leads to efficient gradient methods with fast rates of convergence (Nesterov, 2014; Boyd and Vandenberghe, 2004) and it offers a convenient framework for the study of the gradient flows associated to the functional (Ambrosio et al., 2008; Santambrogio, 2017). We thus focus in the first part of this thesis on the derivation of strong convexity estimates for the Kantorovich functional. This functional is not defined on \mathbb{R}^d but on the set of functions from $\Omega \subseteq \mathbb{R}^d$ to \mathbb{R} . As such, the notion of strong convexity of Definition 1.14 does not apply directly and we will need to adapt it. Before we mention existing strong convexity estimates and present our contributions, let us observe under what conditions we can hope to derive such estimates.

1.2.1 Conditions for strong convexity

In general, the Kantorovich functional is not better than convex. This comes from the fact that for any function $\psi : \mathbb{R}^d \rightarrow \mathbb{R}$ and constant $c \in \mathbb{R}$, $(\psi + c)^* = \psi^* - c$. Thus for $\rho \in \mathcal{P}_2(\mathbb{R}^d)$ and $c \in \mathbb{R}$,

$$\mathcal{K}_\rho(\psi + c) = \mathcal{K}_\rho(\psi) - c,$$

and \mathcal{K}_ρ is affine in t on the path $(\psi + tc)_{t \in [0,1]}$. The notion of strong convexity we look for \mathcal{K}_ρ should thus work *up to additive constants* on ψ . Being aware of this, it is still possible to find examples where \mathcal{K}_ρ fails in general to satisfy a strong convexity estimate. This is for instance the case when the source ρ is discrete:

Example 1.17. Let $\mathcal{X} = \{x_i\}_{1 \leq i \leq N}$ be a finite set of $N \geq 1$ points in \mathbb{R}^d and let ρ, μ be probability measures supported over \mathcal{X} . We consider the quadratic optimal transport problem between ρ and μ . In this problem, we know from Remark 1.7 that we may restrict the domain of a candidate Kantorovich potential and its conjugate to \mathcal{X} without loss of generality. Let's build potentials $\psi^0, \psi^1 : \mathcal{X} \rightarrow \mathbb{R}$ such that \mathcal{K}_ρ is affine on the

path $((1-t)\psi^0 + t\psi^1)_{t \in [0,1]}$, so that it does not satisfy a strong convexity estimate. Consider $\psi^0 : \mathcal{X} \rightarrow \mathbb{R}$, defined for all $i \in \{1, \dots, N\}$ by $\psi^0(x_i) = \frac{\|x_i\|^2}{2}$. Then for any $i \in \{1, \dots, N\}$,

$$\psi^{0*}(x_i) = \max_{1 \leq j \leq N} \langle x_i | x_j \rangle - \frac{\|x_j\|^2}{2} = \frac{\|x_i\|^2}{2}.$$

Denote $\varepsilon = \frac{1}{4} \min_{i \neq j} \|x_i - x_j\|^2$. Then, any potential $\psi^1 : \mathcal{X} \rightarrow \mathbb{R}$ that is such that $\|\psi^1 - \psi^0\|_\infty \leq \varepsilon$ satisfies for any $t \in [0, 1]$

$$\mathcal{K}_\rho((1-t)\psi^0 + t\psi^1) = (1-t)\mathcal{K}_\rho(\psi^0) + t\mathcal{K}_\rho(\psi^1). \quad (1.3)$$

Indeed, consider such potential ψ^1 . Let $t \in [0, 1]$ and introduce $\psi^t = (1-t)\psi^0 + t\psi^1$. Notice that $\|\psi^t - \psi^0\|_\infty \leq t\varepsilon$. Let's show that for any $i \in \{1, \dots, N\}$,

$$\psi^{t*}(x_i) = \langle x_i | x_i \rangle - \psi^t(x_i).$$

This would imply that $\psi^{t*}(x_i) = (1-t)\psi^{0*}(x_i) + t\psi^{1*}(x_i)$ and consequently (1.3). For any $i \in \{1, \dots, N\}$, one has

$$\langle x_i | x_i \rangle - \psi^t(x_i) = \langle x_i | x_i \rangle - \psi^0(x_i) - t(\psi^1 - \psi^0)(x_i) \geq \frac{\|x_i\|^2}{2} - t\varepsilon.$$

But by definition of ε , for any $j \neq i$, $\frac{\|x_i\|^2}{2} \geq \langle x_i | x_j \rangle - \frac{\|x_j\|^2}{2} + 2\varepsilon$, so that for any $j \neq i$,

$$\langle x_i | x_i \rangle - \psi^t(x_i) \geq \langle x_i | x_j \rangle - \frac{\|x_j\|^2}{2} + (2-t)\varepsilon \geq \langle x_i | x_j \rangle - \frac{\|x_j\|^2}{2} + t\varepsilon \geq \langle x_i | x_j \rangle - \psi^t(x_j).$$

This shows that $\psi^{t*}(x_i) = \langle x_i | x_i \rangle - \psi^t(x_i)$, which entails (1.3).

This example precludes from allowing discrete ρ in general. This motivates the choice of an absolutely continuous source measure with bounds on its density, that prevent from approaching the pathological discrete case by approximation arguments. We will make the following assumption on the source measure:

Assumption 1.18. ρ is absolutely continuous w.r.t. the Lebesgue measure and its density is bounded away from zero and infinity on its support.

Another necessary assumption on the source measure ρ can be established from the following fact: strong convexity of \mathcal{K}_ρ is equivalent to strong convexity of the functional being minimized in (DP'). Because strong convexity of a convex function implies in general uniqueness of its minimizers (see e.g. Lemma 1.15), we have to be in a setting where uniqueness (up to additive constants) of the solution to the dual quadratic optimal transport problem (DP') between $\rho \in \mathcal{P}_{2,a.c.}(\Omega)$ and any $\mu \in \mathcal{P}(\Omega)$ is ensured. As noticed in Remark 1.13, this is guaranteed whenever the support of ρ corresponds to a connected domain. In contrast, if ρ admits a support with several connected components, the Brenier potentials from Theorem 1.12 are defined uniquely up to additive constants *only on these connected components*. Thus for a given Brenier potential, adding to it a different constant on each different connected component in such a way that the potential can still be extended to a convex function on \mathbb{R}^d does not affect its optimality, and potentials are thus not unique up to a global additive constant. This motivates the assumption that the source measure has a connected support, with bounds on the *connectedness* of this

support that prevent from approaching with approximation arguments the case where the this support has several connected components. This can be ensured by assuming that the support of ρ has a non-zero (lower bounded) Cheeger constant, or equivalently (Attouch et al., 2014) that ρ satisfies a Poincaré-Wirtinger inequality:

Assumption 1.19. ρ satisfies a Poincaré-Wirtinger inequality: for some $p \geq 1$, there exists a constant $C_{PW}(\rho, p) \in (0, +\infty)$ depending only on ρ and p such that for any $f \in \mathcal{C}^1(\mathbb{R}^d)$,

$$\|f - \mathbb{E}_\rho(f)\|_{L^p(\rho)} \leq C_{PW}(\rho, p) \|\nabla f\|_{L^p(\rho, \mathbb{R}^d)}.$$

We note that the necessity of this assumption may be justified in another manner with considerations from Chapter 2. In this chapter, we will consider the setting where the target measures μ are taken discrete with finite fixed support of size $N \in \mathbb{N}^*$ (which corresponds to the semi-discrete setting). In this setting, we will show that \mathcal{K}_ρ can be seen as a \mathcal{C}^2 function on \mathbb{R}^N , whose second derivative features an Hessian matrix that corresponds to the Laplacian matrix of a weighted graph over the support of μ . The weights of this graph are proportional to the size of the intersections between cells of a certain tessellation of the domain of ρ . Finding a strong convexity estimate in this setting will thus reduce to finding a lower bound on the second smallest eigenvalue of this Laplacian matrix, which will be non zero only if the considered weighted graph is connected, i.e. only if the support domain of ρ is connected (see Section 5 of (Kitagawa et al., 2019) for a similar discussion). Finally, we note that Assumption 1.19 already appeared necessary in (Gunsilius, 2022) for the derivation of statistical convergence rates for the estimation of Kantorovich potentials from empirical samples of ρ and μ (a question that is closely related the strong convexity of \mathcal{K}_ρ , see Chapter 5).

1.2.2 A known case

Under Assumptions 1.18 and 1.19 of above, some authors have already established strong convexity estimates for \mathcal{K}_ρ working near regular enough potentials. The following strong convexity estimate was established in (Hütter and Rigollet (2021), Proposition 10) in the context of the statistical estimation of smooth optimal transport maps. Note that similar computations, due to Ambrosio, were reported in (Gigli (2011), Proposition 3.3) to show the 1/2-Hölder behavior of optimal transport maps w.r.t. their target measure for regular enough source and targets.

Proposition 1.20 (Gigli (2011); Hütter and Rigollet (2021)). *Let $\rho \in \mathcal{P}_{2,a.c.}(\mathbb{R}^d)$ satisfying Assumption 1.19 with $p = 2$. Let $\alpha > 0$ and let $\psi^0, \psi^1 \in \mathcal{C}^1(\mathbb{R}^d)$. Assume that ψ^0 is convex and that ψ^1 is α -strongly convex (Definition 1.14). Then, denoting ϕ^0, ϕ^1 the convex conjugates of ψ^0, ψ^1 respectively, \mathcal{K}_ρ verifies*

$$\frac{\alpha}{2C_\rho} \text{Var}_\rho(\phi^1 - \phi^0) \leq \frac{\alpha}{2} \|\nabla \phi^1 - \nabla \phi^0\|_{L^2(\rho, \mathbb{R}^d)}^2 \leq \mathcal{K}_\rho(\psi^1) - \mathcal{K}_\rho(\psi^0) + \langle \psi^1 - \psi^0 | (\nabla \phi^0)_\# \rho \rangle,$$

where $C_\rho = C_{PW}(\rho, 2)$ is the Poincaré constant from Assumption 1.19.

Remark 1.21. We have not rigorously defined what it means for the Kantorovich functional to be strongly convex and yet we refer to the estimate of Proposition 1.20 as a strong convexity estimate for this functional. This can be justified from the fact that,

with the notation of this proposition, we already knew from Lemma 1.8 that $-(\nabla\psi^0)_{\#}\rho$ satisfies the subdifferential inequality for \mathcal{K}_ρ at ψ^0 , that is for any $\psi^1 \in \mathcal{C}^1(\mathbb{R}^d)$,

$$0 \leq \mathcal{K}_\rho(\psi^1) - \mathcal{K}_\rho(\psi^0) + \langle \psi^1 - \psi^0 | (\nabla\phi^0)_{\#}\rho \rangle.$$

The estimate of Proposition 1.20 quantifies the gap in this subdifferential inequality and as such, it is reminiscent of the definition of strong convexity given in (1.2) for functions defined on \mathbb{R}^d . The presence of a variance instead of a squared L^2 norm is to be expected because of the contravariance of the Kantorovich functional with respect to addition of a constant to its argument. The issue we face for a rigorous definition (or more exactly verification) of the strong convexity of the Kantorovich functional is that in the non-compact setting, we have not described the subdifferential of \mathcal{K}_ρ at $\psi^0 : \mathbb{R}^d \rightarrow \mathbb{R}$ (essentially because in such setting, there is no natural normed space to take as domain of definition of the Kantorovich functional, see Section 1.1.3). Nonetheless, in the rest of this thesis, the strong convexity estimates we will derive for the Kantorovich functional will all take the form of the one presented in Proposition 1.20. Though not entirely rigorously establishing strong convexity of \mathcal{K}_ρ , these estimates will find rigorous applications in Part II of this thesis.

Proof of Proposition 1.20. Using that ρ satisfies Assumption 1.19 with $p = 2$ and that ϕ^0, ϕ^1 are differentiable ρ -almost everywhere as convex functions, we have:

$$\text{Var}_\rho(\phi^1 - \phi^0) \leq C_{PW}(\rho, 2) \int \|\nabla\phi^1 - \nabla\phi^0\|^2 d\rho.$$

In turn, the strong convexity assumption on ψ^1 ensures the bound:

$$\frac{\alpha}{2} \int_{\mathbb{R}^d} \|\nabla\phi^1 - \nabla\phi^0\|^2 d\rho \leq \int (\psi^1(\nabla\phi^0) - \psi^1(\nabla\phi^1) - \langle \nabla\psi^1(\nabla\phi^1) | \nabla\phi^0 - \nabla\phi^1 \rangle) d\rho.$$

Now using the Fenchel-Young equality (Remark 1.6), $\nabla\psi^1(\nabla\phi^1) = \text{id}$ and $\phi^1 = \langle \text{id} | \nabla\phi^1 \rangle - \psi^1(\nabla\phi^1)$. We thus have:

$$\frac{\alpha}{2} \int_{\mathbb{R}^d} \|\nabla\phi^1 - \nabla\phi^0\|^2 d\rho \leq \int (\psi^1(\nabla\phi^0) + \phi^1 - \langle \text{id} | \nabla\phi^0 \rangle) d\rho.$$

Again, by Fenchel-Young, $\langle \text{id} | \nabla\phi^0 \rangle = \phi^0 + \psi^0(\nabla\phi^0)$, so that:

$$\begin{aligned} \frac{\alpha}{2} \int_{\mathbb{R}^d} \|\nabla\phi^1 - \nabla\phi^0\|^2 d\rho &\leq \int (\psi^1(\nabla\phi^0) + \phi^1 - \phi^0 - \psi^0(\nabla\phi^0)) d\rho \\ &= \mathcal{K}_\rho(\psi^1) - \mathcal{K}_\rho(\psi^0) + \langle \psi^1 - \psi^0 | (\nabla\phi^0)_{\#}\rho \rangle. \quad \square \end{aligned}$$

Remark 1.22. Brenier's theorem (Theorem 1.12) applies here: for $k \in \{0, 1\}$, $\nabla\phi^k$ corresponds to the optimal transport map between ρ and $(\nabla\phi^k)_{\#}\rho$. The proof of Proposition 1.20 thus establishes an upper-bound on the $L^2(\rho, \mathbb{R}^d)$ distance between the transport maps $\nabla\phi^0$ and $\nabla\phi^1$, which can in turn be leveraged to establish the stability of these transport maps w.r.t. their target measures (Gigli, 2011). The Poincaré inequality allows here to translate this bound to a bound on the (conjugate of) the potentials ψ^0, ψ^1 . In the coming chapters, estimates will rather be directly obtained in terms of ψ^0, ψ^1 or their conjugates and a Gagliardo-Nirenberg type inequality will be necessary in order to translate these bounds to bounds on the optimal transport maps $\nabla\phi^0, \nabla\phi^1$ (see Chapter 5).

Proposition 1.20 is straightforward to obtain, but it might not be satisfactory in practice. Indeed, it relies on the strong assumption that ψ^1 is strongly convex, which, seeing ϕ^1 as a Brenier potential, is equivalent to require that the optimal transport map $\nabla\phi^1$ is Lipschitz continuous (Lemma 1.15). Since we assumed the source measure to be supported on a connected set in Assumption 1.19, having $\nabla\phi^1$ Lipschitz continuous requires at least that the target measure $\mu^1 = (\nabla\phi^1)_\# \rho$ is also supported on a connected set. Such an assumption is rarely satisfied in practice since it does not apply for instance when μ^1 is discretely supported. In addition, to prove that $\nabla\phi^1$ is Lipschitz, one must invoke the regularity theory for optimal transport maps, which requires very strong assumptions on the measure μ^1 , in particular that its support is convex. Note however that one may be able to prove weaker regularity results on the Brenier potentials under weaker assumptions on the target measures: for instance, \mathcal{C}^1 -type regularity results are derived in (Jabin et al., 2021) for the Kantorovich potentials between possibly discrete probability measures in the plane. This might constitute another direction for the derivation of the strong convexity of the Kantorovich functional from regularity estimates on the potentials.

In the following chapters, we show that more general strong convexity estimates for \mathcal{K}_ρ can be derived under less stringent assumptions on the involved potentials. Let us summarize our results in the next subsection.

1.2.3 Contributions and outline of Part I

In the following chapters, several new strong convexity estimates are derived for the Kantorovich functional and a variant. We quote here the most representative and general of these estimates, proven in Chapter 3.

Theorem (Theorem 3.1). *Let ρ be a probability density over a compact convex set \mathcal{X} , satisfying $0 < m_\rho \leq \rho \leq M_\rho$. Let $\mu^0, \mu^1 \in \mathcal{P}_2(\mathbb{R}^d)$. For $k \in \{0, 1\}$, denote ϕ^k a Brenier potential for the optimal transport between ρ and μ^k . Assume that*

$$\forall k \in \{0, 1\}, \quad -\infty < m_\phi \leq \min_{\mathcal{X}} \phi^k \leq \max_{\mathcal{X}} \phi^k \leq M_\phi < +\infty.$$

Then the convex conjugates ψ^0 and ψ^1 of ϕ^0 and ϕ^1 verify

$$\frac{1}{C_d(M_\phi - m_\phi)} \frac{m_\rho^2}{M_\rho^2} \text{Var}_{\frac{1}{2}(\mu^0 + \mu^1)}(\psi^1 - \psi^0) \leq \mathcal{K}_\rho(\psi^1) - \mathcal{K}_\rho(\psi^0) + \langle \psi^1 - \psi^0 | \mu^0 \rangle, \quad (1.4)$$

where $C_d = e(d+1)2^{d-1}$.

Remark 1.23 (Convexity of the support of the source measure). The convexity of the support of the source ρ is a very restrictive assumption. However, we show in Corollary 1.31 presented below that an estimate of the form of (1.4) may be derived when ρ satisfies a Poincaré-Wirtinger inequality and its support corresponds to a finite union of convex sets. This result is proven properly in Section 1.4 where explicit bounds on the constants are given. Note that we conjecture that strong-convexity estimates of the form of (1.4) actually hold for any absolutely continuous source measure ρ with density bounded away from zero and infinity on its support and satisfying a Poincaré-Wirtinger inequality.

Corollary (Corollary 1.31). *Let ρ be a probability density over a compact set \mathcal{X} , satisfying $0 < m_\rho \leq \rho \leq M_\rho$. Assume that ρ satisfies a Poincaré-Wirtinger inequality with*

$p = 1$ (see Assumption 1.19) and that \mathcal{X} corresponds to a connected finite union of convex sets. Let $\mu^0, \mu^1 \in \mathcal{P}_2(\mathbb{R}^d)$. For $k \in \{0, 1\}$, denote ϕ^k a Brenier potential for the optimal transport between ρ and μ^k . Assume that

$$\forall k \in \{0, 1\}, \quad -\infty < m_\phi \leq \min_{\mathcal{X}} \phi^k \leq \max_{\mathcal{X}} \phi^k \leq M_\phi < +\infty.$$

Then ψ^0 and ψ^1 verify

$$\frac{1}{C_d C_\rho (M_\phi - m_\phi)} \frac{m_\rho^2}{M_\rho^2} \text{Var}_\rho(\phi^1 - \phi^0) \leq \mathcal{K}_\rho(\psi^1) - \mathcal{K}_\rho(\psi^0) + \langle \psi^1 - \psi^0 | (\nabla \phi^0)_{\#} \rho \rangle.$$

where $C_d = e(d+1)2^d$ and C_ρ is a constant depending on the density of ρ and the geometry of its support.

Remark 1.24 (Variance). The left-hand side of (1.4) involves the variance of $\psi^1 - \psi^0$ instead of a squared L^2 norm. As already mentioned in Remark 1.21, this is to be expected because of the contravariance of the Kantorovich functional under addition of a constant. The choice of $\frac{1}{2}(\mu^0 + \mu^1)$ as the reference measure for the variance term in inequality (1.4) may seem unnatural, but we note that there is no natural reference measure on the target. The choice of $\frac{1}{2}(\mu^0 + \mu^1)$ as the reference measures proves relevant for establishing the strong convexity estimate in terms of the Brenier potentials ϕ^0, ϕ^1 . Indeed, Proposition 1.30 presented in Section 1.4 below especially asserts that $\text{Var}_{\frac{1}{2}(\mu^0 + \mu^1)}(\psi^1 - \psi^0) \geq \frac{1}{2} \text{Var}_\rho(\phi^1 - \phi^0)$. We also note that, as detailed in the proof of Theorem 3.1, the left-hand side of inequality (1.4) could actually be replaced by the quantity

$$\frac{1}{e(M_\phi - m_\phi)} \frac{m_\rho}{M_\rho} \int_0^1 \text{Var}_{\mu^t}(\psi^1 - \psi^0) dt,$$

where for $t \in [0, 1]$, $\mu^t = \nabla((1-t)\psi^0 + t\psi^1)_{\#} \rho$ interpolates between μ^0 and μ^1 . This inequality is tighter, but the interpolation $t \mapsto \mu_t$ has no simple interpretation and is quite difficult to manipulate. In particular, this curve is *not* a generalized geodesic in the sense of Ambrosio, Gigli, Savaré (Ambrosio et al., 2008).

Remark 1.25 (Optimality of exponents). Estimate (1.4) is optimal in term of exponent of $\text{Var}_{\frac{1}{2}(\mu^0 + \mu^1)}(\psi^1 - \psi^0)$. Indeed in dimension $d = 1$, for $\varepsilon \geq 0$, denote $\psi^\varepsilon : y \mapsto \frac{1}{2}(y - \varepsilon)^2$. Then for ρ the uniform probability measure on the segment $[0, 1]$ and $\mu^\varepsilon = (\nabla(\psi^\varepsilon))_{\#} \rho$, one can show that for $\varepsilon \leq 1$, both $\text{Var}_{\frac{1}{2}(\mu^0 + \mu^\varepsilon)}(\psi^\varepsilon - \psi^0)$ and $\mathcal{K}_\rho(\psi^\varepsilon) - \mathcal{K}_\rho(\psi^0) + \langle \psi^\varepsilon - \psi^0 | \mu^0 \rangle$ are of the order of ε^2 .

Remark 1.26 (Oscillation of the Brenier potentials). In Chapter 5, we will observe that assuming for $k \in \{0, 1\}$ that μ^k admits a finite moment of order $p > d$ constitutes a sufficient condition for ensuring that the Brenier potential ϕ^k has a bounded oscillation on \mathcal{X} (Proposition 5.7). This result is found in (Berman and Berndtsson, 2013) and is a direct consequence of Morrey's inequality (Theorem 11.34 and Theorem 12.15 in (Leoni, 2009)), that ensures for any $p > d$ the embedding $W^{1,p}(\mathcal{X}) \subset C^{0,1-\frac{d}{p}}(\mathcal{X})$. We will also notice in Remark 5.10 that such assumption is nearly optimal, since for any $p < d$, one can find a measure μ with finite moment of order p and with unbounded Brenier potential between ρ and μ .

Results of the form of Theorem 3.1 were obtained using different techniques and in diverse settings in these three works: (Mérigot et al., 2020), (Delalande and Mérigot, 2021) and (Delalande, 2022). We present the results of these articles in the three following chapters:

- In Chapter 2, we present a preliminary version of Theorem 3.1 that was derived in (Mérigot et al., 2020) and that only worked for compactly supported target measures. The proof of this result relies on the Brunn-Minkowski inequality, a discrete Poincaré-Wirtinger inequality and an approximation argument that consists in taking the measures μ^0, μ^1 discretely supported (hence we name this chapter *A semi-discrete approach*). Even though the result from this chapter is weaker than Theorem 3.1, we present it because its arguments are less technical.
- In Chapter 3, we give the proof of Theorem 3.1 that was derived in (Delalande and Mérigot, 2021). This result is deduced from the Brascamp-Lieb concentration inequality and an approximation argument that consists in taking the target measures μ^0, μ^1 absolutely continuous and regular enough (hence we name this chapter *A continuous approach*).
- In Chapter 4, we present a strong convexity estimate for the *entropic* Kantorovich functional, that appears in the dual of the entropy-regularized optimal transport problem. This estimate follows from the Prékopa-Leindler inequality and was derived in (Delalande, 2022). It might be useful per se for the study of entropy-regularized transport problems (we will use it for instance in Chapter 7). We mention however in Chapter 4 how one can recover the strong convexity estimate of Theorem 3.1 by letting the entropic regularization parameter go to zero in this estimate (hence we name this chapter *An entropic approach*).

Note from this outline that the main ingredients of the proofs of the different strong convexity estimates to be presented are the Brunn-Minkowski inequality, the (concentration) Brascamp-Lieb inequality and the Prékopa-Leindler inequality. Because of the importance of these geometric and functional inequalities for these estimates, we present them in the next section.

1.3 Some geometric and functional inequalities

In this section, we report some well-known geometric and functional inequalities that constitute the main ingredients for the proofs of the strong-convexity estimates that we present in Chapters 2, 3 and 4. There are many reviews in the literature that gather possibly more general statements of these inequalities, proofs and potential consequences in other areas of mathematics. Here, we limit ourselves to the statement of these inequalities in the settings of our interest and discuss succinctly these results. We refer to the following surveys for more extended treatments: (Ball, 1997), (Gardner, 2001), (Schneider, 2013), (Bakry et al., 2014).

1.3.1 Brunn-Minkowski inequality

The Brunn-Minkowski inequality (Brunn, 1887; Minkowski, 1896; Lusternik, 1935) is a geometric inequality that ensures the log-concavity of the Lebesgue measure $|\cdot|$ on \mathbb{R}^d :

Theorem 1.27 (Brunn-Minkowski inequality). *Let A, B be two nonempty compact subsets of \mathbb{R}^d . Then, denoting $A + B$ the Minkowski sum of A and B , it holds*

$$|A + B|^{1/d} \geq |A|^{1/d} + |B|^{1/d},$$

or, equivalently, for any $\lambda \in [0, 1]$,

$$|(1 - \lambda)A + \lambda B| \geq |A|^{1-\lambda} |B|^\lambda.$$

When A and B are two axis-aligned boxes, e.g. $A = \prod_{i=1}^d [0, a_i]$ and $B = \prod_{i=1}^d [0, b_i]$, one has easily $\prod (a_i + b_i)^{1/d} \geq \prod a_i^{1/d} + \prod b_i^{1/d}$ from the inequality of arithmetic and geometric means. A possible proof of Theorem 1.27 consists in using this fact and approximate any nonempty compact sets A and B with finite unions of axis-aligned boxes (Gardner, 2001). The Brunn-Minkowski inequality admits many consequences: it gives much insight on the geometry of convex bodies (see e.g. Grünbaum's inequality (Grünbaum, 1960), that ensures that hyperplanes passing through the centroid of any convex body divide it into not too small portions) and it can be used to prove isoperimetric inequalities or concentration of measure results (Gardner, 2001).

1.3.2 Prékopa-Leindler inequality

The Prékopa-Leindler inequality (Prékopa, 1971; Leindler, 1972; Prékopa, 1973) corresponds to the functional form of the Brunn-Minkowski inequality:

Theorem 1.28 (Prékopa-Leindler inequality). *Let $0 < \lambda < 1$ and $f, g, h : \mathbb{R}^d \rightarrow \mathbb{R}_+$ be such that for any $x, y \in \mathbb{R}^d$,*

$$h((1 - \lambda)x + \lambda y) \geq f(x)^{1-\lambda} g(y)^\lambda.$$

Then it holds

$$\int_{\mathbb{R}^d} h(x) dx \geq \left(\int_{\mathbb{R}^d} f(x) dx \right)^{1-\lambda} \left(\int_{\mathbb{R}^d} g(x) dx \right)^\lambda.$$

The Prékopa-Leindler inequality obviously entails the Brunn-Minkowski inequality by taking $f = \chi_A, g = \chi_B$ and $h = \chi_{(1-\lambda)A + \lambda B}$, where χ_S denotes the indicator function of a set $S \subset \mathbb{R}^d$. However, one possible way to prove the Prékopa-Leindler inequality is to proceed by induction on the dimension d , where the base case $d = 1$ results from the one-dimensional Brunn-Minkowski inequality (Ball, 1997). These two inequalities can therefore be considered to be equivalent. The Prékopa-Leindler inequality finds applications in the theory of log-concave measures (where it ensures that log-concavity is preserved by marginalization) and can also be used to show concentration of measure results (Maurey, 1991, 2004) or transportation inequalities (Bobkov and Ledoux, 2000).

1.3.3 Brascamp-Lieb inequality

The Brascamp-Lieb *concentration* inequality (Brascamp and Lieb, 1976) corresponds to a Poincaré-type estimate. We cite here a version of this inequality that will be adapted to our context, i.e. that concerns log-concave probability measures supported over a compact and convex set \mathcal{X} . This statement is a special case of Corollary 1.3 of (Le Peutrec, 2017), where \mathcal{X} is a convex subset of a Riemannian manifold. We also refer to Section 3.1.1 of (Kolesnikov and Milman, 2017).

Theorem 1.29 (Brascamp-Lieb inequality). *Let $\mathcal{X} \subset \mathbb{R}^d$ be a compact convex set. Let $\phi \in \mathcal{C}^2(\mathcal{X})$ be a strictly convex function. Let ρ be the probability measure defined by*

$d\rho = \frac{1}{Z_\phi} \exp(-\phi)dx$ with $Z_\phi = \int_{\mathcal{X}} \exp(-\phi)dx$. Then every smooth function s on \mathcal{X} verifies:

$$\mathrm{Var}_\rho(s) \leq \mathbb{E}_\rho \langle \nabla s | (D^2\phi)^{-1} \cdot \nabla s \rangle.$$

This inequality was originally formulated in a non-compact setting ($\mathcal{X} = \mathbb{R}^d$) in (Brascamp and Lieb, 1976) and was shown to be a consequence of the Prékopa-Leindler inequality in (Bobkov and Ledoux, 2000). We develop here a possible argument² in this direction, that will inspire our approach to derive a strong-convexity estimate for entropic optimal transport in Chapter 4. This argument is the following: for $u : \mathbb{R}^d \rightarrow \mathbb{R}$ a convex function, define the functional

$$I : u \mapsto \log \int e^{-u^*}.$$

(Note the similarity between I and $\mathcal{K}_\rho : \psi \mapsto \int \psi^* d\rho$.) Then for $\lambda \in [0, 1]$ and u, v convex functions from \mathbb{R}^d to \mathbb{R} , one can deduce from the Prékopa-Leindler inequality that $I((1 - \lambda)u + \lambda v) \geq (1 - \lambda)I(u) + \lambda I(v)$, i.e. I is a concave functional. For $f : \mathbb{R}^d \rightarrow \mathbb{R}$, assuming all the regularity necessary, this entails that the second derivative of $t \mapsto I(u + tf)$ is non-positive. Computing this second derivative, setting $u = \phi^*$ and $f = s \circ \nabla \phi^*$ and using that this derivative is non-positive then yields Theorem 1.29.

We also note that conversely, the Prékopa-Leindler inequality may be deduced from the Brascamp-Lieb inequality by a local computation (Cordero-Erausquin, 2005). Thus, quoting (Carlen et al., 2013), we may state that the Brascamp-Lieb inequality can be seen as the local form of the Brunn-Minkowski inequality for convex bodies.

The Brascamp-Lieb inequality can finally also be seen as an extension of the Poincaré inequality. Indeed for a strongly log-concave measure ρ in Theorem 1.29, i.e. assuming that there exists $\alpha > 0$ such that $D^2\phi \succeq \alpha \mathrm{Id}$, one recovers from the Brascamp-Lieb inequality the following Poincaré inequality:

$$\mathrm{Var}_\rho(s) \leq \frac{1}{\alpha} \mathbb{E}_\rho \|\nabla s\|^2.$$

As the two preceding inequalities, the Brascamp-Lieb inequality and its generalizations can be used to prove concentration and isoperimetry results, see e.g. (Bakry et al., 2014).

1.3.4 Links with optimal transport

Several works have observed that (quadratic) optimal transport can serve as a tool to prove the Brunn-Minkowski and Prékopa-Leindler inequalities. Indeed, McCann recovered the Brunn-Minkowski inequality in (McCann, 1997) as a consequence of the displacement convexity (i.e. convexity along a Wasserstein geodesic, see Chapter A) of a well-chosen energy functional on the space of probability measures. In (McCann, 1997, 1994) and (Barthe, 1997), the Prékopa-Leindler inequality was also deduced from displacement interpolation and the log-concavity of the determinant on the set of non-negative symmetric matrices (deduced in turn from the arithmetic-geometric inequality). This line of works has allowed to find Prékopa-Leindler type inequalities working on the

²Extracted from Exercise 2.2.11 of Bo'az Klartag's lecture notes *Regularity through convexity in high dimensions* (Klartag, 2013).

sphere (Cordero-Erausquin, 1999) or more generally on Riemannian manifolds (Cordero-Erausquin et al., 2001, 2006). We also note that optimal transport approaches have allowed to derive and extend other types of geometric and functional inequalities, such as (geometric) Brascamp-Lieb and Sobolev type inequalities – we refer to Chapter 6 of (Villani, 2003) for a more extended treatment.

In the coming chapters, instead of using optimal transport as a tool to derive geometric and functional inequalities, we take the other way around and use the Brunn-Minkowski, Brascamp-Lieb and Prékopa-Leindler inequalities as a tool to derive properties of optimal transport, namely strong-convexity of the Kantorovich functional.

1.4 Extension to source measures with non-convex support (Proof of Corollary 1.31)

We first mention the following convex analysis fact, that allows to translate the estimate (1.4) expressed in terms of the dual potentials ψ^0, ψ^1 into an estimate expressed in terms of the Brenier potentials ϕ^0, ϕ^1 .

Proposition 1.30. *Let ρ be a probability density over a compact set \mathcal{X} , and let ϕ^0, ϕ^1 be convex functions on \mathcal{X} . For $k \in \{0, 1\}$, denote ψ^k the convex conjugate of ϕ^k and μ^k the image of ρ under $\nabla\phi^k$. Then for any $p > 0$,*

$$\|\phi^1 - \phi^0\|_{L^p(\rho)} \leq \|\psi^1 - \psi^0\|_{L^p(\mu^0 + \mu^1)}.$$

In particular,

$$\frac{1}{2} \text{Var}_\rho(\phi^1 - \phi^0) \leq \text{Var}_{\frac{1}{2}(\mu^0 + \mu^1)}(\psi^1 - \psi^0).$$

Proof of Proposition 1.30. Let $A = \{x \in \mathcal{X} \mid \phi^1(x) \geq \phi^0(x)\}$ and let $x \in A$ where ϕ^1 is differentiable. The Fenchel-Young inequality (and equality) give:

$$\psi^0(\nabla\phi^1(x)) \geq \langle x \mid \nabla\phi^1(x) \rangle - \phi^0(x) = \psi^1(\nabla\phi^1(x)) + \phi^1(x) - \phi^0(x),$$

which thus ensures that for almost every $x \in A$,

$$\psi^0(\nabla\phi^1(x)) - \psi^1(\nabla\phi^1(x)) \geq \phi^1(x) - \phi^0(x) \geq 0.$$

Similarly, for almost every $x \in \mathcal{X} \setminus A$, we have

$$\psi^1(\nabla\phi^0(x)) - \psi^0(\nabla\phi^0(x)) \geq \phi^0(x) - \phi^1(x) \geq 0.$$

From this, we deduce the first statement of the proposition:

$$\begin{aligned} \|\psi^1 - \psi^0\|_{L^p(\mu^0 + \mu^1)}^p &= \int_{\mathcal{X}} |\psi^1(\nabla\phi^0) - \psi^0(\nabla\phi^0)|^p d\rho + \int_{\mathcal{X}} |\psi^1(\nabla\phi^1) - \psi^0(\nabla\phi^1)|^p d\rho \\ &\geq \int_{\mathcal{X} \setminus A} (\psi^1(\nabla\phi^0) - \psi^0(\nabla\phi^0))^p d\rho + \int_A (\psi^0(\nabla\phi^1) - \psi^1(\nabla\phi^1))^p d\rho \\ &\geq \int_{\mathcal{X} \setminus A} (\phi^0 - \phi^1)^p d\rho + \int_A (\phi^1 - \phi^0)^p d\rho = \|\phi^1 - \phi^0\|_{L^p(\rho)}^p. \end{aligned}$$

Let $c \in \mathbb{R}$. Having established the previous inequality for any convex functions ϕ^0, ϕ^1 on \mathcal{X} , we may replace ϕ^0 with $\phi^0 - c$ in this inequality, and consequently replace ψ^0 with $(\phi^0 - c)^* = \psi^0 + c$. This yields thus for any $c \in \mathbb{R}$,

$$\|\phi^1 - \phi^0 + c\|_{L^p(\rho)} \leq \|\psi^1 - \psi^0 - c\|_{L^p(\mu^0 + \mu^1)}.$$

Taking c that achieves the minimum on the right-hand side, for $p = 2$, we get

$$\begin{aligned} \text{Var}_\rho(\phi^1 - \phi^0) &\leq \|\phi^1 - \phi^0 + c\|_{L^2(\rho)}^2 \\ &\leq 2 \|\psi^1 - \psi^0 - c\|_{L^2(\frac{1}{2}(\mu^0 + \mu^1))}^2 = 2\text{Var}_{\frac{1}{2}(\mu^0 + \mu^1)}(\psi^1 - \psi^0). \quad \square \end{aligned}$$

Under the notation and assumptions of Theorem 3.1, Proposition 1.30 directly entails the following strong convexity estimate:

$$\frac{1}{2C_d(M_\phi - m_\phi)} \frac{m_\rho^2}{M_\rho^2} \text{Var}_\rho(\phi^1 - \phi^0) \leq \mathcal{K}_\rho(\psi^1) - \mathcal{K}_\rho(\psi^0) + \langle \psi^1 - \psi^0 | \mu^0 \rangle.$$

A consequence of this new estimate is the following corollary, which is an extension of Theorem 3.1 to a non-convexly supported source measure ρ . The idea of this result is to upper-bound the variance with respect to ρ by a sum of local variances with respect to convexly supported measures: this is allowed when we assume that ρ satisfies a Poincaré-Wirtinger inequality (Assumption 1.19) and that its support corresponds to a union of convex sets. This result was originally obtained in (Carlier et al., 2022).

Corollary 1.31. *Let ρ be a probability density over a connected compact set $\mathcal{X} \subset B(0, R_\mathcal{X})$, satisfying $0 < m_\rho \leq \rho \leq M_\rho$. Assume that ρ satisfies a Poincaré-Wirtinger inequality with $p = 1$ (see assumption 1.19) and that there exists $N \geq 1$ distinct convex sets $(C_i)_{1 \leq i \leq N}$ such that $\mathcal{X} = \bigcup_{i=1}^N C_i$. Also assume that*

$$\varepsilon := \min \left(\min_{i,j | C_i \cap C_j \neq \emptyset} \rho(C_i \cap C_j), \min_i \rho(C_i \setminus \bigcup_{j \neq i} C_j) \right) > 0.$$

Let $\mu^0, \mu^1 \in \mathcal{P}_2(\mathbb{R}^d)$. For $k \in \{0, 1\}$, denote ϕ^k a Brenier potential for the optimal transport between ρ and μ^k . Assume that

$$\forall k \in \{0, 1\}, \quad -\infty < m_\phi \leq \min_{\mathcal{X}} \phi^k \leq \max_{\mathcal{X}} \phi^k \leq M_\phi < +\infty.$$

Then ψ^0 and ψ^1 verify

$$\frac{1}{C_d C_\rho (M_\phi - m_\phi)} \frac{m_\rho^2}{M_\rho^2} \text{Var}_\rho(\phi^1 - \phi^0) \leq \mathcal{K}_\rho(\psi^1) - \mathcal{K}_\rho(\psi^0) + \langle \psi^1 - \psi^0 | (\nabla \phi^0)_\# \rho \rangle.$$

where $C_d = e(d+1)2^d$ and $C_\rho = \left(N^2 + N^7 M_\rho^3 s_{d-1}^3 R_\mathcal{X}^{3(d-1)} C_{PW}(\rho, 1)^3 \varepsilon^{-6} \right)$.

Proof. Let's denote for now $f = \phi^1 - \phi^0$. We will first exploit a discrete Laplacian over \mathcal{X} in order to upper bound $\text{Var}_\rho(f)$ by a sum of variances of f w.r.t. probability measures supported over the convex sets $(C_i)_i$. We will then use Theorem 3.1 to conclude.

For any $i \in \{1, \dots, N\}$, we denote $\rho_i = \frac{1}{\rho(C_i)}\rho|_{C_i}$ and $m_i = \int_{C_i} f d\rho_i$. Then one has the following bound:

$$\begin{aligned}
\text{Var}_\rho(f) &= \frac{1}{2} \int_{\mathcal{X} \times \mathcal{X}} (f(x) - f(y))^2 d\rho(x) d\rho(y) \\
&\leq \frac{1}{2} \sum_{i,j} \int_{C_i \times C_j} (f(x) - f(y))^2 d\rho(x) d\rho(y) \\
&= \frac{1}{2} \sum_{i,j} \int_{C_i \times C_j} (f(x) - m_i + m_i - m_j + m_j - f(y))^2 d\rho(x) d\rho(y) \\
&= \left(\sum_i \rho(C_i) \right) \sum_i \int_{C_i} (f(x) - m_i)^2 d\rho(x) + \frac{1}{2} \sum_{i,j} (m_i - m_j)^2 \rho(C_i) \rho(C_j) \\
&= \left(\sum_i \rho(C_i) \right) \sum_i \rho(C_i) \text{Var}_{\rho_i}(f) + \frac{1}{2} \sum_{i,j} (m_i - m_j)^2 \rho(C_i) \rho(C_j). \tag{1.5}
\end{aligned}$$

Introduce the graph $G = (\{C_i\}_{1 \leq i \leq N}, \{w_{ij}\}_{1 \leq i, j \leq N})$ with vertices $\{C_i\}_{1 \leq i \leq N}$ and weighted edges $\{w_{ij}\}_{1 \leq i, j \leq N}$ defined by

$$\forall i, j \in \{1, \dots, N\}, \quad w_{ij} = \rho(C_i \cap C_j).$$

By construction, this graph has a single connected component. We introduce the weighted Laplacian matrix $L \in \mathbb{R}^{N \times N}$ of G as follows:

$$\forall i, j \in \{1, \dots, N\}, \quad L_{ij} = \begin{cases} \sum_k w_{ik} & \text{if } i = j, \\ -w_{ij} & \text{else.} \end{cases}$$

Then L is a symmetric and positive semi-definite matrix. Its null space is made of constant vectors and we denote $\lambda_2(L)$ its second smallest eigenvalue, which is non-zero. Denoting $m = (m_i)_{1 \leq i \leq N} \in \mathbb{R}^N$, we introduce $\bar{m} = \left(\frac{1}{N} \sum_i m_i\right) \mathbf{1}_N \in \mathbb{R}^N$ the constant vector whose coordinates equal the mean of m (we use $\mathbf{1}_N = (1)_{1 \leq i \leq N} \in \mathbb{R}^N$). Notice that $m - \bar{m}$ is in the orthogonal to the null space of L , ensuring the following bound:

$$\begin{aligned}
\frac{1}{2} \sum_{i,j} (m_i - m_j)^2 \rho(C_i) \rho(C_j) &\leq N^2 \frac{1}{2} \sum_{i,j} (m_i - m_j)^2 \frac{1}{N^2} \\
&= N \|m - \bar{m}\|^2 \\
&\leq \frac{N}{\lambda_2(L)} \langle m - \bar{m} | L(m - \bar{m}) \rangle \\
&= \frac{N}{\lambda_2(L)} \sum_{i,j} w_{ij} (m_i^2 - m_i m_j) \\
&= \frac{N}{\lambda_2(L)} \sum_{i,j} \frac{w_{ij}}{2} (m_i - m_j)^2. \tag{1.6}
\end{aligned}$$

But for any i, j such that $w_{ij} > 0$, denoting $m_{i \cap j} = \frac{1}{\rho(C_i \cap C_j)} \int_{C_i \cap C_j} f d\rho$, one has

$$\frac{1}{2} (m_i - m_j)^2 \leq (m_{i \cap j} - m_i)^2 + (m_{i \cap j} - m_j)^2.$$

And for such i, j ,

$$\begin{aligned} (m_{i \cap j} - m_i)^2 &= \left(\frac{1}{\rho(C_i \cap C_j)} \int_{C_i \cap C_j} (f - m_i) d\rho \right)^2 \\ &\leq \frac{1}{\rho(C_i \cap C_j)} \int_{C_i} (f - m_i)^2 d\rho \\ &= \frac{\rho(C_i)}{w_{ij}} \mathbb{V}\text{ar}_{\rho_i}(f), \end{aligned}$$

where we used Jensen's inequality and the fact that $C_i \cap C_j \subset C_i$. A similar bound can be shown for $(m_{i \cap j} - m_j)^2$, and plugging these into (1.6) yields

$$\begin{aligned} \frac{1}{2} \sum_{i,j} (m_i - m_j)^2 \rho(C_i) \rho(C_j) &\leq \frac{N}{\lambda_2(L)} \sum_i \sum_{j|C_i \cap C_j \neq \emptyset} (\rho(C_i) \mathbb{V}\text{ar}_{\rho_i}(f) + \rho(C_j) \mathbb{V}\text{ar}_{\rho_j}(f)) \\ &\leq \frac{2N^2}{\lambda_2(L)} \sum_i \rho(C_i) \mathbb{V}\text{ar}_{\rho_i}(f). \end{aligned}$$

Injecting this into (1.5) yields

$$\mathbb{V}\text{ar}_{\rho}(f) \leq \left(N + \frac{2N^2}{\lambda_2(L)} \right) \sum_i \rho(C_i) \mathbb{V}\text{ar}_{\rho_i}(f). \quad (1.7)$$

Now recalling that $f = \phi^1 - \phi^0$, we have by Theorem 3.1 and Proposition 1.30 for any $i \in \{1, \dots, N\}$ that

$$\frac{1}{C_d(M_\phi - m_\phi)} \frac{m_\rho^2}{M_\rho^2} \mathbb{V}\text{ar}_{\rho_i}(\phi^1 - \phi^0) \leq \mathcal{K}_{\rho_i}(\psi^1) - \mathcal{K}_{\rho_i}(\psi^0) + \langle \psi^1 - \psi^0 | (\nabla \phi^0)_{\#} \rho_i \rangle,$$

where $C_d = e(d+1)2^d$. Weighting this last inequality with $\rho(C_i)$ and summing over $i \in \{1, \dots, N\}$, this raises

$$\frac{1}{NC_d(M_\phi - m_\phi)} \frac{m_\rho^2}{M_\rho^2} \sum_{i=1}^N \rho(C_i) \mathbb{V}\text{ar}_{\rho_i}(\phi^1 - \phi^0) \leq \mathcal{K}_\rho(\psi^1) - \mathcal{K}_\rho(\psi^0) + \langle \psi^1 - \psi^0 | (\nabla \phi^0)_{\#} \rho \rangle,$$

Using (1.7) eventually gives

$$\frac{1}{C_d C_{N,L}(M_\phi - m_\phi)} \frac{m_\rho^2}{M_\rho^2} \mathbb{V}\text{ar}_{\rho}(\phi^1 - \phi^0) \leq \mathcal{K}_\rho(\psi^1) - \mathcal{K}_\rho(\psi^0) + \langle \psi^1 - \psi^0 | (\nabla \phi^0)_{\#} \rho \rangle. \quad (1.8)$$

where $C_{N,L} = \left(N^2 + \frac{2N^3}{\lambda_2(L)} \right)$.

Finally, we lower bound the eigenvalue $\lambda_2(L)$ in terms of the Poincaré-Wirtinger constant $C_{PW}(\rho, 1)$ of the probability measure ρ . For this, we will rely on an intermediate quantity: the weighted Cheeger constant ρ , defined by

$$h(\rho) = \inf_{A \subset \mathcal{X}} \frac{|\partial A|_\rho}{\min(\rho(A), \rho(\mathcal{X} \setminus A))},$$

where $|\partial A|_\rho = \int_{\partial A \cap \text{int}(\mathcal{X})} \rho(x) d\mathcal{H}^{d-1}(x)$ and where the infimum is taken over Lipschitz domains $A \subset \text{int}(\text{spt}(\rho))$ with boundary of finite \mathcal{H}^{d-1} -measure. Quoting (Kitagawa

et al., 2019) (Lemma 5.3), we note that this constant is positive because ρ admits a finite Poincaré-Wirtinger constant $C_{PW}(\rho, 1)$, by the inequality

$$h(\rho) \geq \frac{2}{C_{PW}(\rho, 1)}, \quad (1.9)$$

which is a result coming from properties of functions with bounded variations (Attouch et al., 2014). We first use the following result from (Friedland and Nabben, 2000) (Corollary 2.2):

$$\lambda_2(L) \geq \frac{1}{2} \left(\min_i \sum_j w_{ij} \right) i(G)^2, \quad (1.10)$$

where $i(G)$ corresponds to the weighted Cheeger constant of G :

$$i(G) = \min_{\emptyset \neq U \subseteq \{1, \dots, N\}} \frac{\sum_{i \in U, j \in \bar{U}} w_{ij}}{\min \left(\sum_{i \in U} \sum_{j \neq i} w_{ij}, \sum_{i \in \bar{U}} \sum_{j \neq i} w_{ij} \right)},$$

where $\bar{U} = \{C_i\}_i \setminus U$. There are two terms in (1.10) that need to be lower bounded in terms of $h(\rho)$. We first lower bound $\min_i \sum_j w_{ij}$. Notice that for any $i \in \{1, \dots, N\}$,

$$\sum_j w_{ij} = \sum_j \rho(C_i \cap C_j).$$

By assumption, we have for any j such that $C_i \cap C_j \neq \emptyset$, $\rho(C_i \cap C_j) \geq \varepsilon$ (and such a j always exists since \mathcal{X} is assumed to be connected). On the other hand, for any such j , noticing that C_i is a convex subset of $B(0, R_{\mathcal{X}})$, one has

$$|\partial C_i \cap C_j|_{\rho} \leq M_{\rho} s_{d-1} R_{\mathcal{X}}^{d-1},$$

where s_{d-1} denotes the surface area of the unit sphere in \mathbb{R}^d . Therefore, for j such that $w_{ij} > 0$, we get the bound

$$\rho(C_i \cap C_j) \geq \frac{\varepsilon}{M_{\rho} s_{d-1} R_{\mathcal{X}}^{d-1}} |\partial C_i \cap C_j|_{\rho}.$$

Hence for any $i \in \{1, \dots, N\}$,

$$\sum_j w_{ij} \geq \frac{\varepsilon}{M_{\rho} s_{d-1} R_{\mathcal{X}}^{d-1}} \sum_{j | C_i \cap C_j \neq \emptyset} |\partial C_i \cap C_j|_{\rho} = \frac{\varepsilon}{M_{\rho} s_{d-1} R_{\mathcal{X}}^{d-1}} |\partial C_i|_{\rho}.$$

Now using that $|\partial C_i|_{\rho} \geq h(\rho) \min(\rho(C_i), \rho(\mathcal{X} \setminus C_i)) \geq h(\rho)\varepsilon$, we have the bound

$$\min_i \sum_j w_{ij} \geq \frac{\varepsilon^2}{M_{\rho} s_{d-1} R_{\mathcal{X}}^{d-1}} h(\rho). \quad (1.11)$$

Similarly, we now lower bound $i(G)$ in terms of $h(\rho)$. For a given subset $\emptyset \neq U \subseteq \{1, \dots, N\}$, define $A = \bigcup_{i \in U} C_i$. Then, as in what precedes we can notice that

$$\begin{aligned} \sum_{i \in U, j \in \bar{U}} w_{ij} &= \sum_{i \in U, j \in \bar{U}} \rho(C_i \cap C_j) \geq \sum_{i \in U, j \in \bar{U}} \frac{\varepsilon}{M_{\rho} s_{d-1} R_{\mathcal{X}}^{d-1}} |\partial C_i \cap C_j|_{\rho} \\ &= \frac{\varepsilon}{M_{\rho} s_{d-1} R_{\mathcal{X}}^{d-1}} |\partial A|_{\rho}. \end{aligned} \quad (1.12)$$

Then, we have

$$\sum_{i \in U} \sum_{j \neq i} w_{ij} = \sum_{i \in U} \sum_{j \neq i} \rho(C_i \cap C_j) \leq N \sum_{i \in U} \rho(C_i) \leq N^2 \rho(A).$$

Similarly, one can show

$$\sum_{i \in \bar{U}} \sum_{j \neq i} w_{ij} = \sum_{i \in \bar{U}} \sum_{j \neq i} \rho(C_i \cap C_j) \leq N^2 \leq \frac{N^2}{\varepsilon} \rho(\mathcal{X} \setminus A),$$

where we used the assumption that $\min_i \rho(C_i \setminus \cup_{j \neq i} C_j) \geq \varepsilon$. Using that $\varepsilon \leq 1$ the two last bounds yield

$$\min \left(\sum_{i \in U} \sum_{j \neq i} w_{ij}, \sum_{i \in \bar{U}} \sum_{j \neq i} w_{ij} \right) \leq \frac{N^2}{\varepsilon} \min(\rho(A), \rho(\mathcal{X} \setminus A)). \quad (1.13)$$

Hence combining (1.12) and (1.13), we have the bound

$$\begin{aligned} \frac{\sum_{i \in U, j \in \bar{U}} w_{ij}}{\min \left(\sum_{i \in U} \sum_{j \neq i} w_{ij}, \sum_{i \in \bar{U}} \sum_{j \neq i} w_{ij} \right)} &\geq \frac{\varepsilon^2}{M_\rho s_{d-1} R_{\mathcal{X}}^{d-1} N^2} \frac{|\partial A|_\rho}{\min(\rho(A), \rho(\mathcal{X} \setminus A))} \\ &\geq \frac{\varepsilon^2}{M_\rho s_{d-1} R_{\mathcal{X}}^{d-1} N^2} h(\rho). \end{aligned}$$

Minimizing over U on the left-hand side then yields

$$i(G) \geq \frac{\varepsilon^2}{M_\rho s_{d-1} R_{\mathcal{X}}^{d-1} N^2} h(\rho). \quad (1.14)$$

Combining (1.11) and (1.14) into (1.10) finally yields

$$\lambda_2(L) \geq \frac{\varepsilon^6}{2M_\rho^3 s_{d-1}^3 R_{\mathcal{X}}^{3(d-1)} N^4} h(\rho)^3.$$

Using finally the comparison (1.9) between the weighted Cheeger constant of ρ and its Poincaré-Wirtinger constant $C_{PW}(\rho, 1)$, this ensures:

$$\lambda_2(L) \geq \frac{4\varepsilon^6}{M_\rho^3 s_{d-1}^3 R_{\mathcal{X}}^{3(d-1)} N^4 C_{PW}(\rho, 1)^3}.$$

This last bound together with (1.8) thus ensures

$$\frac{1}{C_d C_\rho (M_\phi - m_\phi)} \frac{m_\rho^2}{M_\rho^2} \text{Var}_\rho(\phi^1 - \phi^0) \leq \mathcal{K}_\rho(\psi^1) - \mathcal{K}_\rho(\psi^0) + \langle \psi^1 - \psi^0 | (\nabla \phi^0)_\# \rho \rangle,$$

where $C_\rho = \left(N^2 + \frac{N^7}{2\varepsilon^6} M_\rho^3 s_{d-1}^3 R_{\mathcal{X}}^{3(d-1)} C_{PW}(\rho, 1)^3 \right)$. \square

A semi-discrete approach

Abstract

This chapter reports the proof of a first strong convexity estimate for the Kantorovich functional, established in (Mérigot et al., 2020). This estimate holds when the Kantorovich functional is evaluated on Kantorovich potentials associated to compactly supported target measures. The proof of this strong convexity estimate is first derived in the semi-discrete setting where the targets are assumed to be discrete, and it is generalized by a density argument. The main proof ingredients are the Brunn-Minkowski inequality and a discrete Poincaré-Wirtinger inequality that emanates from the stability analysis of finite volumes discretizations of elliptic PDEs.

2.1 Introduction

In this chapter, we prove a first strong convexity estimate for the Kantorovich functional that appears in the dual quadratic optimal transport problem (see Sections 1.1.3 and 1.2). We recall that the Kantorovich functional associated to a source measure $\rho \in \mathcal{P}_2(\mathbb{R}^d)$ reads

$$\mathcal{K}_\rho : \begin{cases} (\mathbb{R}^d \rightarrow \bar{\mathbb{R}}) & \rightarrow \bar{\mathbb{R}}, \\ \psi & \mapsto \int_{\mathbb{R}^d} \psi^* d\rho. \end{cases}$$

The following estimate was established in (Mérigot et al., 2020):

Theorem 2.1. *Let ρ be a probability density over a compact convex set \mathcal{X} , satisfying $0 < m_\rho \leq \rho \leq M_\rho$. Let \mathcal{Y} be a compact subset of \mathbb{R}^d and let $\mu^0, \mu^1 \in \mathcal{P}(\mathcal{Y})$. For $k \in \{0, 1\}$, denote ϕ^k a Brenier potential between ρ and μ^k . Then the convex conjugates ψ^0 and ψ^1 of ϕ^0 and ϕ^1 verify*

$$\frac{1}{C_{d,\mathcal{X},\mathcal{Y}}} \frac{m_\rho^2}{M_\rho^3} \text{Var}_{\frac{1}{2}(\mu^0 + \mu^1)}(\psi^1 - \psi^0) \leq \mathcal{K}_\rho(\psi^1) - \mathcal{K}_\rho(\psi^0) + \langle \psi^1 - \psi^0 | \mu_0 \rangle, \quad (2.1)$$

where $C_{d,\mathcal{X},\mathcal{Y}} = \omega_d(d+1)2^{d-1} \text{diam}(\mathcal{X})^{d+1} \text{diam}(\mathcal{Y})$ and ω_d denotes the volume of the d -dimensional unit ball.

This result requires that the probability measures μ^0, μ^1 are supported on a known compact set in order to use a Poincaré-Wirtinger type inequality to find a lower bound on

the *Hessian* of the Kantorovich functional. In terms of the Brenier potentials between ρ and μ^0, μ^1 (Remark 1.13), this is equivalent to require that these potentials are Lipschitz continuous. In Chapter 3, we will show that similar strong convexity estimates can be found under the weaker assumption that these potentials only have a bounded oscillation on the compact support of ρ , which can in turn be obtained from the assumption that these potentials are Hölder continuous. Theorem 2.1 can thus be seen a preliminary version of Theorem 3.1 presented in Chapter 3.

Remark 2.2 (Numerical optimal transport). As mentioned above, estimate (2.1) follows from an explicit lower-bound on the Hessian of the Kantorovich functional (established in the semi-discrete setting and presented in Proposition 2.4). This corresponds to guarantees on the conditioning on this Hessian, which may be leveraged to improve the analysis of Newton's methods used to solve semi-discrete optimal transport problems (Kitagawa et al., 2019).

Outline. Our strategy to prove Theorem 2.1 is the following: first, we assume that the probability measures μ^0, μ^1 are discretely supported on a common finite set of N points. This places ourselves in the context of semi-discrete optimal transport, with an absolutely continuous source ρ and discrete targets μ^0, μ^1 . In this context, the Kantorovich functional can be seen as a convex function on \mathbb{R}^N for which the gradient and Hessian are known (§2.2). We show that an explicit lower-bound on the smallest non-zero eigenvalue of this Hessian can be found, from which the strong convexity estimate of Theorem 2.1 is deduced in the semi-discrete setting (§2.3). We finally prove Theorem 2.1 in the general case by a density argument (§2.4).

2.2 Semi-discrete optimal transport

In this section and the following, we work in the *semi-discrete* setting, assuming that for $k \in \{0, 1\}$, the measure μ^k of Theorem 2.1 is supported on a (fixed) finite set $\mathcal{Y} = \{y_1, \dots, y_N\}$. For $R_{\mathcal{Y}}$ such that $\mathcal{Y} \subset B(0, R_{\mathcal{Y}})$, the Brenier potential ϕ^k is $R_{\mathcal{Y}}$ -Lipschitz so that it is differentiable ρ -almost everywhere. Fenchel-Young equality ensures then that for ρ -a.e. $x \in \mathcal{X}$,

$$\phi^k(x) = \langle x | \nabla \phi^k(x) \rangle - \psi^k(\nabla \phi^k(x)).$$

The assumption that $\mu^k = (\nabla \phi^k)_{\#} \rho$ is supported on the set \mathcal{Y} then entails that for ρ -a.e. $x \in \mathcal{X}$, there exists $i \in \{1, \dots, N\}$ such that $\nabla \phi^k(x) = y_i$ so that

$$\phi^k(x) = \langle x | y_i \rangle - \psi^k(y_i).$$

Fenchel-Young inequality then ensures that for ρ -a.e. $x \in \mathcal{X}$,

$$\phi^k(x) = \max_{1 \leq i \leq N} \langle x | y_i \rangle - \psi^k(y_i).$$

Note that in this context, ϕ^k is piece-wise affine. To simplify the notation, we will conflate in this section and the following the function ψ^k with the vector $\boldsymbol{\psi}^k \in \mathbb{R}^N$ defined by $\boldsymbol{\psi}_i^k = \psi^k(y_i)$. This vector $\boldsymbol{\psi}^k$ defines a partition of the domain \mathcal{X} into so-called Laguerre cells, described for all $1 \leq i \leq N$ by

$$\text{Lag}_i(\boldsymbol{\psi}^k) = \{x \in \mathcal{X} \mid \forall j, \boldsymbol{\psi}_j^k \geq \boldsymbol{\psi}_i^k + \langle y_j - y_i | x \rangle\}.$$




Figure 2.1: (Left) The source ρ is the Lebesgue measure on $\mathcal{X} = [0, 1] \times [0, 1]$ and the target measure μ is the uniform discrete measure with $N = 4$ support points positioned at the blue crosses. (Middle) Voronoi tessellation of \mathcal{X} with respect to the support points of μ . (Right) Laguerre tessellation of \mathcal{X} with respect to the support points of μ for an optimal potential (the area of each cell is $\frac{1}{4}$).

This partition is such that

$$\mathcal{K}_\rho(\psi^k) = \sum_{i=1}^N \int_{\text{Lag}_i(\psi^k)} (\langle x | y_i \rangle - \psi_i^k) d\rho(x),$$

so that the Kantorovich functional evaluated in ψ^k only depends on the finite dimensional vector $\psi^k \in \mathbb{R}^N$ (as already observed more generally in Remark 1.7). The Kantorovich functional thus verifies $\mathcal{K}_\rho(\psi^k) = \mathcal{K}_\rho(\psi^k)$, where

$$\mathcal{K}_\rho : \psi \in \mathbb{R}^N \mapsto \sum_{i=1}^N \int_{\text{Lag}_i(\psi)} (\langle x | y_i \rangle - \psi_i) d\rho(x) \in \mathbb{R}. \quad (2.2)$$

Notice that the above defined Laguerre cells verify for all $1 \leq i \leq N$,

$$\text{Lag}_i(\psi^k) = \{x \in \mathcal{X} \mid \forall j, \|x - y_i\|^2 + 2\psi_i^k - \|y_i\|^2 \leq \|x - y_j\|^2 + 2\psi_j^k - \|y_j\|^2\}.$$

As such, one can think of the Laguerre cells as modified Voronoi cells, which are defined for each $i \in \{1, \dots, N\}$ by

$$\text{Vor}_i = \{x \in \mathcal{X} \mid \forall j, \|x - y_i\| \leq \|x - y_j\|\}.$$

The Voronoi tessellation of \mathcal{X} with respect to the finite set \mathcal{Y} partitions \mathcal{X} into cells consisting of all points of \mathcal{X} closer to a given point in \mathcal{Y} than to any other. The Laguerre tessellation modifies this partition by inflating or deflating the distance to each point of \mathcal{Y} . See Figure 2.1 for an illustration.

Gradient of \mathcal{K}_ρ . By Theorem 1.1 in (Kitagawa et al., 2019) (see also (Aurenhammer et al., 1998)), the Kantorovich functional (2.2) seen as a function on \mathbb{R}^N has the following gradient:

$$\forall \psi \in \mathbb{R}^N, \quad \nabla \mathcal{K}_\rho(\psi) = -(\rho(\text{Lag}_i(\psi)))_{1 \leq i \leq N} \in \mathbb{R}^N. \quad (2.3)$$

Note that this corresponds to the gradient formula of Lemma 1.11. Given a potential $\psi \in \mathbb{R}^N$, we introduce the corresponding probability measure:

$$\mu_\psi = \sum_{1 \leq i \leq N} \rho(\text{Lag}_i(\psi)) \delta_{y_i}.$$

With this notation, note that we recover the measures $\mu^k = \mu_{\psi^k} = (\nabla\phi^k)_{\#}\rho$ for $k \in \{0, 1\}$. In particular, the mass that ρ gives to $\text{Lag}_i(\psi^k)$ is equal to $\mu^k(y_i)$ (this corresponds to the first order condition $\nabla\mathcal{K}_\rho(\psi) + \boldsymbol{\mu}^k = 0$, where $\boldsymbol{\mu}^k = (\mu^k(y_i))_{1 \leq i \leq N}$).

Hessian of \mathcal{K}_ρ . It is also possible to compute the Hessian of \mathcal{K}_ρ . For this, we consider the set $S_+ \subseteq \mathbb{R}^N$ of potentials ψ such that all Laguerre cells $\text{Lag}_i(\psi)$ contain some mass, defined by

$$S_+ = \{\psi \in \mathbb{R}^N \mid \forall i, \rho(\text{Lag}_i(\psi)) > 0\}. \quad (2.4)$$

From Theorems 1.3 and 4.1 in (Kitagawa et al., 2019), we know that the map $\nabla\mathcal{K}_\rho$ is \mathcal{C}^1 on the set S_+ . By Theorem 1.3 in (Kitagawa et al., 2019), if $\psi \in S_+$, its derivatives are given by

$$\begin{cases} \frac{\partial^2 \mathcal{K}_\rho}{\partial \psi_i \partial \psi_j}(\psi) = - \int_{\text{Lag}_i(\psi) \cap \text{Lag}_j(\psi)} \frac{\rho(x)}{\|y_j - y_i\|} d\mathcal{H}^{d-1}(x) & \text{for } i \neq j, \\ \frac{\partial^2 \mathcal{K}_\rho}{\partial^2 \psi_i}(\psi) = - \sum_{j \neq i} \frac{\partial^2 \mathcal{K}_\rho}{\partial \psi_i \partial \psi_j}(\psi). \end{cases} \quad (2.5)$$

Notice from this definition that the Hessian matrix $\nabla^2 \mathcal{K}_\rho(\psi)$ corresponds to the Laplacian matrix of a weighted graph, where the vertices correspond to the support points $(y_i)_{1 \leq i \leq N}$ and the weight w_{ij} between two different vertices y_i and y_j corresponds to the (normalized) *size* of the intersection between the Laguerre cells $\text{Lag}_i(\psi)$ and $\text{Lag}_j(\psi)$ according to ρ :

$$w_{ij} = \frac{1}{\|y_j - y_i\|} \int_{\text{Lag}_i(\psi) \cap \text{Lag}_j(\psi)} \rho(x) d\mathcal{H}^{d-1}(x).$$

2.3 Proof of Theorem 2.1 in the semi-discrete case

We prove here Theorem 2.1, assuming that for $k \in \{0, 1\}$, the measures $\mu^k = (\nabla\phi^k)_{\#}\rho$ are supported on the fixed and common set $\mathcal{Y} = \{y_1, \dots, y_N\}$. We will require two preliminary results. The next lemma follows from the Brunn-Minkowski inequality (see Section 1.3 and in particular Theorem 1.27). This inequality has already appeared in the numerical analysis of Monge-Ampère equations, see (Benamou et al., 2016; Nchetto and Zhang, 2019).

Lemma 2.3. *Let $\psi^0, \psi^1 \in S^+$ and consider $\psi^t = (1-t)\psi^0 + t\psi^1$. Define the probability measure $\mu^t = \mu_{\psi^t} \in \mathcal{P}(\mathcal{Y})$. Then μ^t satisfies*

$$\mu^t \geq \min(1-t, t)^d \frac{m_\rho}{M_\rho} (\mu^0 + \mu^1).$$

In particular, $\psi^t \in S^+$.

Proof. Let $i \in \{1, \dots, N\}$ and $x^0 \in \text{Lag}_i(\psi^0)$ and $x^1 \in \text{Lag}_i(\psi^1)$. Then, for all $j \in \{1, \dots, N\}$,

$$\begin{cases} \psi_j^0 \geq \psi_i^0 + \langle y_j - y_i | x^0 \rangle, \\ \psi_j^1 \geq \psi_i^1 + \langle y_j - y_i | x^1 \rangle. \end{cases}$$

Taking the convex combination of these inequalities we get for all $j \in \{1, \dots, N\}$,

$$\psi_j^t \geq \psi_i^t + \langle y_j - y_i | (1-t)x^0 + tx^1 \rangle.$$

This shows that $(1-t)x^0 + tx^1 \in \text{Lag}_i(\boldsymbol{\psi}^t)$ (note that we use the convexity of \mathcal{X} here). Thus,

$$(1-t)\text{Lag}_i(\boldsymbol{\psi}^0) + t\text{Lag}_i(\boldsymbol{\psi}^1) \subseteq \text{Lag}_i(\boldsymbol{\psi}^t).$$

Taking the measure ρ on both sides and applying Brunn-Minkowski's inequality gives

$$\begin{aligned} (\mu^t(y_i))^{1/d} &= \rho(\text{Lag}_i(\boldsymbol{\psi}^t))^{1/d} \geq m_\rho^{1/d} |(1-t)\text{Lag}_i(\boldsymbol{\psi}^0) + t\text{Lag}_i(\boldsymbol{\psi}^1)|^{1/d} \\ &\geq m_\rho^{1/d} \left((1-t) |\text{Lag}_i(\boldsymbol{\psi}^0)|^{1/d} + t |\text{Lag}_i(\boldsymbol{\psi}^1)|^{1/d} \right) \\ &\geq \left(\frac{m_\rho}{M_\rho} \right)^{1/d} \left((1-t) (\mu^0(y_i))^{1/d} + t (\mu^1(y_i))^{1/d} \right). \end{aligned}$$

This finally yields:

$$\begin{aligned} \mu^t(y_i) &\geq \frac{m_\rho}{M_\rho} \left(\min(1-t, t) \left((\mu^0(y_i))^{1/d} + (\mu^1(y_i))^{1/d} \right) \right)^d \\ &\geq \min(1-t, t)^d \frac{m_\rho}{M_\rho} (\mu^0(y_i) + \mu^1(y_i)). \quad \square \end{aligned}$$

The next proposition gives an explicit lower bound on the smallest non-zero eigenvalue of the Hessian matrix

$$\nabla^2 \mathcal{K}_\rho(\boldsymbol{\psi}) = \left(\frac{\partial^2 \mathcal{K}_\rho}{\partial \psi_i \partial \psi_j}(\boldsymbol{\psi}) \right)_{1 \leq i, j \leq N} \in \mathbb{R}^{N \times N},$$

whose terms have been defined in (2.5). The proof of this proposition follows from the stability analysis of finite volumes discretization of elliptic PDEs, see Lemma 3.7 in (Eymard et al., 2000). We report this proof at the end of this section, with very minor adaptations to our case. The compactness assumption made on the target measures is used in this result (note the explicit dependence of the constant on the diameter of the target domain \mathcal{Y}).

Proposition 2.4 (Discrete Poincaré-Wirtinger inequality). *Consider $\boldsymbol{\psi} \in S_+$ and $v \in \mathbb{R}^N$. Then,*

$$\text{Var}_{\mu_\boldsymbol{\psi}}(v) \leq C_{d, \mathcal{X}, \mathcal{Y}, \rho} \langle \nabla^2 \mathcal{K}_\rho(\boldsymbol{\psi}) v | v \rangle$$

where $C_{d, \mathcal{X}, \mathcal{Y}, \rho} = \omega_d \text{diam}(\mathcal{Y}) \text{diam}(\mathcal{X})^{d+1} \frac{M_\rho^2}{m_\rho}$, with ω_d the volume of the d -dimensional unit ball.

Remark 2.5. As expected, $\nabla^2 \mathcal{K}_\rho(\boldsymbol{\psi})$ is positive semidefinite, since its smallest non-zero eigenvalue is greater than a variance. This can also be seen from the definition of $\nabla^2 \mathcal{K}_\rho(\boldsymbol{\psi})$ as a Laplacian matrix, or simply from Gershgorin's circle theorem and the explicit formula for $\nabla^2 \mathcal{K}_\rho(\boldsymbol{\psi})$ recalled in (2.5).

With these two results at hand, we can now show the strong convexity estimate of Theorem 2.1 in the semi-discrete case, that we phrase as the following proposition:

Proposition 2.6. *With the notation and assumption of Theorem 2.1, assume additionally that the set \mathcal{Y} is finite and that μ^0, μ^1 give mass to all points in \mathcal{Y} . Then, (2.1) holds.*

Proof. We denote $\mathcal{Y} = \{y_1, \dots, y_N\}$. With the notation of Section 2.2, the assumptions entail that for $k \in \{0, 1\}$, the vector $\boldsymbol{\psi}^k = (\psi^k(y_i))_{1 \leq i \leq N}$ is in S_+ . Again with the notation of this section, one has:

$$\mathcal{K}_\rho(\boldsymbol{\psi}^1) - \mathcal{K}_\rho(\boldsymbol{\psi}^0) + \langle \mu_0 | \boldsymbol{\psi}^1 - \boldsymbol{\psi}^0 \rangle = \mathcal{K}_\rho(\boldsymbol{\psi}^1) - \mathcal{K}_\rho(\boldsymbol{\psi}^0) - \langle \nabla \mathcal{K}_\rho(\boldsymbol{\psi}^0) | \boldsymbol{\psi}^1 - \boldsymbol{\psi}^0 \rangle.$$

Denote $v = \boldsymbol{\psi}^1 - \boldsymbol{\psi}^0 \in \mathbb{R}^N$ and $\boldsymbol{\psi}^t = \boldsymbol{\psi}^0 + tv \in \mathbb{R}^N$ for $t \in [0, 1]$. Using that for all $t \in [0, 1]$, $\boldsymbol{\psi}^t \in S_+$ by Lemma 2.3 and that \mathcal{K}_ρ is \mathcal{C}^2 on S_+ (Section 2.2), one has from the fundamental theorem of calculus:

$$\mathcal{K}_\rho(\boldsymbol{\psi}^1) - \mathcal{K}_\rho(\boldsymbol{\psi}^0) - \langle \nabla \mathcal{K}_\rho(\boldsymbol{\psi}^0) | \boldsymbol{\psi}^1 - \boldsymbol{\psi}^0 \rangle = \int_0^1 \int_0^s \langle \nabla^2 \mathcal{K}_\rho(\boldsymbol{\psi}^t) v | v \rangle dt ds.$$

Using again that for any $t \in [0, 1]$, $\boldsymbol{\psi}^t \in S_+$, Proposition 2.4 ensures

$$\mathcal{K}_\rho(\boldsymbol{\psi}^1) - \mathcal{K}_\rho(\boldsymbol{\psi}^0) - \langle \nabla \mathcal{K}_\rho(\boldsymbol{\psi}^0) | \boldsymbol{\psi}^1 - \boldsymbol{\psi}^0 \rangle \geq \frac{1}{C_{d, \mathcal{X}, \mathcal{Y}, \rho}} \int_0^1 \int_0^s \mathbb{V}\text{ar}_{\mu_{\boldsymbol{\psi}^t}}(v) dt ds,$$

where $C_{d, \mathcal{X}, \mathcal{Y}, \rho} = \omega_d \text{diam}(\mathcal{Y}) \text{diam}(\mathcal{X})^{d+1} \frac{M_\rho^2}{m_\rho}$. From Lemma 2.3, we have for any $t \in [0, 1]$ the comparison

$$\mathbb{V}\text{ar}_{\mu_{\boldsymbol{\psi}^t}}(v) \geq 2 \min(1-t, t)^d \frac{m_\rho}{M_\rho} \mathbb{V}\text{ar}_{\frac{1}{2}(\mu^0 + \mu^1)}(v),$$

so that

$$\int_0^1 \int_0^s \mathbb{V}\text{ar}_{\mu_{\boldsymbol{\psi}^t}}(v) dt ds \geq \frac{1}{(d+1)2^{d-1}} \frac{m_\rho}{M_\rho} \mathbb{V}\text{ar}_{\frac{1}{2}(\mu^0 + \mu^1)}(v).$$

This finally ensures the bound:

$$\mathcal{K}_\rho(\boldsymbol{\psi}^1) - \mathcal{K}_\rho(\boldsymbol{\psi}^0) + \langle \mu_0 | \boldsymbol{\psi}^1 - \boldsymbol{\psi}^0 \rangle \geq \frac{1}{C_{d, \mathcal{X}, \mathcal{Y}}} \frac{m_\rho^2}{M_\rho^3} \mathbb{V}\text{ar}_{\frac{1}{2}(\mu^0 + \mu^1)}(\boldsymbol{\psi}^1 - \boldsymbol{\psi}^0),$$

where

$$C_{d, \mathcal{X}, \mathcal{Y}} = \omega_d (d+1) 2^{d-1} \text{diam}(\mathcal{Y}) \text{diam}(\mathcal{X})^{d+1}. \quad \square$$

We finally proceed with the proof of Proposition 2.4.

Proof of Proposition 2.4. This proof is a straightforward adaptation of a stability result for finite volume discretization of elliptic PDEs, see Lemma 3.7 in (Eymard et al., 2000). We consider the function u on \mathcal{X} defined a.e. by $u|_{\text{Lag}_i(\boldsymbol{\psi})} = v_i$. Then,

$$\begin{aligned} \mathbb{V}\text{ar}_{\mu_\boldsymbol{\psi}}(v) &= \sum_i v_i^2 \mu_\boldsymbol{\psi}(y_i) - \left(\sum_i v_i \mu_\boldsymbol{\psi}(y_i) \right)^2 \\ &= \int_{\mathcal{X}} u^2 d\rho - \left(\int_{\mathcal{X}} u d\rho \right)^2 \\ &= \frac{1}{2} \int_{\mathcal{X} \times \mathcal{X}} (u(x) - u(y))^2 d\rho(x) d\rho(y), \end{aligned}$$

so it suffices to control the right hand side of this equality. Given $(i, j) \in \{1, \dots, N\}$ and $(x, y) \in \mathcal{X}$, we denote

$$\chi_{ij}(x, y) = \begin{cases} 1 & \text{if } \text{Lag}_i(\boldsymbol{\psi}) \cap \text{Lag}_j(\boldsymbol{\psi}) \cap [x, y] \neq \emptyset \text{ and } \langle y_j - y_i | y - x \rangle \geq 0, \\ 0 & \text{if not.} \end{cases}$$

Then, $u(y) - u(x) = \sum_{i \neq j} (v_j - v_i) \chi_{ij}(x, y)$. We introduce

$$d_{ij} = \|y_j - y_i\|, \quad c_{ij,z} = \left| \left\langle \frac{z}{\|z\|} \middle| \frac{y_j - y_i}{\|y_j - y_i\|} \right\rangle \right|,$$

and we apply Cauchy-Schwarz's inequality to get

$$\begin{aligned} (u(y) - u(x))^2 &= \left(\sum_{i \neq j} (v_j - v_i) \chi_{ij}(x, y) \right)^2 \\ &\leq \sum_{i \neq j} \frac{(v_j - v_i)^2}{d_{ij} c_{ij,y-x}} \chi_{ij}(x, y) \sum_{i \neq j} d_{ij} c_{ij,y-x} \chi_{ij}(x, y). \end{aligned}$$

In addition, when $\chi_{ij}(x, y) = 1$, we have $\langle y - x | y_j - y_i \rangle \geq 0$ so that

$$d_{ij} c_{ij,y-x} = \|y_j - y_i\| \left\langle \frac{y - x}{\|y - x\|} \middle| \frac{y_j - y_i}{\|y_j - y_i\|} \right\rangle \geq 0,$$

and

$$\sum_{i \neq j} d_{ij} c_{ij,y-x} \chi_{ij}(x, y) = \sum_{i \neq j} \left\langle \frac{y - x}{\|y - x\|} \middle| y_j - y_i \right\rangle \chi_{ij}(x, y) \leq \text{diam}(\mathcal{Y}).$$

Therefore,

$$\begin{aligned} &\int_{\mathcal{X} \times \mathcal{X}} (u(y) - u(x))^2 d\rho(x) d\rho(y) \\ &\leq \text{diam}(\mathcal{Y}) \int_{\mathcal{X} \times \mathcal{X}} \sum_{i \neq j} \frac{(v_j - v_i)^2}{d_{ij} c_{ij,y-x}} \chi_{ij}(x, y) d\rho(x) d\rho(y) \\ &\leq M_\rho \text{diam}(\mathcal{Y}) \int_{B(0, \text{diam}(\mathcal{X}))} \sum_{i \neq j} \frac{(v_j - v_i)^2}{d_{ij} c_{ij,z}} \left(\int_{\mathcal{X}} \chi_{ij}(x, x+z) d\rho(x) \right) dz. \end{aligned}$$

Moreover, denoting $m_{ij} = \int_{\text{Lag}_i(\psi) \cap \text{Lag}_j(\psi)} \rho(x) d\mathcal{H}^{d-1}(x)$, we get for any $z \in B(0, \text{diam}(\mathcal{X}))$,

$$\begin{aligned} \int_{\mathcal{X}} \chi_{ij}(x, x+z) d\rho(x) &\leq M_\rho \int_{\mathcal{X}} \chi_{ij}(x, x+z) dx \\ &\leq M_\rho \text{vol}^{d-1}(\text{Lag}_i(\psi) \cap \text{Lag}_j(\psi)) \|z\| c_{ij,z} \\ &\leq \frac{M_\rho}{m_\rho} m_{ij} \|z\| c_{ij,z}. \end{aligned}$$

This gives

$$\frac{1}{2} \int_{\mathcal{X} \times \mathcal{X}} (u(y) - u(x))^2 d\rho(x) d\rho(y) \leq \omega_d \text{diam}(\mathcal{Y}) \text{diam}(\mathcal{X})^{d+1} \frac{M_\rho^2}{m_\rho} \frac{1}{2} \sum_{i \neq j} \frac{m_{ij}}{d_{ij}} (v_j - v_i)^2,$$

where ω_d denotes the volume of the d -dimensional unit ball. Define $H \in \mathbb{R}^{N \times N}$ with

coefficients $H_{ij} = -\frac{m_{ij}}{d_{ij}}$, $H_{ii} = -\sum_{j \neq i} H_{ij}$. Then, $\nabla^2 \mathcal{K}_\rho(\psi) = H$, and notice that:

$$\begin{aligned} \frac{1}{2} \sum_{i \neq j} \frac{m_{ij}}{d_{ij}} (v_j - v_i)^2 &= \frac{1}{2} \sum_{i \neq j} -H_{ij} (v_j - v_i)^2 \\ &= \sum_{i \neq j} -H_{ij} v_i (v_i - v_j) \\ &= \sum_i \left(H_{ii} v_i v_i + \sum_{j \neq i} H_{ij} v_i v_j \right) \\ &= \sum_{i,j} H_{ij} v_i v_j \\ &= \langle \nabla \mathcal{K}_\rho(\psi) v | v \rangle. \end{aligned}$$

We finally obtain

$$\text{Var}_{\mu_\psi}(v) \leq \omega_d \text{diam}(\mathcal{Y}) \text{diam}(\mathcal{X})^{d+1} \frac{M_\rho^2}{m_\rho} \langle \nabla \mathcal{K}_\rho(\psi) v | v \rangle. \quad \square$$

2.4 From the semi-discrete case to the general case

We now turn to the proof of Theorem 2.1 in the general case. This results from a simple density argument, which is summarized in the following lemma.

Lemma 2.7 (Semi-discrete approximation). *Under the notation and assumptions of Theorem 2.1 and for $k \in \{0, 1\}$, there exists sequences of Brenier potentials $(\phi_N^k)_{N \geq 1} \in \mathcal{C}^0(\mathcal{X})$ with associated convex conjugates $(\psi_N^k)_{N \geq 1} \in \mathcal{C}^0(\mathcal{Y})$ such that:*

- (i) *The measures $\mu_N^0 = (\nabla \phi_N^0)_\# \rho$ and $\mu_N^1 = (\nabla \phi_N^1)_\# \rho$ have the same finite support,*
- (ii) $\lim_{N \rightarrow +\infty} \mathcal{K}_\rho(\psi_N^1) - \mathcal{K}_\rho(\psi_N^0) = \mathcal{K}_\rho(\psi^1) - \mathcal{K}_\rho(\psi^0),$
- (iii) $\lim_{N \rightarrow +\infty} \langle \psi_N^1 - \psi_N^0 | \mu_N^0 \rangle = \langle \psi^1 - \psi^0 | \mu^0 \rangle,$
- (iv) $\lim_{N \rightarrow +\infty} \text{Var}_{\frac{1}{2}(\mu_N^0 + \mu_N^1)}(\psi_N^1 - \psi_N^0) = \text{Var}_{\frac{1}{2}(\mu^0 + \mu^1)}(\psi^1 - \psi^0).$

Proof. For any $N > 0$, we consider a finite partition $\mathcal{Y} = \sqcup_{1 \leq i \leq N} \mathcal{Y}_i^N$, and we let $\varepsilon_N = \max_i \text{diam}(\mathcal{Y}_i^N)$. We assume that $\lim_{N \rightarrow +\infty} \varepsilon_N = 0$. Then, we define

$$\mu_N^k = \sum_{1 \leq i \leq N} \left[\left(1 - \frac{1}{N} \right) \mu^k(\mathcal{Y}_i^N) + \frac{1}{N^2} \right] \delta_{y_i^N},$$

where $y_i^N \in \mathcal{Y}_i^N$. It is easy to check that the support of the measures μ_N^0 and μ_N^1 is the set $\{y_1^N, \dots, y_N^N\}$. For $k \in \{0, 1\}$ and $N \geq 1$, denote $\phi_N^k \in \mathcal{C}^0(\mathcal{X})$ a Brenier potential for the quadratic optimal transport between ρ and μ_N^k chosen such that $\langle \phi_N^k | \rho \rangle = \langle \phi^k | \rho \rangle$: ϕ_N^k is a convex function that satisfies $\mu_N^k = (\nabla \phi_N^k)_\# \rho$ (see Theorem 1.12). Then ϕ_N^k and its convex conjugate ψ_N^k are Kantorovich potentials for the optimal transport problem between ρ and μ_N^k . Introduce

$$\tilde{\mu}_N^k = \sum_{1 \leq i \leq N} \mu^k(\mathcal{Y}_i^N) \delta_{y_i^N},$$

and notice that

$$\begin{aligned} W_1(\mu_N^k, \mu^k) &\leq W_1(\mu_N^k, \tilde{\mu}_N^k) + W_1(\tilde{\mu}_N^k, \mu^k) \\ &\leq \text{diam}(\mathcal{Y}) \left\| \mu_N^k - \tilde{\mu}_N^k \right\|_{\text{TV}} + \varepsilon_N \\ &\leq \frac{2\text{diam}(\mathcal{Y})}{N} + \varepsilon_N, \end{aligned}$$

so that $\lim_{N \rightarrow +\infty} W_1(\mu_N^k, \mu^k) = 0$ (here, $\|\cdot\|_{\text{TV}}$ denotes the total variation distance and we used the comparison $W_1 \leq \text{diam}(\mathcal{Y}) \|\cdot\|_{\text{TV}}$ valid for probability measures supported on \mathcal{Y}). Standard stability results of optimal transport (Theorem A.10) then ensure that ϕ_N^k, ψ_N^k converge uniformly to Kantorovich potentials $\tilde{\phi}^k, \tilde{\psi}^k$ between ρ and μ^k on \mathcal{X} and \mathcal{Y} respectively. However, in our setting where ρ is supported on the whole convex set \mathcal{X} , the Kantorovich potentials between ρ and μ^k are unique up to additive constants (Proposition 7.18 of (Santambrogio, 2015), see Remark 1.13). Thus the condition

$$\langle \phi^k | \rho \rangle = \lim_{N \rightarrow +\infty} \langle \phi_N^k | \rho \rangle = \langle \tilde{\phi}^k | \rho \rangle,$$

ensures that $\tilde{\phi}^k = \phi^k$, which entails in turn $\tilde{\psi}^k = \psi^k$. Using that for $k \in \{0, 1\}$, ϕ_N^k converges uniformly to ϕ^k on \mathcal{X} then ensures together with the dominated convergence theorem the limit $\int (\phi_N^1 - \phi_N^0) d\rho \rightarrow \int (\phi^1 - \phi^0) d\rho$ so that $\mathcal{K}_\rho(\psi_N^1) - \mathcal{K}_\rho(\psi_N^0) \rightarrow \mathcal{K}_\rho(\psi^1) - \mathcal{K}_\rho(\psi^0)$. Then notice that

$$\begin{aligned} |\langle \psi_N^1 - \psi_N^0 | \mu_N^0 \rangle - \langle \psi^1 - \psi^0 | \mu^0 \rangle| &\leq |\langle \psi_N^1 - \psi_N^0 | \mu_N^0 - \mu^0 \rangle| \\ &\quad + |\langle (\psi_N^1 - \psi_N^0) - (\psi^1 - \psi^0) | \mu^0 \rangle|. \end{aligned}$$

As Legendre transform of $\phi_N^k \in \mathcal{C}^0(\mathcal{X})$, ψ_N^k is $R_{\mathcal{X}}$ -Lipschitz where $R_{\mathcal{X}}$ is such that $\mathcal{X} \subset B(0, R_{\mathcal{X}})$ (see Remarks 1.6 and 1.7). Therefore by the Kantorovich-Rubinstein duality result (Proposition A.8),

$$|\langle \psi_N^1 - \psi_N^0 | \mu_N^0 - \mu^0 \rangle| \leq 2R_{\mathcal{X}} W_1(\mu_N^0, \mu^0) \rightarrow 0.$$

One has the limit $|\langle (\psi_N^1 - \psi_N^0) - (\psi^1 - \psi^0) | \mu^0 \rangle| \rightarrow 0$ from the uniform convergence of ψ_N^k to ψ^k on \mathcal{Y} for $k \in \{0, 1\}$ and the dominated convergence theorem. The exact same arguments allow to show the limit $\langle \psi_N^1 - \psi_N^0 | \frac{1}{2}(\mu_N^0 + \mu_N^1) \rangle \rightarrow \langle \psi^1 - \psi^0 | \frac{1}{2}(\mu^0 + \mu^1) \rangle$. Finally, one shows very similarly the limit $\langle (\psi_N^1 - \psi_N^0)^2 | \frac{1}{2}(\mu_N^0 + \mu_N^1) \rangle \rightarrow \langle (\psi^1 - \psi^0)^2 | \frac{1}{2}(\mu^0 + \mu^1) \rangle$, using that $(\psi_N^1 - \psi_N^0)^2$ is Lipschitz continuous since $\psi_N^1 - \psi_N^0$ is $2R_{\mathcal{X}}$ -Lipschitz continuous and uniformly bounded (since it converges uniformly to the function $\psi^1 - \psi^0$, which is Lipschitz on the compact set \mathcal{Y} and thus bounded). \square

Proof of Theorem 2.1. Let $(\phi_N^k)_{N \geq 1} \in \mathcal{C}^0(\mathcal{X})$ be the sequence of Brenier potentials from Lemma 2.7 that are such that the measures $\mu_N^0 = (\nabla \phi_N^0)_{\#} \rho$ and $\mu_N^1 = (\nabla \phi_N^1)_{\#} \rho$ have the same finite support. Then for all $N \geq 1$, Proposition 2.6 ensures the inequality

$$\frac{1}{C_{d, \mathcal{X}, \mathcal{Y}}} \frac{m_\rho^2}{M_\rho^3} \text{Var}_{\frac{1}{2}(\mu_N^0 + \mu_N^1)}(\psi_N^1 - \psi_N^0) \leq \mathcal{K}_\rho(\psi_N^1) - \mathcal{K}_\rho(\psi_N^0) + \langle (\nabla \phi_N^0)_{\#} \rho | \psi_N^1 - \psi_N^0 \rangle.$$

By Lemma 2.7, taking N to $+\infty$ in the preceding inequality establishes (2.1) in the limit. \square

A continuous approach

Abstract

This chapter details the proof of a second strong convexity estimate for the Kantorovich functional originally derived in (Delalande and Mérigot, 2021). This result holds when the Kantorovich functional is evaluated on Kantorovich potentials whose convex conjugates present a bounded oscillation. The proof of this strong convexity estimate is first derived under strong regularity assumptions on the evaluated potentials, and it is generalized by a density argument. The main elements of proof are deduced from the Brascamp-Lieb inequality and the log-concavity of the determinant on the set of non-negative symmetric matrices.

3.1 Introduction

Once again, we prove in this chapter a strong convexity estimate for the Kantorovich functional that appears in the dual quadratic optimal transport problem (see Sections 1.1.3 and 1.2). We recall that the Kantorovich functional associated to a source measure $\rho \in \mathcal{P}_2(\mathbb{R}^d)$ reads

$$\mathcal{K}_\rho : \begin{cases} (\mathbb{R}^d \rightarrow \bar{\mathbb{R}}) & \rightarrow \bar{\mathbb{R}}, \\ \psi & \mapsto \int_{\mathbb{R}^d} \psi^* d\rho. \end{cases}$$

We prove the following result, originally established in (Delalande and Mérigot, 2021):

Theorem 3.1. *Let ρ be a probability density over a compact convex set \mathcal{X} , satisfying $0 < m_\rho \leq \rho \leq M_\rho$. Let $\mu^0, \mu^1 \in \mathcal{P}_2(\mathbb{R}^d)$. For $k \in \{0, 1\}$, denote ϕ^k a Brenier potential between ρ and μ^k . Assume that*

$$\forall k \in \{0, 1\}, \quad -\infty < m_\phi \leq \min_{\mathcal{X}} \phi^k \leq \max_{\mathcal{X}} \phi^k \leq M_\phi < +\infty.$$

Then the convex conjugates ψ^0 and ψ^1 of ϕ^0 and ϕ^1 verify

$$\frac{1}{C_d(M_\phi - m_\phi)} \frac{m_\rho^2}{M_\rho^2} \text{Var}_{\frac{1}{2}(\mu^0 + \mu^1)}(\psi^1 - \psi^0) \leq \mathcal{K}_\rho(\psi^1) - \mathcal{K}_\rho(\psi^0) + \langle \psi^1 - \psi^0 | \mu^0 \rangle, \quad (3.1)$$

where $C_d = e(d+1)2^{d-1}$.

The main assumption in this result is that the probability measures μ^0, μ^1 are such that their Brenier potentials for the transport between ρ and themselves have a bounded oscillation on the compact set \mathcal{X} . As noted in Remark 1.26, by Morrey's inequality, such assumption is satisfied whenever the measures μ^0, μ^1 admit a finite p -th moment for some $p > d$.

Outline. In order to prove Theorem 3.1, we take the following steps: first, we assume that the probability measures μ^0, μ^1 are such that the Brenier potentials ϕ^0, ϕ^1 are smooth and strongly convex – which requires at least that μ^0 and μ^1 are absolutely continuous and compactly supported. In this regular context, we compute the first and second derivatives of \mathcal{K}_ρ along the path $((1-t)\psi^0 + t\psi^1)_{t \in [0,1]}$ (§3.2). We then show that the Brascamp-Lieb inequality gives an explicit lower-bound on the second derivative of \mathcal{K}_ρ , from which we deduce Theorem 3.1 in the regular context (§3.3). Finally, we prove that Theorem 3.1 holds in the general case by density arguments (§3.4).

3.2 Derivatives of the Kantorovich functional for regular potentials

The strong convexity estimate (3.1) will be derived from a local estimate – a Poincaré-Wirtinger type inequality for the second derivative of \mathcal{K}_ρ – which will be shown to be a consequence of the Brascamp-Lieb inequality in the following section. To make the connection with the Brascamp-Lieb inequality clearer, we first compute the first and second order derivatives of \mathcal{K}_ρ along the path $((1-t)\psi^0 + t\psi^1)_{t \in [0,1]}$, under regularity and strong convexity hypotheses.

Proposition 3.2. *Let $\phi^0, \phi^1 \in \mathcal{C}^2(\mathbb{R}^d)$ be strongly convex functions. Define $\psi^0 = (\phi^0)^*$, $\psi^1 = (\phi^1)^*$ and $v = \psi^1 - \psi^0$. For $t \in [0, 1]$, define $\psi^t = \psi^0 + tv$ and finally $\phi^t = (\psi^t)^*$. Then, ϕ^t is a strongly convex function, belongs to $\mathcal{C}^2(\mathbb{R}^d)$, and*

$$\frac{d}{dt} \mathcal{K}_\rho(\psi^t) = - \int_{\mathcal{X}} v(\nabla \phi^t(x)) d\rho(x), \quad (3.2)$$

$$\frac{d^2}{dt^2} \mathcal{K}_\rho(\psi^t) = \int_{\mathcal{X}} \langle \nabla v(\nabla \phi^t(x)) | D^2 \phi^t(x) \cdot \nabla v(\nabla \phi^t(x)) \rangle d\rho(x). \quad (3.3)$$

Proof of Proposition 3.2. We assume that ϕ^0, ϕ^1 are both α -strongly convex and belong to $\mathcal{C}^2(\mathbb{R}^d)$. Then, the convex conjugates $\psi^0 = (\phi^0)^*$, $\psi^1 = (\phi^1)^*$ are \mathcal{C}^2 with $1/\alpha$ -Lipschitz gradients and satisfy $D^2\psi^0 > 0, D^2\psi^1 > 0$ everywhere on \mathbb{R}^d . Hence their linear interpolates $\psi^t = (1-t)\psi^0 + t\psi^1$ enjoy the same properties. This in turn implies that for all $t \in [0, 1]$, the convex conjugate ϕ^t of ψ^t belongs to $\mathcal{C}^2(\mathbb{R}^d)$ and is α -strongly convex.

We will now prove that the map $G : (t, x) \mapsto \nabla \phi^t(x)$ has class \mathcal{C}^1 . Let $F : [0, 1] \times \mathbb{R}^d \times \mathbb{R}^d \rightarrow \mathbb{R}^d$ be the continuously differentiable function defined by $F(t, x, y) = \nabla \psi^t(y) - x$. A well-known property of the convex conjugate is that $\nabla \phi^t$ is the inverse of $\nabla \psi^t$, implying that $G(t, x)$ is uniquely characterized by $F(t, x, G(t, x)) = 0$. Since $D^2\psi^t > 0$, the Jacobian $D_y F(t, x, y) = D^2\psi^t(y)$ is invertible and the implicit function theorem thus implies that G has class \mathcal{C}^1 . Differentiating the relation $F(t, x, G(t, x)) = 0$ with respect

to time, we get

$$\frac{d}{dt}\nabla\phi^t(x) = -D^2\phi^t(x) \cdot \nabla v(\nabla\phi^t(x)). \quad (3.4)$$

By Fenchel-Young's equality case, one has for any $x \in \mathcal{X}$ and $t \in [0, 1]$,

$$\phi^t(x) = \langle x | \nabla\phi^t(x) \rangle - \psi^t(\nabla\phi^t(x)),$$

so that ϕ^t is at least \mathcal{C}^1 with respect to time. We can actually differentiate this equation with respect to time twice and using (3.4) we get

$$\begin{aligned} \frac{d}{dt}\phi^t(x) &= \langle x | \frac{d}{dt}\nabla\phi^t(x) \rangle - v(\nabla\phi^t(x)) - \langle \nabla\psi^t(\nabla\phi^t) | \frac{d}{dt}\nabla\phi^t(x) \rangle = -v(\nabla\phi^t(x)), \\ \frac{d^2}{dt^2}\phi^t(x) &= -\langle \nabla v(\nabla\phi^t(x)) | \frac{d}{dt}\nabla\phi^t(x) \rangle = \langle \nabla v(\nabla\phi^t(x)) | D^2\phi^t(x) \cdot \nabla v(\nabla\phi^t(x)) \rangle. \end{aligned}$$

Since $\mathcal{K}_\rho(\psi^t) = \int_{\mathcal{X}} \phi^t(x) d\rho(x)$, we get the result by differentiating twice under the integral. \square

3.3 Proof of Theorem 3.1 in the regular case

In this section we find a positive lower-bound on the second-order derivative of \mathcal{K}_ρ expressed in equation (3.3) using the Brascamp-Lieb inequality ((Brascamp and Lieb, 1976), see Theorem 1.29). This allows to prove Theorem 3.1 under additional regularity assumptions on the Brenier potentials ϕ^0, ϕ^1 , a result that we summarize in the following statement.

Proposition 3.3. *In addition to the assumptions of Theorem 3.1, assume that the functions ϕ^0, ϕ^1 are strongly convex, belong to $\mathcal{C}^2(\mathbb{R}^d)$, and that $\nabla\phi^0$ and $\nabla\phi^1$ induce diffeomorphisms between \mathcal{X} and a closed ball \mathcal{Y} . Then, inequality (3.1) holds.*

Proof. Under the assumptions on ϕ^0, ϕ^1 , Proposition 3.2 ensures that the function ϕ^t it defines is strongly convex and belongs to $\mathcal{C}^2(\mathbb{R}^d)$ for any $t \in [0, 1]$. Proposition 3.2 ensures that

$$\left. \frac{d}{dt}\mathcal{K}_\rho(\psi^t) \right|_{t=0} = -\langle \psi^1 - \psi^0 | \mu^0 \rangle.$$

By the fundamental theorem of calculus, again with the notation of Proposition 3.2, we thus have:

$$\mathcal{K}_\rho(\psi^1) - \mathcal{K}_\rho(\psi^0) + \langle \psi^1 - \psi^0 | \mu^0 \rangle = \int_0^1 \int_0^s \frac{d^2}{dt^2}\mathcal{K}_\rho(\psi^t) dt ds. \quad (3.5)$$

From Proposition 3.2, we have the following expression for the second derivative of \mathcal{K}_ρ :

$$\frac{d^2}{dt^2}\mathcal{K}_\rho(\psi^t) = \mathbb{E}_\rho \langle \nabla v(\nabla\phi^t) | (D^2\phi^t) \cdot \nabla v(\nabla\phi^t) \rangle.$$

We introduce $\tilde{v}^t = v(\nabla\phi^t)$ for any $t \in [0, 1]$, which belongs to $\mathcal{C}^1(\mathbb{R}^d)$ as the composition of $v = \psi^1 - \psi^0 \in \mathcal{C}^2(\mathbb{R}^d)$ and $\nabla\phi^t \in \mathcal{C}^1(\mathbb{R}^d)$. We have $\nabla\tilde{v}^t = D^2\phi^t \cdot \nabla v(\nabla\phi^t)$, where $(D^2\phi^t)$ is invertible by strong convexity. Thus,

$$\frac{d^2}{dt^2}\mathcal{K}_\rho(\psi^t) = \mathbb{E}_\rho \langle \nabla\tilde{v}^t | (D^2\phi^t)^{-1} \cdot \nabla\tilde{v}^t \rangle. \quad (3.6)$$

We now introduce $\tilde{\rho}^t = \exp(-\phi^t)/Z_t$ where $Z_t = \int_{\mathcal{X}} \exp(-\phi^t(x))dx$, which is the density of a log-concave probability measure supported on \mathcal{X} . The Brascamp-Lieb inequality, recalled in Theorem 1.29, then ensures that

$$\mathrm{Var}_{\tilde{\rho}^t}(\tilde{v}^t) \leq \mathbb{E}_{\tilde{\rho}^t} \langle \nabla \tilde{v}^t | (D^2 \phi^t)^{-1} \cdot \nabla \tilde{v}^t \rangle. \quad (3.7)$$

We assumed that for any $k \in \{0, 1\}$ and $x \in \mathcal{X}$, $m_\phi \leq \phi^k(x) \leq M_\phi$. We claim that this property is transferred to ϕ^t for any $t \in [0, 1]$. Indeed, on the one hand for all $t \in [0, 1]$,

$$\phi^t = ((1-t)\psi^0 + t\psi^1)^* \leq (1-t)(\psi^0)^* + t(\psi^1)^* = (1-t)\phi^0 + t\phi^1 \leq M_\phi,$$

where we used the convexity of the convex conjugation. On the other hand, for any $x \in \mathcal{X}$, we have by definition:

$$\phi^t(x) = \sup_{y \in \mathbb{R}^d} \langle x | y \rangle - \psi^t(y) \geq -\psi^t(0) = -(1-t)\psi^0(0) - t\psi^1(0).$$

By definition, $\psi^k = (\phi^k)^*$. Thus, for $k \in \{0, 1\}$, $\psi^k(0) = \sup_{x \in \mathcal{X}} -\phi^k(x) \leq -m_\phi$, ensuring that $\phi^t \geq m_\phi$ for all $t \in [0, 1]$. The inequality $m_\phi \leq \phi^t \leq M_\phi$ allows us to compare the densities ρ and $\tilde{\rho}^t$:

$$\left(\frac{\exp(-M_\phi)}{M_\rho Z_t} \right) \rho \leq \tilde{\rho}^t \leq \left(\frac{\exp(-m_\phi)}{m_\rho Z_t} \right) \rho.$$

This comparison and equation (3.7) thus give:

$$\left(\frac{\exp(-M_\phi)}{M_\rho Z_t} \right) \mathrm{Var}_\rho(\tilde{v}^t) \leq \left(\frac{\exp(-m_\phi)}{m_\rho Z_t} \right) \mathbb{E}_\rho \langle \nabla \tilde{v}^t | (D^2 \phi^t)^{-1} \cdot \nabla \tilde{v}^t \rangle,$$

where we used that for any absolutely continuous $\rho_1, \rho_2 \in \mathcal{P}_2(\mathbb{R}^d)$ and $f : \mathbb{R}^d \rightarrow \mathbb{R}$, the density comparison $\rho_1 \leq C\rho_2$ for some $C > 0$ yields

$$\mathrm{Var}_{\rho_1}(f) = \min_{c \in \mathbb{R}} \|f - c\|_{L^2(\rho_1)}^2 \leq C \min_{c \in \mathbb{R}} \|f - c\|_{L^2(\rho_2)}^2 = C \mathrm{Var}_{\rho_2}(f).$$

Therefore, using $\tilde{v}^t = v(\nabla \phi^t)$, $\mu^t = (\nabla \phi^t)_{\#} \rho$, $v = \psi^1 - \psi^0$ and expression (3.6):

$$\mathrm{Var}_{\mu^t}(\psi^1 - \psi^0) \leq \frac{M_\rho}{m_\rho} \exp(M_\phi - m_\phi) \frac{d^2}{dt^2} \mathcal{K}_\rho(\psi^t). \quad (3.8)$$

Note that $\mu^t = \nabla((1-t)\psi^0 + t\psi^1)_{\#} \rho$ interpolates between μ^0 in $t = 0$ and μ^1 in $t = 1$, but this interpolation is neither a displacement interpolation in the sense of McCann (McCann, 1997) nor a generalized geodesic in the sense of Ambrosio, Gigli, Savaré (Ambrosio et al., 2008). Recalling equation (3.5), this equation is similar to that of (3.1), except that we would like to replace μ^t by $\frac{1}{2}(\mu_0 + \mu_1)$. For this purpose, we will prove that

$$\mu^t \geq \frac{m_\rho}{M_\rho} \min(t, 1-t)^d (\mu^0 + \mu^1). \quad (3.9)$$

This will be done using an explicit expression for μ^t . By smoothness and strong convexity of the function ϕ^t , the restriction of $\nabla \phi^t$ to \mathcal{X} is a diffeomorphism on its image. This implies that μ^t is absolutely continuous with respect to the Lebesgue measure. Moreover, by e.g. (Villani, 2003, p.9), for any $x \in \mathcal{X}$ the density of μ^t with respect to Lebesgue,

also denoted μ^t , is given by $\mu^t(\nabla\phi^t(x)) \det(D^2\phi^t(x)) = \rho(x)$. Setting $y = \nabla\phi^t(x)$ in this formula, we get

$$\forall y \in \nabla\phi^t(\mathcal{X}), \quad \mu^t(y) = \rho(\nabla\psi^t(y)) \det(D^2\psi^t(y)).$$

By assumption, for $k \in \{0, 1\}$, $\nabla\phi^k$ is a diffeomorphism from \mathcal{X} to \mathcal{Y} and so is $\nabla\psi^k$ from \mathcal{Y} to \mathcal{X} . Thus by convexity of \mathcal{X} , $\nabla\psi^t(\mathcal{Y}) \subset \mathcal{X}$, which entails $\mathcal{Y} \subset \nabla\phi^t(\mathcal{X})$. The equality above then gives

$$\forall k \in \{0, 1\}, \forall y \in \mathcal{Y}, \quad \mu^k(y) \leq M_\rho \det(D^2\psi^k(y)).$$

On the other hand, the same equality gives

$$\forall t \in [0, 1], \forall y \in \mathcal{Y}, \quad \mu^t(y) \geq m_\rho \det(D^2\psi^t(y)).$$

Using the two inequalities above and the concavity of $\det^{1/d}$ over the set of non-negative symmetric matrices, we get for every $y \in \mathcal{Y}$,

$$\begin{aligned} \mu^t(y) &\geq m_\rho \det(D^2\psi^t(y)) \\ &\geq m_\rho \left((1-t) \det(D^2\psi^0)^{1/d} + t \det(D^2\psi^1)^{1/d} \right)^d \\ &\geq m_\rho \min(t, 1-t)^d (\det(D^2\psi^0(y)) + \det(D^2\psi^1(y))) \\ &\geq \frac{m_\rho}{M_\rho} \min(t, 1-t)^d (\mu^0(y) + \mu^1(y)). \end{aligned}$$

Using that $\text{spt}(\mu^0) = \text{spt}(\mu^1) = \mathcal{Y}$, this directly implies (3.9), which in turn gives us

$$\text{Var}_{\mu^t}(v) \geq 2 \min(t, 1-t)^d \frac{m_\rho}{M_\rho} \text{Var}_{\frac{1}{2}(\mu^0 + \mu^1)}(v).$$

Combined with inequality (3.8), this gives:

$$2 \min(t, 1-t)^d \text{Var}_{\frac{1}{2}(\mu^0 + \mu^1)}(v) \leq \frac{M_\rho^2}{m_\rho^2} \exp(M_\phi - m_\phi) \frac{d^2}{dt^2} \mathcal{K}_\rho(\psi^t).$$

After integrating over $t \in [0, s]$ and $s \in [0, 1]$, this ensures with (3.5) the inequality:

$$\frac{1}{(d+1)2^{d-1}} \text{Var}_{\frac{1}{2}(\mu^0 + \mu^1)}(v) \leq \frac{M_\rho^2}{m_\rho^2} \exp(M_\phi - m_\phi) (\mathcal{K}_\rho(\psi^1) - \mathcal{K}_\rho(\psi^0) + \langle v | \mu^0 \rangle), \quad (3.10)$$

where we recall that $v = \psi^1 - \psi^0$. We finally leverage an in-homogeneity in the scale of the Brenier potentials ϕ^0, ϕ^1 in the last inequality in order to improve the dependence on $M_\phi - m_\phi$. For any $\lambda > 0$, introduce for $k \in \{0, 1\}$ the Brenier potential $\phi_\lambda^k = \lambda\phi^k$ and denote $\mu_\lambda^k = (\nabla\phi_\lambda^k)_{\#}\rho$ the corresponding probability measure and $\psi_\lambda^k = (\phi_\lambda^k)^*$ its convex conjugate. Then using the formula $\psi_\lambda^k = \lambda\psi^k(\cdot/\lambda)$, one can notice that for any $\lambda > 0$,

$$\begin{aligned} \text{Var}_{\frac{1}{2}(\mu_\lambda^0 + \mu_\lambda^1)}(\psi_\lambda^1 - \psi_\lambda^0) &= \lambda^2 \text{Var}_{\frac{1}{2}(\mu^0 + \mu^1)}(\psi^1 - \psi^0), \\ \mathcal{K}_\rho(\psi_\lambda^1) - \mathcal{K}_\rho(\psi_\lambda^0) + \langle \psi_\lambda^1 - \psi_\lambda^0 | \mu_\lambda^0 \rangle &= \lambda (\mathcal{K}_\rho(\psi^1) - \mathcal{K}_\rho(\psi^0) + \langle \psi^1 - \psi^0 | \mu^0 \rangle), \\ \forall x \in \mathcal{X}, \forall k \in \{0, 1\}, \quad \lambda m_\phi &\leq \phi_\lambda^k(x) \leq \lambda M_\phi. \end{aligned}$$

Thus applying inequality (3.10) to $\psi_\lambda^0, \psi_\lambda^1$ and the associated quantities yields for any $\lambda > 0$

$$\frac{1}{(d+1)2^{d-1}} \text{Var}_{\frac{1}{2}(\mu^0 + \mu^1)}(\psi^1 - \psi^0) \leq \frac{M_\rho^2}{m_\rho^2} \frac{\exp(\lambda(M_\phi - m_\phi))}{\lambda} (\mathcal{K}_\rho(\psi^1) - \mathcal{K}_\rho(\psi^0) + \langle v | \mu^0 \rangle).$$

Choosing $\lambda = \frac{1}{M_\phi - m_\phi}$ in the last inequality finally gives

$$\frac{1}{e(d+1)2^{d-1}} \mathbb{V}\text{ar}_{\frac{1}{2}(\mu^0 + \mu^1)}(\psi^1 - \psi^0) \leq \frac{M_\rho^2}{m_\rho^2} (M_\phi - m_\phi) (\mathcal{K}_\rho(\psi^1) - \mathcal{K}_\rho(\psi^0) + \langle \psi^1 - \psi^0 | \mu^0 \rangle).$$

□

3.4 From the regular case to the general case

To deduce the general case of Theorem 3.1, we need to approximate the convex potentials ϕ^0, ϕ^1 on \mathcal{X} with strongly convex potentials ϕ_n^0, ϕ_n^1 that belong to $\mathcal{C}^2(\mathbb{R}^d)$ and that are such that their gradients $\nabla\phi_n^0, \nabla\phi_n^1$ induce diffeomorphisms between \mathcal{X} and a closed ball \mathcal{Y}_n . A regularization that uses a (standard) convolution does not seem directly feasible. Indeed, ϕ^k is defined on \mathcal{X} only, and its gradient may explode on the boundary of \mathcal{X} when μ^k has non-compact support, so that any convex extension of ϕ^k to \mathbb{R}^d has to take value $+\infty$ in this case.

Our strategy is as follows. First, we resort to Moreau-Yosida's regularization to approximate the functions ϕ^0, ϕ^1 by regular convex functions defined on \mathbb{R}^d . Then, we regularize the target probability measures associated to these approximated potentials and resort to Caffarelli's regularity theory to guarantee smoothness and strong convexity. Caffarelli's regularity theory results require smoothness assumptions on the source probability measure and strong convexity and smoothness assumption on the domain. We make these assumptions in the next proposition, but we will later show that these can be relaxed to get the general case of Theorem 3.1.

Proposition 3.4. *Let \mathcal{X} be a compact, smooth and strongly convex set, let ρ be a smooth probability density on \mathcal{X} and assume that ρ is bounded away from zero and infinity on this set. Let $\mu^0, \mu^1 \in \mathcal{P}_2(\mathbb{R}^d)$. For $k \in \{0, 1\}$, denote ϕ^k the Brenier potentials for the quadratic optimal transport from ρ to μ^k , and assume that there exists $m_\phi, M_\phi \in \mathbb{R}$ such for $k \in \{0, 1\}$ and any $x \in \mathcal{X}$,*

$$m_\phi \leq \phi^k(x) \leq M_\phi.$$

Then there exists a sequence of strongly convex functions $(\phi_n^0)_{n \in \mathbb{N}}, (\phi_n^1)_{n \in \mathbb{N}}$ in $\mathcal{C}^2(\mathbb{R}^d)$ such that if one introduces $\mu_n^k = (\nabla\phi_n^k)_{\#}\rho$, $\psi_n^k = (\phi_n^k)^$ and $\mu^k = (\nabla\phi^k)_{\#}\rho$, then:*

(i) *let $m_{\phi_n} = \min_{\mathcal{X}} \min_k \phi_n^k$, and $M_{\phi_n} = \max_{\mathcal{X}} \max_k \phi_n^k$. Then,*

$$m_\phi \leq \liminf_{n \rightarrow +\infty} m_{\phi_n} \leq \limsup_{n \rightarrow +\infty} M_{\phi_n} \leq M_\phi,$$

(ii) $\lim_{n \rightarrow +\infty} \mathcal{K}_\rho(\psi_n^1) - \mathcal{K}_\rho(\psi_n^0) + \langle \psi_n^1 - \psi_n^0 | \mu_n^0 \rangle = \mathcal{K}_\rho(\psi^1) - \mathcal{K}_\rho(\psi^0) + \langle \psi^1 - \psi^0 | \mu^0 \rangle,$

(iii) $\lim_{n \rightarrow +\infty} \mathbb{V}\text{ar}_{\frac{1}{2}(\mu_n^0 + \mu_n^1)}(\psi_n^1 - \psi_n^0) = \mathbb{V}\text{ar}_{\frac{1}{2}(\mu^0 + \mu^1)}(\psi^1 - \psi^0),$

(iv) *there exists a closed ball \mathcal{Y}_n such that for $k \in \{0, 1\}$, $\nabla\phi_n^k$ is a diffeomorphism between \mathcal{X} and \mathcal{Y}_n .*

Before proving this proposition, we recall some facts regarding Moreau-Yosida's regularization of convex functions. Quoting Section 3.4 of (Attouch, 1984), the Moreau-Yosida regularization of parameter $\lambda > 0$ of a closed and proper convex function

$f : \mathbb{R}^d \rightarrow \mathbb{R} \cup \{+\infty\}$ is defined for all $x \in \mathbb{R}^d$ by infimum convolution of the function f with $\frac{1}{2\lambda} \|\cdot\|^2$:

$$f_\lambda(x) = \min_{u \in \mathbb{R}^d} f(u) + \frac{1}{2\lambda} \|u - x\|^2.$$

The next lemma gathers a few properties of the Moreau-Yosida regularisation.

Lemma 3.5. *Let $f : \mathbb{R}^d \rightarrow \mathbb{R} \cup \{+\infty\}$ be a closed and proper convex function and let $\lambda > 0$. Then,*

- (i) $f_\lambda = (f^* + \frac{\lambda}{2} \|\cdot\|^2)^*$,
- (ii) for all $x \in \mathbb{R}^d$, $\lim_{\lambda \rightarrow 0} f_\lambda(x) = f(x)$,
- (iii) $f_\lambda \in \mathcal{C}^{1,1}(\mathbb{R}^d)$ and more precisely, ∇f_λ is $\frac{1}{\lambda}$ -Lipschitz,
- (iv) if f is differentiable at $x \in \mathbb{R}^d$, then $\lim_{\lambda \rightarrow 0} \nabla f_\lambda(x) = \nabla f(x)$,
- (v) if f is differentiable at $x \in \mathbb{R}^d$, then $\|\nabla f_\lambda(x)\| \leq \|\nabla f(x)\|$.

Proof. Point (i) is found in Proposition 3.3 of (Attouch, 1984), points (ii) and (iii) are found in Theorem 3.24 of (Attouch, 1984) and (iv) and (v) can be found in Proposition 2.6 of (Brézis, 1973). We include here the proof of (v) because this property is not very well-known. Let x be a point of \mathbb{R}^d where f is differentiable. By Theorem 3.24 of (Attouch, 1984), there exists a unique point $x_\lambda \in \mathbb{R}^d$ such that

$$f_\lambda(x) = f(x_\lambda) + \frac{1}{2\lambda} \|x_\lambda - x\|^2,$$

which satisfies $g_\lambda := \nabla f_\lambda(x) = \frac{1}{\lambda}(x - x_\lambda) \in \partial f(x_\lambda)$. By monotonicity of the subdifferential of f , this gives

$$\begin{aligned} 0 &\leq \langle x - x_\lambda | \nabla f(x) - g_\lambda \rangle = \langle x - x_\lambda | \nabla f(x) \rangle - \frac{1}{\lambda} \|x - x_\lambda\|^2 \\ &\leq \|x - x_\lambda\| \|\nabla f(x)\| - \frac{1}{\lambda} \|x - x_\lambda\|^2 = \|x - x_\lambda\| (\|\nabla f(x)\| - \|\nabla f_\lambda(x)\|). \quad \square \end{aligned}$$

Proof of Proposition 3.4. First regularization and truncation. Let $k \in \{0, 1\}$. We will first approximate the convex function ϕ^k with elements of $\mathcal{C}^{1,1}(\mathbb{R}^d)$. To do so, we denote by ϕ_α^k the Moreau-Yosida regularization of ϕ^k with parameter α . We let $\mu_\alpha^k = (\nabla \phi_\alpha^k) \# \rho$ and define ψ_α^k as the convex conjugate of ϕ_α^k . By Lemma 3.5, $\nabla \phi_\alpha^k$ is Lipschitz on the bounded domain \mathcal{X} , implying that for $k \in \{0, 1\}$, the images of $\nabla \phi_\alpha^k(\mathcal{X})$ are contained in a closed ball $\mathcal{Y}_\alpha = B(0, R_\alpha)$. We now prove the claimed convergences (i)-(iii), relying mainly on the dominated convergence theorem. We first note that if $f : \mathbb{R}^d \rightarrow \mathbb{R}$ satisfies the growth condition $|f(x)| \leq C(1 + \|x\|^2)$ for some constant C ,

$$\langle f | \mu_\alpha^k \rangle = \int_{\mathcal{X}} f(\nabla \phi_\alpha^k) d\rho \xrightarrow{\alpha \rightarrow 0} \int_{\mathcal{X}} f(\nabla \phi^k) d\rho = \langle f | \mu^k \rangle. \quad (3.11)$$

Indeed, Lemma 3.5.(iv) ensures that for every point $x \in \mathcal{X}$ where ϕ^k is differentiable, thus for ρ -almost every point x , one has $\lim_{\alpha \rightarrow 0} \nabla \phi_\alpha^k(x) = \nabla \phi^k(x)$. Besides, for all such x , Lemma 3.5.(v) gives

$$\left| f(\nabla \phi_\alpha^k(x)) \right| \leq C \left(1 + \left\| \nabla \phi_\alpha^k(x) \right\|^2 \right) \leq C \left(1 + \left\| \nabla \phi^k(x) \right\|^2 \right).$$

Moreover,

$$\int_{\mathcal{X}} \left(1 + \|\nabla \phi^k(x)\|^2\right) d\rho(x) \leq 1 + M_2(\nabla \phi^k(x)_{\#}\rho) = 1 + M_2(\mu^k) < +\infty.$$

Thus, the dominated convergence theorem ensures that (3.11) holds.

(i) We note that for $k \in \{0, 1\}$ and $x \in \mathcal{X}$, $m_\phi \leq \phi_\alpha^k(x) \leq M_\phi$. This is a simple consequence of the definition of the Moreau-Yosida regularization ϕ_α^k as an infimum convolution. Indeed for any $x \in \mathcal{X}$, we have on one hand:

$$\phi_\alpha^k(x) = \inf_{x' \in \mathcal{X}} \left(\phi^k(x') + \frac{1}{2\alpha} \|x - x'\|^2 \right) \geq \inf_{x' \in \mathcal{X}} \phi^k(x') + \inf_{x' \in \mathcal{X}} \frac{1}{2\alpha} \|x - x'\|^2 \geq m_\phi.$$

On the other hand,

$$\phi_\alpha^k(x) = \inf_{x' \in \mathcal{X}} \left(\phi^k(x') + \frac{1}{2\alpha} \|x - x'\|^2 \right) \leq \phi^k(x) + \frac{1}{2\alpha} \|x - x\|^2 \leq M_\phi.$$

(ii) By Lemma 3.5.(i),

$$\psi_\alpha^0 - \psi_\alpha^1 = \phi_\alpha^{0*} - \phi_\alpha^{1*} = \psi^0 + \frac{\alpha}{2} \|\cdot\|^2 - \psi^1 - \frac{\alpha}{2} \|\cdot\|^2 = \psi^0 - \psi^1.$$

Since ψ^0, ψ^1 are convex conjugates of functions defined on the compact set \mathcal{X} , the functions ψ^0 and ψ^1 are (globally) Lipschitz on \mathbb{R}^d (Remark 1.6). Thus $f = \psi_\alpha^1 - \psi_\alpha^0 = \psi^1 - \psi^0$ is also Lipschitz, and therefore satisfies a growth condition of the form $|f| \leq C(1 + \|x\|)$. By an application of (3.11), we get

$$\lim_{\alpha \rightarrow 0} \langle \psi_\alpha^1 - \psi_\alpha^0 | \mu_\alpha^0 \rangle = \lim_{\alpha \rightarrow 0} \langle \psi^1 - \psi^0 | \mu_\alpha^0 \rangle = \langle \psi^1 - \psi^0 | \mu^0 \rangle.$$

Moreover, for $k \in \{0, 1\}$, $\mathcal{K}_\rho(\psi_\alpha^k) = \int \phi_\alpha^k d\rho$. We just noticed that for all $\alpha \geq 0$, $|\phi_\alpha^k| \leq M_\phi$. Therefore, using the limit $\lim_{\alpha \rightarrow 0} \phi_\alpha^k(x) = \phi^k(x)$ for all $x \in \mathcal{X}$ given in Lemma 3.5.(i), we have with the dominated convergence theorem

$$\mathcal{K}_\rho(\psi_\alpha^k) \xrightarrow{\alpha \rightarrow 0} \mathcal{K}_\rho(\psi^k).$$

(iii) We use $\text{Var}_\mu(f) = \int f^2 d\mu - (\int f d\mu)^2$. Letting f as in the previous item, we get

$$\begin{aligned} \text{Var}_{\frac{1}{2}(\mu_\alpha^0 + \mu_\alpha^1)}(\psi_\alpha^1 - \psi_\alpha^0) &= \langle f^2 | \frac{1}{2}(\mu_\alpha^0 + \mu_\alpha^1) \rangle - \langle f | \frac{1}{2}(\mu_\alpha^0 + \mu_\alpha^1) \rangle^2, \\ \text{Var}_{\frac{1}{2}(\mu^0 + \mu^1)}(\psi^1 - \psi^0) &= \langle f^2 | \frac{1}{2}(\mu^0 + \mu^1) \rangle - \langle \psi^1 - \psi^0 | \frac{1}{2}(\mu^0 + \mu^1) \rangle^2. \end{aligned}$$

Since f is Lipschitz, both f and f^2 satisfy the growth condition allowing us to apply (3.11). We therefore get

$$\lim_{\alpha \rightarrow 0} \text{Var}_{\frac{1}{2}(\mu_\alpha^0 + \mu_\alpha^1)}(\psi_\alpha^1 - \psi_\alpha^0) = \text{Var}_{\frac{1}{2}(\mu^0 + \mu^1)}(\psi^1 - \psi^0).$$

We have all the desired properties (i)-(iii) but the potentials ϕ_α^k are not strongly convex and \mathcal{C}^2 on \mathbb{R}^d : they are merely $\mathcal{C}^{1,1}$. Moreover, the property (iv) does not hold. These properties will be obtained thanks to a second regularization.

Second regularization. From now on, we fix some $\alpha > 0$, and we denote $\mathcal{Y}_\alpha = B(0, R_\alpha)$ a closed ball that contains the supports of μ_α^0 and μ_α^1 . To construct the regularization of ϕ_α^k we will regularize the measures μ_α^k and solve an optimal transport problem. We first note that it is straightforward, e.g. using a simple convolution and truncation, to approximate the probability measures μ_α^k on \mathcal{Y}_α by smooth probability densities $\mu_{\alpha,\beta}^k$ supported on \mathcal{Y}_α , bounded away from zero and infinity on \mathcal{Y}_α and such that $\lim_{\beta \rightarrow 0} W_2(\mu_{\alpha,\beta}^k, \mu_\alpha^k) = 0$. By Caffarelli's regularity theory (e.g. Theorem 3.3 in (De Philippis and Figalli, 2014)), the optimal transport map $T_{\alpha,\beta}$ between ρ and $\mu_{\alpha,\beta}^k$ is the gradient of a strongly convex potential $\phi_{\alpha,\beta}^k$ belonging to $\mathcal{C}^2(\mathcal{X})$ and is actually a diffeomorphism between \mathcal{X} and \mathcal{Y}_α . By Theorem 4.4 in (Yan, 2014), the potential $\phi_{\alpha,\beta}^k$ can be extended into a \mathcal{C}^2 strongly convex function on \mathbb{R}^d . By stability of Kantorovich potentials (Theorem A.10), taking a subsequence if necessary, we can assume that $\phi_{\alpha,\beta}^k$ converges uniformly to ϕ_α^k on \mathcal{X} as $\beta \rightarrow 0$. Since $\nabla \phi_{\alpha,\beta}^k$ sends ρ to the measure $\mu_{\alpha,\beta}^k$, which is supported on $B(0, R_\alpha)$, we get $\|\nabla \phi_{\alpha,\beta}^k\| \leq R_\alpha$. Moreover, since the convex function $\phi_{\alpha,\beta}^k$ converges uniformly to ϕ_α^k as $\beta \rightarrow 0$, we get

$$\text{for a.e. } x \in \mathcal{X}, \lim_{\beta \rightarrow 0} \nabla \phi_{\alpha,\beta}^k(x) = \nabla \phi_\alpha^k(x).$$

This convergence result is also induced by the stability of optimal transport maps (see (Villani, 2008, Corollary 5.21) or Corollary A.11), since $\lim_{\beta \rightarrow 0} W_2(\mu_{\alpha,\beta}^k, \mu_\alpha^k) = 0$ and $\nabla \phi_{\alpha,\beta}^k$ (resp. $\nabla \phi_\alpha^k$) is the optimal transport map between ρ and $\mu_{\alpha,\beta}^k$ (resp. ρ and μ_α^k). From these two properties we get as above the desired convergence properties:

- (i) For $k \in \{0, 1\}$, $m_\phi \leq \liminf_{\beta \rightarrow 0} \min_{\mathcal{X}} \phi_{\alpha,\beta}^k(x) \leq \limsup_{\beta \rightarrow 0} \max_{\mathcal{X}} \phi_{\alpha,\beta}^k(x) \leq M_\phi$.
- (ii) $\lim_{\beta \rightarrow 0} \langle \phi_{\alpha,\beta}^1 - \phi_{\alpha,\beta}^0 | \rho \rangle + \langle (\phi_{\alpha,\beta}^1)^* - (\phi_{\alpha,\beta}^0)^* | \mu_{\alpha,\beta}^0 \rangle = \langle \phi_\alpha^1 - \phi_\alpha^0 | \rho \rangle + \langle (\phi_\alpha^1)^* - (\phi_\alpha^0)^* | \mu_\alpha^0 \rangle$.
- (iii) $\lim_{\beta \rightarrow 0} \text{Var}_{\frac{1}{2}(\mu_{\alpha,\beta}^0 + \mu_{\alpha,\beta}^1)}((\phi_{\alpha,\beta}^1)^* - (\phi_{\alpha,\beta}^0)^*) = \text{Var}_{\frac{1}{2}(\mu_\alpha^0 + \mu_\alpha^1)}((\phi_\alpha^1)^* - (\phi_\alpha^0)^*)$.

The sequence in the statement of the proposition is finally constructed using a diagonal argument. \square

Proposition 3.6. *In addition to the assumptions of Theorem 3.1, assume that \mathcal{X} is a smooth and strongly convex set and that the density ρ is smooth. Then, (3.1) holds.*

Proof. Let ϕ_n^0, ϕ_n^1 be the sequence of \mathcal{C}^2 and strongly convex potentials constructed by Proposition 3.4, converging respectively to ϕ^0 and ϕ^1 , and such that $\nabla \phi_n^0, \nabla \phi_n^1$ are diffeomorphisms from \mathcal{X} to a ball \mathcal{Y}_n . By Proposition 3.3, (3.1) holds for ϕ_n^0, ϕ_n^1 :

$$\text{Var}_{\frac{1}{2}(\mu_n^0 + \mu_n^1)}(\psi_n^1 - \psi_n^0) \leq C_d \frac{M_\rho^2}{m_\rho^2} (M_{\phi_n} - m_{\phi_n}) (\mathcal{K}_\rho(\psi_n^1) - \mathcal{K}_\rho(\psi_n^0) + \langle \psi_n^1 - \psi_n^0 | \mu_n^0 \rangle).$$

By the claims (i)-(iv) in Proposition 3.4, all the terms in this inequality converge as $n \rightarrow +\infty$ and establish (3.1) in the limit. \square

Proof of Theorem 3.1. Let \mathcal{X} be a bounded convex set and assume that ρ is a probability density supported on this set and satisfying $m_\rho \leq \rho \leq M_\rho$ on it. We extend ρ by m_ρ outside of \mathcal{X} . One can construct a sequence \mathcal{X}_n of smooth and strongly convex sets included in \mathcal{X} and converging to \mathcal{X} in the Hausdorff sense as $n \rightarrow +\infty$ (Schneider, 2013,

§3.3). Let K be a smooth, non-negative and compactly supported function, $K_n(x) = n^d K(nx)$ and define

$$\rho_n = \frac{1}{Z_n} (\rho * K_n)|_{\mathcal{X}_n}, m_{\rho_n} = \frac{m_\rho}{Z_n}, M_{\rho_n} = \frac{M_\rho}{Z_n},$$

where Z_n is a constant ensuring that ρ_n belongs to $\mathcal{P}(\mathcal{X}_n)$. We define $\mu_n^k = (\nabla \phi^k)_\# \rho_n$. Applying Proposition 3.6 to (\mathcal{X}_n, ρ_n) and (ϕ^0, ϕ^1) , we have:

$$\mathbb{V}\text{ar}_{\frac{1}{2}(\mu_n^0 + \mu_n^1)}(\psi^1 - \psi^0) \leq C_d \frac{M_\rho^2}{m_{\rho_n}^2} (M_\phi - m_\phi) (\mathcal{K}_\rho(\psi^1) - \mathcal{K}_\rho(\psi^0) + \langle \psi^1 - \psi^0 | \mu_n^0 \rangle). \quad (3.12)$$

By construction, $\lim_{n \rightarrow +\infty} Z_n = 1$ and ρ_n converges to ρ in $L^1(\mathcal{X})$. Thus up to subsequences, ρ_n converges pointwise almost everywhere to ρ . Setting $f = \psi^0 - \psi^1$, we have

$$\langle \psi^0 - \psi^1 | \mu_n^0 \rangle = \int_{\mathcal{X}} f(\nabla \phi^0) \rho_n(x) dx \xrightarrow{n \rightarrow +\infty} \int_{\mathcal{X}} f(\nabla \phi^0) \rho(x) dx = \langle \psi^0 - \psi^1 | \mu^0 \rangle.$$

The limit in the above equation is proven as in Proposition 3.4, using that f is Lipschitz, that $M_2(\mu_0) < +\infty$ and applying the dominated convergence theorem. All the terms can be dealt with in a similar manner. Taking the limit $n \rightarrow +\infty$ in (3.12) gives the desired estimate (3.1). \square

An entropic approach

Abstract

This chapter considers the entropy-regularized quadratic optimal transport problem. The dual of this problem is shown to feature an *entropic* Kantorovich functional, that enjoys a non-trivial strong convexity estimate reminiscent of the estimates presented in Chapters 2 and 3 for the *classical* Kantorovich functional. This estimate, originally obtained in (Delalande, 2022), is derived properly in the semi-discrete setting and is shown to be a consequence of the Prékopa-Leindler inequality. It is finally shown that this estimate may be used, thanks to density arguments, to recover the strong convexity estimate of Chapter 3 for the classical Kantorovich functional.

4.1 Introduction

In this chapter, we consider the entropic regularization of the optimal transport problem. Let \mathcal{X}, \mathcal{Y} be subsets of \mathbb{R}^d . For two probability measures $\rho \in \mathcal{P}_2(\mathcal{X}), \mu \in \mathcal{P}_2(\mathcal{Y})$, the quadratic optimal transport problem between ρ and μ with entropic regularization of weight $\varepsilon \geq 0$ reads

$$\min_{\gamma \in \Gamma(\rho, \mu)} \int_{\mathcal{X} \times \mathcal{Y}} \|x - y\|^2 d\gamma(x, y) + \varepsilon \text{KL}(\gamma | \rho \otimes \mu), \quad (4.1)$$

where $\Gamma(\rho, \mu)$ denotes the set of couplings between ρ and μ and KL denotes the Kullback-Leibler divergence or relative entropy (up to an additive term):

$$\text{KL}(\gamma | \rho \otimes \mu) = \int_{\mathcal{X} \times \mathcal{Y}} \left(\log \left(\frac{d\gamma}{d\rho \otimes \mu}(x, y) \right) - 1 \right) d\gamma(x, y)$$

if $\gamma \ll \rho \otimes \mu$ and $+\infty$ otherwise. When $\varepsilon = 0$, problem (4.1) coincides with the usual quadratic optimal transport problem between ρ and μ introduced in Chapter 1. Problem (4.1) thus corresponds to the optimal transport problem regularized with the 1-strongly-convex *entropic* penalty $\gamma \mapsto \text{KL}(\gamma | \rho \otimes \mu)$, whose weight is controlled by the parameter ε . When $\varepsilon > 0$, problem (4.1) can be shown to be equivalent to the static Schrödinger problem (Léonard, 2014) which was initially considered by Schrödinger in statistical physics (Schrödinger, 1931). Very formally, this problem studies a collection of gas particles in \mathbb{R}^d , whose spatial configurations at two moments $t = 0$ and $t = 1$ are known and given by

the measures ρ and μ respectively. Assuming that the particles evolve at a temperature $\varepsilon > 0$ (i.e. they are subject to a Brownian motion), Schrödinger wondered what were the most probable paths of each particle between these two instants. The (unique) solution of (4.1) convolved with a Brownian bridge gives the path measure solution to this problem (Föllmer, 1988). In the recent years, problem (4.1) was revisited by (Cuturi, 2013), who initiated a fruitful line of works that took benefit from the computational and statistical advantages of problem (4.1) with $\varepsilon > 0$ over classical optimal transport. We refer to Chapter 7 for a greater emphasis on these advantages, and more generally to the monograph (Peyré and Cuturi, 2019) or lecture notes (Nutz, 2022) for an introduction to entropy-regularized optimal transport.

In this chapter, we show that the *entropic* Kantorovich functional that appears in the dual formulation of (4.1) may enjoy a non-trivial strong-convexity estimate that is reminiscent of the estimates of Chapters 2 and 3. For technical and practical reasons justified in Section 4.2, we will limit ourselves to the semi-discrete setting and assume that the source ρ is absolutely continuous and supported on a compact convex $\mathcal{X} \subset \mathbb{R}^d$ and that the target μ is discretely supported on a finite set $\mathcal{Y} = \{y_1, \dots, y_N\} \subset \mathbb{R}^d$. In this setting, we will show that the *semi-discrete entropic Kantorovich functional* $\mathcal{K}_\rho^\varepsilon : \mathbb{R}^N \rightarrow \mathbb{R}$ involved in the dual problem of (4.1) (defined properly in Section 4.3) is twice differentiable and satisfies the following strong convexity estimate, originally established in (Delalande, 2022):

Theorem (Theorem 4.4). *Let ρ be a probability density over a compact convex set \mathcal{X} , satisfying $0 < m_\rho \leq \rho \leq M_\rho$ and let $\mathcal{Y} = \{y_1, \dots, y_N\} \subset \mathbb{R}^d$. Let $\varepsilon > 0$ and let $\psi \in \mathbb{R}^N$. Define*

$$\psi^{c,\varepsilon} : x \in \mathcal{X} \mapsto \varepsilon \log \left(\sum_{i=1}^N e^{\frac{\langle x|y_i \rangle - \psi_i}{\varepsilon}} \right) \in \mathbb{R}, \quad \mathcal{K}_\rho^\varepsilon : \psi \in \mathbb{R}^N \mapsto \int_{\mathcal{X}} \psi^{c,\varepsilon} d\rho \in \mathbb{R},$$

and denote $m_\psi = \min_{\mathcal{X}} \psi^{c,\varepsilon}$ and $M_\psi = \max_{\mathcal{X}} \psi^{c,\varepsilon}$. Then $\mathcal{K}_\rho^\varepsilon$ is \mathcal{C}^2 and there exists a discrete measure $\mu_\psi^\varepsilon \in \mathcal{P}(\mathcal{Y})$ depending only on ψ and ε such that for any $v \in \mathbb{R}^N$,

$$\text{Var}_{\mu_\psi^\varepsilon}(v) \leq \left(e^{(M_\psi - m_\psi)} \frac{M_\rho}{m_\rho} + \varepsilon \right) \langle v | \nabla^2 \mathcal{K}_\rho^\varepsilon(\psi) v \rangle.$$

Remark 4.1. The measure μ_ψ^ε is known explicitly and we refer to Section 4.4 for its definition.

Such strong convexity estimate is interesting in its own right for the study of entropic optimal transport problems (with $\varepsilon > 0$ fixed). We show however at the end of this chapter that taking the limit $\varepsilon \rightarrow 0$ of this estimate also allows to recover the strong convexity estimate of Theorem 3.1 proven in Chapter 3.

Outline. In order to prove Theorem 4.4, we proceed as follows. First, we define the entropic Kantorovich functional and discuss the role played by the target measure in this functional (§4.2). This discussion leads us to the semi-discrete assumption, under which we compute the first and second derivatives of the entropic Kantorovich functional (§4.3). The Prékopa-Leindler inequality is then used to derive a lower-bound on the smallest eigenvalue of the Hessian of the semi-discrete entropic Kantorovich functional, from which Theorem 4.4 is deduced (§4.4). This theorem is then shown to entail the strong convexity estimate of Theorem 3.1 for the *classical* Kantorovich functional (§4.5).

4.2 Entropic Kantorovich functional

Similarly to classical optimal transport, developing the square in the integral term of (4.1) shows that the regularized quadratic optimal transport problem with regularization parameter 2ε is equivalent to the following regularized maximum correlation problem:

$$\max_{\gamma \in \Gamma(\rho, \mu)} \int_{\mathcal{X} \times \mathcal{Y}} \langle x|y \rangle d\gamma(x, y) - \varepsilon \text{KL}(\gamma | \rho \otimes \mu). \quad (4.2)$$

When $\varepsilon > 0$, this problem is strongly-concave (by strong-convexity of $\gamma \mapsto \text{KL}(\gamma | \rho \otimes \mu)$) and thus admits a unique solution. Moreover it admits the following (semi-)dual formulation and strong duality holds (see for instance Sections 2 of (Genevay et al., 2016; Bercu and Bigot, 2020)):

$$\min_{\psi \in \mathcal{C}^0(\mathcal{Y})} \int_{\mathcal{X}} \psi^{c, \varepsilon, \mu} d\rho + \int_{\mathcal{Y}} \psi d\mu + \varepsilon,$$

where $\psi^{c, \varepsilon, \mu}$ corresponds to the following (c, ε, μ) -transform of ψ : $\forall x \in \mathcal{X}$,

$$\psi^{c, \varepsilon, \mu}(x) = \varepsilon \log \left(\int_{\mathcal{Y}} e^{\frac{\langle x|y \rangle - \psi(y)}{\varepsilon}} d\mu(y) \right).$$

Introducing the *entropic* Kantorovich functional $\mathcal{K}_{\rho, \mu}^{\varepsilon} : \psi \mapsto \int_{\mathcal{X}} \psi^{c, \varepsilon, \mu} d\rho + \varepsilon$, the dual formulation can be rewritten

$$\min_{\psi \in \mathcal{C}^0(\mathcal{Y})} \mathcal{K}_{\rho, \mu}^{\varepsilon}(\psi) + \langle \psi | \mu \rangle. \quad (4.3)$$

Using Hölder's inequality, one can show that the mapping $\psi \mapsto \psi^{c, \varepsilon, \mu}(x)$ is convex for any $x \in \mathbb{R}^d$ (see the next section). This ensures that the entropic Kantorovich functional $\mathcal{K}_{\rho, \mu}^{\varepsilon}$ is convex. We now wonder, as we did for the *classical* Kantorovich functional in Chapters 2 and 3, if this functional is strongly convex.

Entropic Kantorovich functional and target measure. In order to study the strong convexity of $\mathcal{K}_{\rho, \mu}^{\varepsilon}$, it is tempting to try to proceed as we did in Chapters 2 and 3: first, consider two different target measures μ^0, μ^1 and select two corresponding minimizers ψ^0, ψ^1 in (4.3). Then, try to differentiate twice $\mathcal{K}_{\rho, \mu}^{\varepsilon}$ along the path $((1-t)\psi^0 + t\psi^1)_{t \in [0,1]}$ and look for a lower-bound on the second derivative of $\mathcal{K}_{\rho, \mu}^{\varepsilon}$ along this path. However, this approach fails when the probability measures μ^0 and μ^1 have incomparable support. Indeed, first notice that the potentials ψ^0 and ψ^1 are defined on the supports of μ^0 and μ^1 respectively, so that the interpolant $(1-t)\psi^0 + t\psi^1$ may not be well-defined. One could extend the domains of ψ^0, ψ^1 by means of double (c, ε, \cdot) -transform-like operations, but there would remain another problem: the target measure actually appears in the definition of the entropic Kantorovich functional $\mathcal{K}_{\rho, \mu}^{\varepsilon}$ (through the (c, ε, μ) -transform): in order to differentiate this functional along $((1-t)\psi^0 + t\psi^1)_{t \in [0,1]}$, it is thus not clear what should be the target measure taken in the definition of the (c, ε, μ) -transform.

A possible workaround is the following: introduce a fixed reference positive measure $\sigma \in \mathcal{M}_+(\mathcal{Y})$ and denote $\mathcal{P}_2^{\sigma}(\mathcal{Y})$ the set of probability measures with finite second moment that are equivalent to σ :

$$\mathcal{P}_2^{\sigma}(\mathcal{Y}) = \{\mu \in \mathcal{P}_2(\mathcal{Y}) | \mu \ll \sigma \text{ and } \sigma \ll \mu\}.$$

Then notice that for any $\rho \in \mathcal{P}_2(\mathcal{X})$ and $\mu \in \mathcal{P}_2^\sigma(\mathcal{Y})$, one has for any $\gamma \in \Gamma(\rho, \mu)$,

$$\text{KL}(\gamma|\rho \otimes \mu) = \text{KL}(\gamma|\rho \otimes \sigma) - \text{KL}(\mu|\sigma).$$

This means that for $\mu \in \mathcal{P}_2^\sigma(\mathcal{Y})$, problem (4.2) is equivalent to the following problem:

$$\max_{\gamma \in \Gamma(\rho, \mu)} \int_{\mathcal{X} \times \mathcal{Y}} \langle x|y \rangle d\gamma(x, y) - \varepsilon \text{KL}(\gamma|\rho \otimes \sigma),$$

which admits the following dual problem:

$$\min_{\psi \in \mathcal{C}^0(\mathcal{Y})} \int_{\mathbb{R}^d} \psi^{c, \varepsilon} d\rho + \int_{\mathbb{R}^d} \psi d\mu + \varepsilon,$$

where $\psi^{c, \varepsilon}$ corresponds to the following (c, ε) -transform of ψ (that depends on the fixed reference σ): $\forall x \in \mathcal{X}$,

$$\psi^{c, \varepsilon}(x) = \varepsilon \log \left(\int_{\mathcal{Y}} e^{\frac{\langle x|y \rangle - \psi(y)}{\varepsilon}} d\sigma(y) \right). \quad (4.4)$$

With a fixed reference σ for the targets, the *entropic* Kantorovich functional may now be denoted $\mathcal{K}_\rho^\varepsilon : \psi \mapsto \int_{\mathcal{X}} \psi^{c, \varepsilon} d\rho + \varepsilon$, so that for all $\mu \in \mathcal{P}_2^\sigma(\mathcal{Y})$, the dual formulation reads

$$\min_{\psi \in \mathcal{C}^0(\mathcal{Y})} \mathcal{K}_\rho^\varepsilon(\psi) + \langle \psi | \mu \rangle. \quad (4.5)$$

Semi-discrete assumption. Now choosing $\mu^0, \mu^1 \in \mathcal{P}_2^\sigma(\mathcal{Y})$, it is possible to interpolate two corresponding minimizers ψ^0, ψ^1 of (4.5) and differentiate twice $\mathcal{K}_\rho^\varepsilon$ along the path $((1-t)\psi^0 + t\psi^1)_{t \in [0,1]}$. One may then lower-bound the second derivative of $\mathcal{K}_\rho^\varepsilon$ along this path using the Prékopa-Leindler inequality and thus derive a strong convexity estimate for the entropic Kantorovich functional. In the following, we do this under the assumption that the set \mathcal{Y} is finite and σ is the counting measure on \mathcal{Y} . As in Chapter 2, this places ourselves in the semi-discrete setting with an absolutely continuous source and a discrete target. We make this assumption for two reasons: these semi-discrete computations will prove useful in Chapter 7; and by density arguments already developed in Chapters 2 and 3, they allow to recover the general strong convexity estimate of Theorem 3.1 for classical optimal transport. We emphasize again that the computations that follow could be done very similarly without a semi-discrete assumption.

4.3 Semi-discrete entropic Kantorovich functional

From now on, we assume that \mathcal{X} is a compact convex subset of \mathbb{R}^d and $\rho \in \mathcal{P}(\mathcal{X})$ is an absolutely continuous probability measure on \mathcal{X} . We consider $\mathcal{Y} = \{y_1, \dots, y_N\} \subset \mathbb{R}^d$ a set of N points in \mathbb{R}^d and we let σ be the counting measure associated to this set, i.e. $\sigma = \sum_{i=1}^N \delta_{y_i}$. Let $\mu = \sum_{i=1}^N \mu_i \delta_{y_i} \in \mathcal{P}(\mathcal{Y})$ where for all $i, \mu_i \geq \underline{\mu} > 0$. In this setting, problem (4.5) reads

$$\min_{\psi \in \mathcal{C}^0(\mathcal{Y})} \int_{\mathcal{X}} \varepsilon \log \left(\sum_{i=1}^N e^{\frac{\langle x|y_i \rangle - \psi(y_i)}{\varepsilon}} \right) d\rho(x) + \sum_{i=1}^N \psi(y_i) \mu_i.$$

In this problem, the function ψ is defined and evaluated only at the points of \mathcal{Y} , so that it can be conflated with the vector $\boldsymbol{\psi} = (\psi(y_i))_{1 \leq i \leq N} \in \mathbb{R}^N$. For such $\boldsymbol{\psi}$, we abuse the

notation and denote for all $x \in \mathcal{X}$, $\psi^{c,\varepsilon}(x) = \psi^{c,\varepsilon}(x)$ (defined in (4.4)). This allows us to introduce the *semi-discrete entropic Kantorovich functional*:

$$\mathcal{K}_\rho^\varepsilon : \psi \in \mathbb{R}^N \mapsto \int_{\mathcal{X}} \psi^{c,\varepsilon} d\rho + \varepsilon \in \mathbb{R}.$$

In the considered setting, (4.5) is thus equivalent to the following finite-dimensional problem:

$$\min_{\psi \in \mathbb{R}^N} \mathcal{K}_\rho^\varepsilon(\psi) + \langle \psi | \mu \rangle, \quad (4.6)$$

where μ is conflated with the vector $(\mu_i)_{1 \leq i \leq N} \in \mathbb{R}^N$. For any $\psi, \varphi \in \mathbb{R}^N$ and $\lambda \in (0, 1)$, Hölder's inequality ensures

$$\begin{aligned} \mathcal{K}_\rho^\varepsilon(\lambda\psi + (1-\lambda)\varphi) &= \int_{\mathcal{X}} \varepsilon \log \left(\sum_{i=1}^N \left(e^{\frac{\langle x|y_i\rangle - \psi_i}{\varepsilon}} \right)^\lambda \left(e^{\frac{\langle x|y_i\rangle - \varphi_i}{\varepsilon}} \right)^{1-\lambda} \right) d\rho(x) \\ &\leq \int_{\mathcal{X}} \varepsilon \log \left(\left(\sum_{i=1}^N e^{\frac{\langle x|y_i\rangle - \psi_i}{\varepsilon}} \right)^\lambda \left(\sum_{i=1}^N e^{\frac{\langle x|y_i\rangle - \varphi_i}{\varepsilon}} \right)^{1-\lambda} \right) d\rho(x) \\ &= \lambda \mathcal{K}_\rho^\varepsilon(\psi) + (1-\lambda) \mathcal{K}_\rho^\varepsilon(\varphi), \end{aligned}$$

with equality if and only if there exists a measurable function $\alpha : \mathcal{X} \rightarrow \mathbb{R}_+^*$ such that for ρ -almost-every $x \in \mathcal{X}$, for all $i \in \{1, \dots, N\}$,

$$e^{\frac{\langle x|y_i\rangle - \psi_i}{\varepsilon}} = \alpha(x) e^{\frac{\langle x|y_i\rangle - \varphi_i}{\varepsilon}}.$$

After simplifying by $e^{\frac{\langle x|y_i\rangle}{\varepsilon}}$, we notice that this is possible only if α is constant ρ -a.e., and it is equivalent to require $\psi_i = \varphi_i + \varepsilon \log \alpha$ for all $i \in \{1, \dots, N\}$, that is to require that ψ and φ only differ by a constant vector. The functional $\mathcal{K}_\rho^\varepsilon$ is thus convex on \mathbb{R}^N and strictly convex on $(\mathbb{1}_N)^\perp$, where $\mathbb{1}_N$ denotes the all-ones vector of \mathbb{R}^N . Problem (4.6) thus admits a unique solution up to a constant, and such solution (up to a constant) will be denoted ψ^ε . Note that in the literature, these solutions are sometimes called Schrödinger potentials (in reference to the above-mentioned Schrödinger problem) or Sinkhorn potentials (in reference to Sinkhorn's algorithm that one can use for the numerical resolution of (4.6), see (Peyré and Cuturi, 2019)). It is then possible to characterize ψ^ε with a first-order condition in (4.6). For this, we need to compute the gradient of $\mathcal{K}_\rho^\varepsilon$.

Derivatives of the semi-discrete entropic Kantorovich functional. Let us recall the definition of the semi-discrete entropic Kantorovich functional: for any $\psi \in \mathbb{R}^N$,

$$\mathcal{K}_\rho^\varepsilon(\psi) = \int_{\mathcal{X}} \psi^{c,\varepsilon} d\rho = \int_{\mathcal{X}} \varepsilon \log \left(\sum_{i=1}^N e^{\frac{\langle x|y_i\rangle - \psi_i}{\varepsilon}} \right) d\rho(x).$$

In order to differentiate $\mathcal{K}_\rho^\varepsilon$, one can try to differentiate the mapping

$$\psi \mapsto \psi^{c,\varepsilon}(x) = \varepsilon \log \left(\sum_{i=1}^N e^{\frac{\langle x|y_i\rangle - \psi_i}{\varepsilon}} \right)$$

for all $x \in \mathcal{X}$. The following lemma ensures that this mapping is \mathcal{C}^2 and gives its first and second derivatives.

Lemma 4.2. *Let $\varepsilon > 0$. For any $x \in \mathcal{X}$, the mapping $\boldsymbol{\psi} \mapsto \boldsymbol{\psi}^{c,\varepsilon}(x)$ is \mathcal{C}^2 on \mathbb{R}^N . For any $\boldsymbol{\psi} \in \mathbb{R}^N$ and $x \in \mathcal{X}$, introduce the vector $\gamma_x^\varepsilon(\boldsymbol{\psi}) \in \mathbb{R}^N$ defined for any $i \in \{1, \dots, N\}$ by*

$$\gamma_x^\varepsilon(\boldsymbol{\psi})_i = \frac{e^{\frac{\langle x|y_i \rangle - \psi_i}{\varepsilon}}}{\sum_{j=1}^N e^{\frac{\langle x|y_j \rangle - \psi_j}{\varepsilon}}}.$$

Then one has for any $x \in \mathcal{X}$,

$$\begin{aligned} \nabla_{\boldsymbol{\psi}}[\boldsymbol{\psi}^{c,\varepsilon}(x)] &= -\gamma_x^\varepsilon(\boldsymbol{\psi}), \\ \nabla_{\boldsymbol{\psi}}^2[\boldsymbol{\psi}^{c,\varepsilon}(x)] &= \frac{1}{\varepsilon} \left(\text{diag}(\gamma_x^\varepsilon(\boldsymbol{\psi})) - \gamma_x^\varepsilon(\boldsymbol{\psi})\gamma_x^\varepsilon(\boldsymbol{\psi})^\top \right). \end{aligned}$$

Remark 4.3. Intuitively, $\gamma_x^\varepsilon(\boldsymbol{\psi})_i$ is a smoothed version of the indicator function associated to the i -th Laguerre cell of $\boldsymbol{\psi}$ introduced in Section 2.2 of Chapter 2. It actually corresponds to the ratio of mass sent from x to y_i by the coupling derived from the proposed potential $\boldsymbol{\psi}$ in the ε -entropic transport.

Before proving Lemma 4.2, let us note the following consequence: by this lemma, $\boldsymbol{\psi} \mapsto \mathcal{K}_\rho^\varepsilon(\boldsymbol{\psi})$ is \mathcal{C}^2 on \mathbb{R}^N , and it admits the following derivatives:

$$\nabla \mathcal{K}_\rho^\varepsilon(\boldsymbol{\psi}) = -\mathbb{E}_{x \sim \rho} \gamma_x^\varepsilon(\boldsymbol{\psi}), \quad (4.7)$$

$$\nabla^2 \mathcal{K}_\rho^\varepsilon(\boldsymbol{\psi}) = \frac{1}{\varepsilon} \mathbb{E}_{x \sim \rho} \left(\text{diag}(\gamma_x^\varepsilon(\boldsymbol{\psi})) - \gamma_x^\varepsilon(\boldsymbol{\psi})\gamma_x^\varepsilon(\boldsymbol{\psi})^\top \right). \quad (4.8)$$

This ensures that a minimizer $\boldsymbol{\psi}^\varepsilon$ of (4.6) is characterized by the following first order condition:

$$\mu = \mathbb{E}_{x \sim \rho} \gamma_x^\varepsilon(\boldsymbol{\psi}^\varepsilon), \quad (4.9)$$

where we conflated again $\mu \in \mathcal{P}(\mathcal{Y})$ with the vector $(\mu_i)_{1 \leq i \leq N} \in \mathbb{R}^N$. Interestingly, the probability measure γ^ε on $\mathcal{X} \times \mathcal{Y}$ defined for any $x \in \mathcal{X}$ and $i \in \{1, \dots, N\}$ by

$$\frac{d\gamma^\varepsilon}{d\rho \otimes \sigma}(x, y_i) = \gamma_x^\varepsilon(\boldsymbol{\psi}^\varepsilon)_i$$

is a coupling between ρ and μ . Moreover, this coupling satisfies

$$\int_{\mathcal{X} \times \mathcal{Y}} \langle x|y \rangle d\gamma^\varepsilon(x, y) - \varepsilon \text{KL}(\gamma^\varepsilon | \rho \otimes \sigma) = \mathcal{K}_\rho^\varepsilon(\boldsymbol{\psi}^\varepsilon) + \langle \boldsymbol{\psi}^\varepsilon | \mu \rangle,$$

so that there is no duality gap and γ^ε is the unique solution to the primal problem (4.2).

Proof of Lemma 4.2. For any $x \in \mathcal{X}$, the mapping $\boldsymbol{\psi} \mapsto \boldsymbol{\psi}^{c,\varepsilon}(x)$ is obviously \mathcal{C}^2 on \mathbb{R}^N and one has for any $i \in \{1, \dots, N\}$,

$$\partial_{\psi_i}[\boldsymbol{\psi}^{c,\varepsilon}(x)] = -\frac{e^{\frac{\langle x|y_i \rangle - \psi_i}{\varepsilon}}}{\sum_{j=1}^N e^{\frac{\langle x|y_j \rangle - \psi_j}{\varepsilon}}} = -\gamma_x^\varepsilon(\boldsymbol{\psi})_i,$$

which gives the gradient formula. The second derivative of $\boldsymbol{\psi} \mapsto \boldsymbol{\psi}^{c,\varepsilon}(x)$ for any $x \in \mathcal{X}$ is deduced from the following computations: for any $i \in \{1, \dots, N\}$,

$$\begin{aligned} \partial_{\psi_i}^2[\boldsymbol{\psi}^{c,\varepsilon}(x)] &= \frac{1}{\varepsilon} \left(\frac{e^{\frac{\langle x|y_i \rangle - \psi_i}{\varepsilon}}}{\sum_{j=1}^N e^{\frac{\langle x|y_j \rangle - \psi_j}{\varepsilon}}} - \left(\frac{e^{\frac{\langle x|y_i \rangle - \psi_i}{\varepsilon}}}{\sum_{j=1}^N e^{\frac{\langle x|y_j \rangle - \psi_j}{\varepsilon}}} \right)^2 \right) \\ &= \frac{1}{\varepsilon} (\gamma_x^\varepsilon(\boldsymbol{\psi})_i - \gamma_x^\varepsilon(\boldsymbol{\psi})_i^2). \end{aligned}$$

And for any $i \neq j$,

$$\begin{aligned} \partial_{\psi_i} \partial_{\psi_j} [\psi^{c,\varepsilon}(x)] &= -\frac{1}{\varepsilon} \frac{e^{\frac{\langle x|y_i\rangle - \psi_i}{\varepsilon}}}{\sum_{k=1}^N e^{\frac{\langle x|y_k\rangle - \psi_k}{\varepsilon}}} \frac{e^{\frac{\langle x|y_j\rangle - \psi_j}{\varepsilon}}}{\sum_{k=1}^N e^{\frac{\langle x|y_k\rangle - \psi_k}{\varepsilon}}} \\ &= -\frac{1}{\varepsilon} \gamma_x^\varepsilon(\psi)_i \gamma_x^\varepsilon(\psi)_j. \end{aligned} \quad \square$$

4.4 Strong convexity of the semi-discrete entropic Kantorovich functional

In this section, we derive a strong convexity estimate for the semi-discrete entropic Kantorovich functional $\mathcal{K}_\rho^\varepsilon$. In order to state our estimate, we need to introduce the *target probability measure associated to a potential*: for any $\psi \in \mathbb{R}^N$ and $\varepsilon > 0$, let $\mu_\psi^\varepsilon \in \mathcal{P}(\mathcal{Y})$ be defined for any $y_i \in \mathcal{Y}$ by

$$\mu_\psi^\varepsilon(y_i) = \mathbb{E}_{x \sim \rho} \gamma_x^\varepsilon(\psi)_i.$$

Notice then that from (4.9), $\psi^\varepsilon \in \mathbb{R}^N$ is a minimizer of (4.6) if and only if $\mu_{\psi^\varepsilon}^\varepsilon = \mu$. The preceding computations then allow to show the following strong-convexity estimate for $\mathcal{K}_\rho^\varepsilon$.

Theorem 4.4 (Strong convexity of $\mathcal{K}_\rho^\varepsilon$). *Let ρ be a probability density over a compact convex set \mathcal{X} , satisfying $0 < m_\rho \leq \rho \leq M_\rho$ and let $\mu \in \mathcal{P}(\mathcal{Y})$ where $\mathcal{Y} = \{y_1, \dots, y_N\} \subset \mathbb{R}^d$. Let $\varepsilon > 0$ and let $\psi \in \mathbb{R}^N$. Denote $m_\phi = \min_{\mathcal{X}} \psi^{c,\varepsilon}$ and $M_\phi = \max_{\mathcal{X}} \psi^{c,\varepsilon}$. Then for any $v \in \mathbb{R}^N$,*

$$\text{Var}_{\mu_\psi^\varepsilon}(v) \leq \left(e^{(M_\phi - m_\phi)} \frac{M_\rho}{m_\rho} + \varepsilon \right) \langle v | \nabla^2 \mathcal{K}_\rho^\varepsilon(\psi) v \rangle. \quad (4.10)$$

Remark 4.5 (Bounds on $\psi^{c,\varepsilon}$). Notice that in our compact semi-discrete context, $\psi^{c,\varepsilon}$ is bounded so that $-\infty < m_\phi \leq M_\phi < +\infty$. Indeed, observe that for any $\psi \in \mathbb{R}^N$, the mapping $x \mapsto \psi^{c,\varepsilon}(x)$ is \mathcal{C}^1 and has for gradient

$$\nabla_x [\psi^{c,\varepsilon}(x)] = \sum_{i=1}^N y_i \gamma_x^\varepsilon(\psi)_i \in \mathbb{R}^d.$$

This ensures that $\psi^{c,\varepsilon}$ is a $R_{\mathcal{Y}}$ -Lipschitz function on \mathcal{X} , where $R_{\mathcal{Y}}$ is such that $\mathcal{Y} \subset B(0, R_{\mathcal{Y}})$. The mapping $x \mapsto \psi^{c,\varepsilon}(x)$ is thus bounded on the compact set \mathcal{X} and satisfies

$$M_\phi - m_\phi \leq R_{\mathcal{Y}} \text{diam}(\mathcal{X}).$$

Note that estimate (4.10) without constants involving the geometry of \mathcal{Y} will be useful in Section 4.5 to retrieve the non-entropic and non-compact strong convexity result of Theorem 3.1.

The strong convexity estimate presented in Theorem 4.4 is reminiscent of the one presented in Theorem 3.1 for the classical Kantorovich functional. We recall that the main tool to derive this estimate was the Brascamp-Lieb inequality (Theorem 1.29), which almost readily gave a lower-bound on the second derivative of the classical Kantorovich functional. In our semi-discrete entropic setting, the Hessian of the Kantorovich functional is also available (4.8), but this second derivative does not look like the second term in the Brascamp-Lieb inequality. However, we saw in Section 1.3 that it is possible to

retrieve the Brascamp-Lieb inequality from the Prékopa-Leindler inequality. In particular, the Brascamp-Lieb inequality may be deduced from the concavity of a well-chosen functional I , the concavity of I being showed in turn with the Prékopa-Leindler inequality (Exercise 2.2.11 of (Klartag, 2013)). Here, we derive the estimate of Theorem 4.4 in a similar fashion. Let's first define the functional I , show its concavity thanks to the Prékopa-Leindler inequality and compute its Hessian.

Proposition 4.6. *The functional $I : \mathbb{R}^N \rightarrow \mathbb{R}, \boldsymbol{\psi} \mapsto \log(\int_{\mathcal{X}} e^{-\boldsymbol{\psi}^{c,\varepsilon}})$ is \mathcal{C}^2 and concave. In particular, its Hessian is negative semi-definite:*

$$\begin{aligned} \nabla^2 I(\boldsymbol{\psi}) &= \mathbb{E}_{x \sim \rho_{\boldsymbol{\psi}}^{\varepsilon}} \gamma_x^{\varepsilon}(\boldsymbol{\psi}) \gamma_x^{\varepsilon}(\boldsymbol{\psi})^{\top} - \mathbb{E}_{x \sim \rho_{\boldsymbol{\psi}}^{\varepsilon}} \gamma_x^{\varepsilon}(\boldsymbol{\psi}) \mathbb{E}_{x \sim \rho_{\boldsymbol{\psi}}^{\varepsilon}} \gamma_x^{\varepsilon}(\boldsymbol{\psi})^{\top} \\ &\quad - \frac{1}{\varepsilon} \mathbb{E}_{x \sim \rho_{\boldsymbol{\psi}}^{\varepsilon}} (\text{diag}(\gamma_x^{\varepsilon}(\boldsymbol{\psi})) - \gamma_x^{\varepsilon}(\boldsymbol{\psi}) \gamma_x^{\varepsilon}(\boldsymbol{\psi})^{\top}) \preceq 0, \end{aligned}$$

where $\rho_{\boldsymbol{\psi}}^{\varepsilon} := \frac{e^{-(\boldsymbol{\psi})^{c,\varepsilon}}}{\int_{\mathcal{X}} e^{-(\boldsymbol{\psi})^{c,\varepsilon}}}$.

Proof. By Lemma 4.2, we know that for any $x \in \mathcal{X}$ the mapping $\boldsymbol{\psi} \mapsto \boldsymbol{\psi}^{c,\varepsilon}(x)$ is \mathcal{C}^2 . This ensures that $I : \boldsymbol{\psi} \mapsto \log(\int_{\mathcal{X}} e^{-\boldsymbol{\psi}^{c,\varepsilon}})$ is a \mathcal{C}^2 function on \mathbb{R}^N . Its derivatives read:

$$\nabla I(\boldsymbol{\psi}) = \frac{-\int_{\mathcal{X}} \nabla_{\boldsymbol{\psi}} [\boldsymbol{\psi}^{c,\varepsilon}(x)] e^{-\boldsymbol{\psi}^{c,\varepsilon}(x)}}{\int_{\mathcal{X}} e^{-\boldsymbol{\psi}^{c,\varepsilon}}},$$

and

$$\begin{aligned} \nabla^2 I(\boldsymbol{\psi}) &= \frac{-\int_{\mathcal{X}} \nabla_{\boldsymbol{\psi}}^2 [\boldsymbol{\psi}^{c,\varepsilon}(x)] e^{-\boldsymbol{\psi}^{c,\varepsilon}(x)}}{\int_{\mathcal{X}} e^{-\boldsymbol{\psi}^{c,\varepsilon}}} + \frac{\int_{\mathcal{X}} \nabla_{\boldsymbol{\psi}} [\boldsymbol{\psi}^{c,\varepsilon}(x)] \nabla_{\boldsymbol{\psi}} [\boldsymbol{\psi}^{c,\varepsilon}(x)]^{\top} e^{-\boldsymbol{\psi}^{c,\varepsilon}(x)}}{\int_{\mathcal{X}} e^{-\boldsymbol{\psi}^{c,\varepsilon}}} \\ &\quad - \left(\frac{\int_{\mathcal{X}} \nabla_{\boldsymbol{\psi}} [\boldsymbol{\psi}^{c,\varepsilon}(x)] e^{-\boldsymbol{\psi}^{c,\varepsilon}(x)}}{\int_{\mathcal{X}} e^{-\boldsymbol{\psi}^{c,\varepsilon}}} \right) \left(\frac{\int_{\mathcal{X}} \nabla_{\boldsymbol{\psi}} [\boldsymbol{\psi}^{c,\varepsilon}(x)] e^{-\boldsymbol{\psi}^{c,\varepsilon}(x)}}{\int_{\mathcal{X}} e^{-\boldsymbol{\psi}^{c,\varepsilon}}} \right)^{\top}, \end{aligned}$$

which entails the claimed expression for $\nabla^2 I(\boldsymbol{\psi})$ using the formulas of Lemma 4.2.

We now show that I is a concave function on \mathbb{R}^N . Let $\boldsymbol{\psi}, \boldsymbol{\varphi} \in \mathbb{R}^N$. Let $0 < \lambda < 1$. Notice that for any $u, v \in \mathcal{X}$ we have:

$$\begin{aligned} (\lambda \boldsymbol{\psi} + (1-\lambda) \boldsymbol{\varphi})^{c,\varepsilon}(\lambda u + (1-\lambda)v) &= \varepsilon \log \left(\sum_{i=1}^N e^{\frac{\langle \lambda u + (1-\lambda)v | y_i \rangle - (\lambda \boldsymbol{\psi} + (1-\lambda) \boldsymbol{\varphi}) \cdot (y_i)}{\varepsilon}} \right) \\ &= \varepsilon \log \left(\sum_{i=1}^N \left(e^{\frac{\langle u | y_i \rangle - \boldsymbol{\psi} \cdot (y_i)}{\varepsilon}} \right)^{\lambda} \left(e^{\frac{\langle v | y_i \rangle - \boldsymbol{\varphi} \cdot (y_i)}{\varepsilon}} \right)^{1-\lambda} \right) \\ &\leq \varepsilon \log \left[\left(\sum_{i=1}^N e^{\frac{\langle u | y_i \rangle - \boldsymbol{\psi} \cdot (y_i)}{\varepsilon}} \right)^{\lambda} \left(\sum_{i=1}^N e^{\frac{\langle v | y_i \rangle - \boldsymbol{\varphi} \cdot (y_i)}{\varepsilon}} \right)^{1-\lambda} \right] \\ &= \lambda \boldsymbol{\psi}^{c,\varepsilon}(u) + (1-\lambda) \boldsymbol{\varphi}^{c,\varepsilon}(v), \end{aligned}$$

where the inequality corresponds to Hölder's inequality. Denoting

$$\begin{aligned} h(u) &= e^{-(\lambda \boldsymbol{\psi} + (1-\lambda) \boldsymbol{\varphi})^{c,\varepsilon}(u)}, \\ f(u) &= e^{-\boldsymbol{\psi}^{c,\varepsilon}(u)}, \\ g(u) &= e^{-\boldsymbol{\varphi}^{c,\varepsilon}(u)}, \end{aligned}$$

we thus have shown that

$$h(\lambda u + (1 - \lambda)v) \geq f(u)^\lambda g(v)^{1-\lambda}.$$

Using that \mathcal{X} is convex, the Prékopa–Leindler inequality (see Section 1.3 and in particular Theorem 1.28) then ensures that

$$\int_{\mathcal{X}} h \geq \left(\int_{\mathcal{X}} f \right)^\lambda \left(\int_{\mathcal{X}} g \right)^{1-\lambda}.$$

This leads to the concavity of I :

$$\begin{aligned} I(\lambda\psi + (1 - \lambda)\varphi) &= \log \left(\int_{\mathcal{X}} h \right) \geq \lambda \log \left(\int_{\mathcal{X}} f \right) + (1 - \lambda) \log \left(\int_{\mathcal{X}} g \right) \\ &= \lambda I(\psi) + (1 - \lambda) I(\varphi). \end{aligned} \quad \square$$

Proposition 4.6 implies Theorem 4.4. Indeed, in the expression of $\nabla^2 I(\psi^\varepsilon)$, one can notice that the last term *almost* corresponds to $-\nabla^2 \mathcal{K}_\rho^\varepsilon(\psi^\varepsilon)$, while the sum of the first and second terms *almost* corresponds to a p.s.d. matrix whose associated bilinear form corresponds to the covariance w.r.t. $\mu_{\psi^\varepsilon}^\varepsilon = \mathbb{E}_{x \sim \rho_{\psi^\varepsilon}^\varepsilon} \gamma_x^\varepsilon(\psi^\varepsilon)$. The difference with those terms resides in the presence of $\rho_{\psi^\varepsilon}^\varepsilon$ instead of ρ . The bounds on $\psi^{c,\varepsilon}$ allow to bound $\rho_{\psi^\varepsilon}^\varepsilon$ in terms of ρ , giving Theorem 4.4.

Proof of Theorem 4.4. Let $v \in \mathbb{R}^N$. Notice that

$$\begin{aligned} \langle v | \frac{1}{\varepsilon} \mathbb{E}_{x \sim \rho_{\psi^\varepsilon}^\varepsilon} \left(\text{diag}(\gamma_x^\varepsilon(\psi)) - \gamma_x^\varepsilon(\psi) \gamma_x^\varepsilon(\psi)^\top \right) v \rangle &= \int_{\mathcal{X}} \frac{1}{\varepsilon} \text{Var}_{\gamma_x^\varepsilon(\psi)}(v) d\rho_{\psi^\varepsilon}^\varepsilon(x), \\ \langle v | \mathbb{E}_{x \sim \rho_{\psi^\varepsilon}^\varepsilon} \gamma_x^\varepsilon(\psi) \gamma_x^\varepsilon(\psi)^\top v \rangle - \langle v | \mathbb{E}_{x \sim \rho_{\psi^\varepsilon}^\varepsilon} \gamma_x^\varepsilon(\psi) \mathbb{E}_{x \sim \rho_{\psi^\varepsilon}^\varepsilon} \gamma_x^\varepsilon(\psi)^\top v \rangle &= \text{Var}_{x \sim \rho_{\psi^\varepsilon}^\varepsilon} (\mathbb{E}_{\gamma_x^\varepsilon(\psi)}(v)). \end{aligned}$$

Thus Proposition 4.6 ensures that

$$\text{Var}_{x \sim \rho_{\psi^\varepsilon}^\varepsilon} (\mathbb{E}_{\gamma_x^\varepsilon(\psi)}(v)) \leq \int_{\mathcal{X}} \frac{1}{\varepsilon} \text{Var}_{\gamma_x^\varepsilon(\psi)}(v) d\rho_{\psi^\varepsilon}^\varepsilon(x), \quad (4.11)$$

where we recall

$$\rho_{\psi^\varepsilon}^\varepsilon = \frac{e^{-(\psi)^{c,\varepsilon}}}{\int_{\mathcal{X}} e^{-(\psi)^{c,\varepsilon}}} = \frac{e^{-(\psi)^{c,\varepsilon}}}{Z}.$$

The bound

$$-\infty < m_\phi \leq \psi^{c,\varepsilon} \leq M_\phi < +\infty$$

then gives the control

$$\frac{e^{-M_\phi}}{Z} \leq \rho_{\psi^\varepsilon}^\varepsilon \leq \frac{e^{-m_\phi}}{Z}.$$

Recalling that $m_\rho \leq \rho \leq M_\rho$ we thus have:

$$\frac{e^{-M_\phi}}{Z M_\rho} \rho \leq \rho_{\psi^\varepsilon}^\varepsilon \leq \frac{e^{-m_\phi}}{Z m_\rho} \rho.$$

This control, (4.8) and (4.11) thus give

$$\begin{aligned} \text{Var}_{x \sim \rho} (\mathbb{E}_{\gamma_x^\varepsilon(\psi)}(v)) &\leq e^{(M_\phi - m_\phi)} \frac{M_\rho}{m_\rho} \int_{\mathcal{X}} \frac{1}{\varepsilon} \text{Var}_{\gamma_x^\varepsilon(\psi)}(v) d\rho(x) \\ &= e^{(M_\phi - m_\phi)} \frac{M_\rho}{m_\rho} \langle v | \nabla^2 \mathcal{K}_\rho^\varepsilon(\psi) v \rangle. \end{aligned} \quad (4.12)$$

Recall the following expression of μ_{ψ}^{ε} :

$$\mu_{\psi}^{\varepsilon} = \mathbb{E}_{x \sim \rho} \gamma_x^{\varepsilon}(\psi).$$

Hence using the associativity of variances we have:

$$\text{Var}_{\mu_{\psi}^{\varepsilon}}(v) = \text{Var}_{x \sim \rho}(\mathbb{E}_{\gamma_x^{\varepsilon}(\psi)}(v)) + \int_{\mathcal{X}} \text{Var}_{\gamma_x^{\varepsilon}(\psi)}(v) d\rho(x),$$

so that using again (4.8),

$$\text{Var}_{x \sim \rho}(\mathbb{E}_{\gamma_x^{\varepsilon}(\psi)}(v)) = \text{Var}_{\mu_{\psi}^{\varepsilon}}(v) - \varepsilon \langle v | \nabla^2 \mathcal{K}_{\rho}^{\varepsilon}(\psi) v \rangle.$$

Injecting this last equality into (4.12) yields the desired result. \square

4.5 From the semi-discrete entropic case to the classical case

We finally mention briefly how the strong convexity estimate for the semi-discrete entropic Kantorovich functional of Theorem 4.4 may allow to recover the general strong-convexity estimate of Theorem 3.1 for the (classical) Kantorovich functional. This is simply the result of a series of three approximation arguments, allowing to go from the semi-discrete entropic case to the semi-discrete non-regularized case, then from the semi-discrete case to the compact case and finally from the compact case to the general case. Since we have already developed the last two approximation arguments in the preceding chapters (in Lemma 2.7 and Proposition 3.4 precisely), we only mention here how to go from the entropic semi-discrete estimate to the non-regularized semi-discrete estimate.

The transition from the entropic case to the non-regularized case can be done using a recent result from (Altschuler et al., 2022). This result ensures that the semi-discrete Schrödinger potentials for a regularization parameter ε converge, as ε goes to 0, to Kantorovich potentials and this convergence happens at a speed greater than ε . Note that we will quantify this rate of convergence in Chapter 7, relying in part on the strong convexity estimate of Theorem 4.4.

Proposition 4.7. *With the notation and assumptions of Theorem 3.1, assume additionally that the target measures μ^0, μ^1 are commonly supported on a finite set $\mathcal{Y} = \{y_1, \dots, y_N\} \subset \mathbb{R}^d$. Then (3.1) holds.*

Proof. For $\varepsilon > 0$, denote $\psi^{\varepsilon,0}$ and $\psi^{\varepsilon,1}$ minimizers of (4.6) with target measures μ^0 and μ^1 respectively. For $k \in \{0, 1\}$, denote $\phi^{\varepsilon,k}$ the (c, ε) -transform of $\psi^{\varepsilon,k}$ and assume that $\psi^{\varepsilon,k}$ is chosen such that $\langle \phi^{\varepsilon,k} | \rho \rangle = \langle \phi^k | \rho \rangle$. For $v^{\varepsilon} = \psi^{\varepsilon,1} - \psi^{\varepsilon,0}$ and $t \in [0, 1]$, introduce $\psi^{\varepsilon,t} = \psi^{\varepsilon,0} + tv^{\varepsilon}$. Then the fundamental theorem of calculus ensures:

$$\int_0^1 \int_0^s \langle \nabla^2 \mathcal{K}_{\rho}^{\varepsilon}(\psi^{\varepsilon,t}) v^{\varepsilon} | v^{\varepsilon} \rangle dt ds = \mathcal{K}_{\rho}^{\varepsilon}(\psi^{\varepsilon,1}) - \mathcal{K}_{\rho}^{\varepsilon}(\psi^{\varepsilon,0}) - \langle \nabla \mathcal{K}_{\rho}^{\varepsilon}(\psi^{\varepsilon,0}) | \psi^{\varepsilon,1} - \psi^{\varepsilon,0} \rangle.$$

Denote $m_{\phi^{\varepsilon}} = \min_{k \in \{0,1\}} \min_{\mathcal{X}} \phi^{\varepsilon,k}$ and $M_{\phi^{\varepsilon}} = \max_{k \in \{0,1\}} \max_{\mathcal{X}} \phi^{\varepsilon,k}$. Theorem 4.4 together with the first order condition $\nabla \mathcal{K}_{\rho}^{\varepsilon}(\psi^{\varepsilon,0}) = -\mu^0$ then ensure the inequality

$$\left(e^{(M_{\phi^{\varepsilon}} - m_{\phi^{\varepsilon}}) \frac{M_{\rho}}{m_{\rho}}} + \varepsilon \right)^{-1} \int_0^1 \int_0^s \text{Var}_{\mu_{\psi^{\varepsilon,t}}^{\varepsilon}}(v^{\varepsilon}) dt ds \leq \langle \phi^{\varepsilon,1} - \phi^{\varepsilon,0} | \rho \rangle + \langle \mu^0 | \psi^{\varepsilon,1} - \psi^{\varepsilon,0} \rangle. \quad (4.13)$$

Since ρ, μ^0 and μ^1 are compactly supported, the Arzelà-Ascoli theorem ensures that for $k \in \{0, 1\}$, $\phi^{\varepsilon, k}$ converges uniformly to ϕ^k and $\psi^{\varepsilon, k}$ converges to ψ^k as ε goes to 0 (where $\psi^k = (\psi^k(y_i))_{1 \leq i \leq N}$) (see e.g. (Nutz and Wiesel, 2021)). This ensures the limits:

$$\begin{aligned} M_{\phi^\varepsilon} - m_{\phi^\varepsilon} &\xrightarrow{\varepsilon \rightarrow 0} M_\phi - m_\phi, \\ \langle \phi^{\varepsilon, 1} - \phi^{\varepsilon, 0} | \rho \rangle &\xrightarrow{\varepsilon \rightarrow 0} \langle \phi^1 - \phi^0 | \rho \rangle, \\ \langle \mu^0 | \psi^{\varepsilon, 1} - \psi^{\varepsilon, 0} \rangle &\xrightarrow{\varepsilon \rightarrow 0} \langle \mu^0 | \psi^1 - \psi^0 \rangle. \end{aligned}$$

There now only remains to handle the limit of the measure $\mu_{\psi^{\varepsilon, t}}^\varepsilon$ as $\varepsilon \rightarrow 0$. Notice that by definition, for any $i \in \{1, \dots, N\}$ and $t \in [0, 1]$,

$$\mu_{\psi^{\varepsilon, t}}^\varepsilon(y_i) = \int_{\mathcal{X}} e^{\frac{\langle x | y_i \rangle - \psi_i^{\varepsilon, t} - \phi^{\varepsilon, t}(x)}{\varepsilon}} d\rho(x) \geq \int_{\text{Lag}_i(\psi^t)} e^{\frac{\langle x | y_i \rangle - \psi_i^{\varepsilon, t} - \phi^{\varepsilon, t}(x)}{\varepsilon}} d\rho(x),$$

where $\psi^t = (1-t)\psi^0 + t\psi^1$ and $\text{Lag}_i(\psi^t)$ denotes the i -th Laguerre cell of ψ^t , defined in Section 2.2. Using the inequality $e^x \geq 1+x$ then gives the bound

$$\mu_{\psi^{\varepsilon, t}}^\varepsilon(y_i) \geq \rho(\text{Lag}_i(\psi^t)) + \frac{1}{\varepsilon} \int_{\text{Lag}_i(\psi^t)} (\langle x | y_i \rangle - \psi_i^{\varepsilon, t} - \phi^{\varepsilon, t}(x)) d\rho(x).$$

Notice that for $x \in \text{Lag}_i(\psi^t)$, one has $\langle x | y_i \rangle = \phi^t(x) + \psi_i^t$ where $\phi^t = ((1-t)\psi^0 + t\psi^1)^*$. This leads to

$$\frac{1}{\varepsilon} \int_{\text{Lag}_i(\psi^t)} (\langle x | y_i \rangle - \psi_i^{\varepsilon, t} - \phi^{\varepsilon, t}(x)) d\rho(x) = \frac{1}{\varepsilon} \int_{\text{Lag}_i(\psi^t)} (\phi^t(x) - \phi^{\varepsilon, t}(x) + \psi_i^t - \psi_i^{\varepsilon, t}) d\rho(x)$$

Let's show that the limit of this quantity is 0. Theorem 1.3 of (Altschuler et al., 2022) ensures the limits

$$\frac{1}{\varepsilon} \|\psi^t - \psi^{\varepsilon, t}\|_\infty \xrightarrow{\varepsilon \rightarrow 0} 0 \quad \text{and} \quad \frac{1}{\varepsilon} |\phi^t - \phi^{\varepsilon, t}| \xrightarrow{\varepsilon \rightarrow 0} 0,$$

where the last limit is point-wise. In order to apply the dominated convergence theorem, there only remains to show that $\frac{1}{\varepsilon} |\phi^t - \phi^{\varepsilon, t}|$ is bounded by an integrable function on $\text{Lag}_i(\psi^t)$: for all $x \in \text{Lag}_i(\psi^t)$,

$$\begin{aligned} \frac{1}{\varepsilon} |\phi^t(x) - \phi^{\varepsilon, t}(x)| &\leq \frac{1}{\varepsilon} |\phi^t(x) - (\psi^t)^{c, \varepsilon}(x)| + \frac{1}{\varepsilon} |(\psi^t)^{c, \varepsilon}(x) - \phi^{\varepsilon, t}(x)| \\ &= \left| \log \left(\sum_i e^{\frac{\langle x | y_i \rangle - \psi_i^t - \phi^t(x)}{\varepsilon}} \right) \right| + \left| \log \left(\frac{\sum_i e^{\frac{\langle x | y_i \rangle - \psi_i^t}{\varepsilon}}}{\sum_i e^{\frac{\langle x | y_i \rangle - \psi_i^{\varepsilon, t}}{\varepsilon}}} \right) \right| \\ &\leq \log N + \frac{1}{\varepsilon} \|\psi^t - \psi^{\varepsilon, t}\|_\infty. \end{aligned}$$

We can thus apply the dominated convergence theorem and give the limit

$$\frac{1}{\varepsilon} \int_{\text{Lag}_i(\psi^t)} (\langle x | y_i \rangle - \psi_i^{\varepsilon, t} - \phi^{\varepsilon, t}(x)) d\rho(x) \xrightarrow{\varepsilon \rightarrow 0} 0.$$

We thus has proven that for all $i \in \{1, \dots, N\}$,

$$\lim_{\varepsilon \rightarrow 0} \mu_{\psi^{\varepsilon, t}}^\varepsilon(y_i) \geq \rho(\text{Lag}_i(\psi^t)) := \mu^t(y_i).$$

Taking the limit $\varepsilon \rightarrow 0$ of (4.13) and using the comparison $\mu^t \geq \min(1-t, t) \frac{d m_\rho}{M_\rho} (\mu^0 + \mu^1)$ of Lemma 2.3 then gives the estimate:

$$\frac{1}{C_d e^{(M_\phi - m_\phi)}} \frac{m_\rho^2}{M_\rho^2} \text{Var}_{\frac{1}{2}(\mu^0 + \mu^1)}(\psi^1 - \psi^0) \leq \mathcal{K}_\rho(\psi^1) - \mathcal{K}_\rho(\psi^0) + \langle \psi^1 - \psi^0 | \mu^0 \rangle,$$

where $C_d = (d+1)2^{d-1}$. Finally, using the same homogeneity argument as in the end of the proof of Proposition 3.3, we know that we can replace $e^{(M_\phi - m_\phi)}$ with $e(M_\phi - m_\phi)$ in the last inequality, which gives (3.1). \square

Part II

Consequences for the stability of solutions to optimal transport problems

Quantitative stability of optimal transport maps with respect to the target measure

Abstract

This chapter, extracted from (Delalande and Mérigot, 2021), studies the quantitative stability of the quadratic optimal transport map between a fixed probability density ρ and a probability measure μ on \mathbb{R}^d , which we denote T_μ . Assuming that the source density ρ is bounded from above and below on a compact convex set, we prove that the map $\mu \mapsto T_\mu$ is bi-Hölder continuous w.r.t. the 2-Wasserstein metric on large families of probability measures, such as the set of probability measures whose moment of order $p > d$ is bounded by some constant. These stability estimates show that the *linearized optimal transport* metric $W_{2,\rho}(\mu, \nu) = \|T_\mu - T_\nu\|_{L^2(\rho, \mathbb{R}^d)}$ is bi-Hölder equivalent to the 2-Wasserstein distance on such sets, justifying its use in applications.

5.1 Introduction

In Part I, we have introduced the quadratic optimal problem between ρ and μ , which corresponds to the following minimization problem:

$$\min_{\gamma \in \Gamma(\rho, \mu)} \int_{\mathbb{R}^d \times \mathbb{R}^d} \|x - y\|^2 d\gamma(x, y).$$

The square root of the value of this problem is called the *2-Wasserstein distance* between ρ and μ and is denoted $W_2(\rho, \mu)$ (see Chapter A). Brenier's theorem (Brenier, 1991) (reported in Theorem 1.12) asserts that if ρ is absolutely continuous with respect to the Lebesgue measure, the minimizer of the optimal transport problem is unique, and is induced by a map $T = \nabla\phi$, where ϕ is a convex function that verifies $\nabla\phi_{\#}\rho = \mu$. We recall that $T_{\#}\rho$ denotes the image measure of ρ under the map T . In our precise setting, where the density ρ is bounded from above and below on a compact convex set, the potential ϕ is uniquely defined in $L^2(\rho)$ up to an additional constant (see Remark 1.13 and note that square-summability of ϕ follows from the Poincaré-Wirtinger inequality on \mathcal{X}).

Definition 5.1 (Potentials and maps). We fix a probability measure $\rho \in \mathcal{P}_2(\mathbb{R}^d)$, which we assume to be absolutely continuous with respect to the Lebesgue measure and supported over a compact convex set \mathcal{X} . We assume that the density of ρ is bounded from above and below by positive constants on \mathcal{X} . Given $\mu \in \mathcal{P}_2(\mathbb{R}^d)$, we call

- *Monge or Brenier map* and denote T_μ the (unique) optimal transport map between ρ and μ ;
- *Brenier potential* the unique lower semi-continuous convex function $\phi_\mu \in L^2(\mathcal{X})$ such that $T_\mu = \nabla \phi_\mu$ and which satisfies $\int_{\mathcal{X}} \phi_\mu d\rho = 0$;
- *dual potential* the convex conjugate of ϕ_μ , denoted ψ_μ :

$$\forall y \in \mathbb{R}^d, \quad \psi_\mu(y) = \max_{x \in \mathcal{X}} \langle x|y \rangle - \phi_\mu(x),$$

where the maximum is attained by lower semi-continuity of the convex function ϕ_μ on the compact convex set \mathcal{X} .

Since μ is the image of ρ under T_μ , the mapping

$$\begin{cases} (\mathcal{P}_2(\mathbb{R}^d), W_2) & \rightarrow L^2(\rho, \mathbb{R}^d), \\ \mu & \mapsto T_\mu, \end{cases}$$

is obviously injective. Using that $(T_\mu, T_\nu)_\# \rho$ is a coupling between μ and ν , one can actually prove that this mapping increases distances, namely

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

This mapping is also continuous: if a sequence of probability measures $(\mu_n)_n$ converges to some μ in $(\mathcal{P}_2(\mathbb{R}^d), W_2)$, then T_{μ_n} converges to T_μ in $L^2(\rho, \mathbb{R}^d)$. This continuity property is for instance implied by Corollary 5.21 in (Villani, 2008) (reported in Corollary A.11), together with the dominated convergence theorem. However, we note that the arguments used to prove this general continuity result are non-quantitative.

5.1.1 Linearized Optimal Transport

The continuous and reverse-Lipschitz behavior of the map $\mu \mapsto T_\mu$ motivated its use to embed the metric space $(\mathcal{P}_2(\mathbb{R}^d), W_2)$ into the Hilbert space $L^2(\rho, \mathbb{R}^d)$ (Wang et al., 2013). This approach is often referred to as the *Linearized Optimal Transport* (LOT) framework and has shown great results in applications to image processing:

- (Wang et al., 2013; Kolouri et al., 2016; Basu et al., 2014; Cai et al., 2020) used this idea to perform pattern recognition in images for various tasks, including discrimination of nuclear chromatin patterns in cancer cells; detection of differences in facial expressions, bird species, galaxy morphologies, sub-cellular protein distributions; detection and visualization of cell phenotype differences from microscopy images; or finally jets tagging of collider data in collider physics.
- (Park and Thorpe, 2018) considered this framework for generative modelling of images, with experiments showcasing the generative modelling of digits and faces images, PET scans in the context of Alzheimer's disease neuroimaging, or thyroid nuclei images.

- (Kolouri and Rohde, 2015) followed this approach for improving the resolution of faces images.

At this stage, the good practical behavior of the linearized optimal transport framework is not justified from a mathematical viewpoint. A practical benefit of the embedding is to enable the use of the classical Hilbertian statistical toolbox on families of probability measures while keeping some features of the Wasserstein geometry. A particularly nice feature of the embedding $\mu \mapsto T_\mu$ is that its image in $L^2(\rho, \mathbb{R}^d)$ is convex, i.e. barycenters of optimal transport maps are optimal transport maps. Working with this embedding is equivalent to replacing the Wasserstein distance by the distance

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

We note that the geodesic curves with respect to the distance $W_{2,\rho}$ are called the *generalized geodesics* in the book of Ambrosio, Gigli, Savaré (Ambrosio et al., 2008). The choice of the Brenier map between a reference measure ρ and a measure μ as an embedding of μ may also be motivated by the Riemannian interpretation of the Wasserstein geometry (Otto, 2001; Ambrosio et al., 2008). In this interpretation, the tangent space to $(\mathcal{P}_2(\mathbb{R}^d), W_2)$ at ρ is included in $L^2(\rho, \mathbb{R}^d)$ and defined by

$$\mathcal{T}_\rho \mathcal{P}_2(\mathbb{R}^d) = \overline{\{\lambda(\nabla\phi - \text{id}) \mid \lambda > 0, \phi \in \mathcal{C}_c^\infty(\mathbb{R}^d), \phi \text{ convex}\}}^{L^2(\rho, \mathbb{R}^d)}.$$

The Brenier map minus the identity, $T_\mu - \text{id}$, can thus be regarded as the vector in the tangent space at ρ which supports the Wasserstein geodesic from ρ to μ . In the Riemannian language again, the map $\mu \mapsto T_\mu - \text{id}$ would be called a *logarithm*, i.e. the inverse of the Riemannian exponential map: it sends a probability measure μ in the (curved) manifold $\mathcal{P}_2(\mathbb{R}^d)$ to a vector $T_\mu - \text{id}$ belonging to the linear space $L^2(\rho, \mathbb{R}^d)$. This establishes a connection between the linearized optimal transport framework idea and similar strategies used to extend statistical inference notions such as principal component analysis to manifold-valued data, e.g. (Fletcher et al., 2004; Cazelles et al., 2018). Finally, we note that the work in (Chernozhukov et al., 2017) also proposes to use optimal transport maps in a statistical context to overcome the lack of a canonical ordering in \mathbb{R}^d for $d > 1$. Notions of vector-quantile, vector-ranks and depth are defined based on the transport maps (and their inverses) between a reference measure defined as the uniform distribution on the unit hyperball and the d -dimensional samples of interest.

It is quite natural to expect that the embedding $\mu \mapsto T_\mu$ retains some of the geometry of the underlying space, or equivalently that the metric $W_{2,\rho}$ is comparable, in some coarse sense, to the Wasserstein distance. The main difficulty, which we study in this chapter, is to establish quantitative (e.g. Hölder) continuity properties for the mappings $\mu \mapsto T_\mu$ and $\mu \mapsto \phi_\mu$. We note that such stability estimates are also important in numerical analysis and in statistics, where a probability measure of interest $\mu \in \mathcal{P}_2(\mathbb{R}^d)$ is often approximated by a sequence of finitely supported measures $(\mu_n)_n$. Consider for instance the problem of computing the map T_μ for some $\mu \in \mathcal{P}_2(\mathbb{R}^d)$: a natural approximation consists in computing instead the map T_{μ_n} for some sequence of measures $(\mu_n)_n$ with finite supports such that $\lim_{n \rightarrow +\infty} W_2(\mu, \mu_n) = 0$. This corresponds to the *semi-discrete approach*, which can be traced back to Minkowski and Alexandrov and was developed in many works from the 1990s (Cullen et al., 1991; Gangbo and McCann, 1996; Olikar and Prussner, 1989; Caffarelli et al., 1999). This approach was revisited and popularized by recent development of efficient algorithms to solve semi-discrete optimal

transport problems (Aurenhammer et al., 1998; Mérigot, 2011; De Goes et al., 2012; Lévy, 2015; Kitagawa et al., 2019; Genevay et al., 2016). In this approach, convergence rates of quantities related to the sequence $(T_{\mu_n})_n$ toward a quantity related to T_μ may be directly deduced from quantitative stability estimates controlling $\|T_{\mu_n} - T_\mu\|_{L^2(\rho, \mathbb{R}^d)}$ with $W_2(\mu_n, \mu)$.

5.1.2 Existing results

We focus here on the already known stability results on the mapping $\mu \mapsto T_\mu$, starting with negative results.

Negative results. We first note that explicit examples show that the mapping $\mu \mapsto T_\mu$ is in general not better than $\frac{1}{2}$ -Hölder, see §4 in (Gigli, 2011) or the following lemma (Lemma 5.1 in (Mérigot et al., 2020)):

Lemma 5.2. *Let ρ be uniform on the unit disc $\mathcal{X} \subseteq \mathbb{R}^2$. Then, there is a curve $\theta \in [0, 2\pi] \rightarrow \mu_\theta \in \mathcal{P}(\mathcal{X})$ such that $\|T_{\mu_\theta} - T_{\mu_0}\|_{L^2(\rho, \mathbb{R}^d)} \geq W_2(\mu_\theta, \mu_0)^{1/2}$.*

Proof. Given $\theta \in \mathbb{R}$, we denote $x_\theta = (\cos \theta, \sin \theta)$ and $\mu_\theta = \frac{1}{2}(\delta_{x_\theta} + \delta_{-x_\theta})$. Then, the optimal transport map between ρ and μ_θ is given by

$$T_{\mu_\theta}(x) = \begin{cases} x_\theta & \text{if } \langle x | x_\theta \rangle \geq 0 \\ -x_\theta & \text{if not.} \end{cases}$$

One can easily check that for $|\theta| \leq \frac{\pi}{2}$ one has $W_2(\mu_0, \mu_\theta) \leq |\theta|$. For $\theta > 0$ we set

$$D_\theta = \{x \in \mathbb{R}^2 \mid \langle x | x_0 \rangle \geq 0 \text{ and } \langle x | x_\theta \rangle \leq 0\}.$$

Then, on D_θ , $T_{\mu_\theta} \equiv x_{-\theta}$ and $T_{\mu_0} \equiv x_0$, giving

$$\|T_{\mu_\theta} - T_{\mu_0}\|_{L^2(\rho, \mathbb{R}^d)}^2 \geq \int_{D_\theta} \|x_{-\theta} - x_0\|^2 dx = |D_\theta| \|x_{-\theta} - x_0\|^2.$$

Moreover, if $|\theta| \leq \frac{\pi}{2}$ one has $\|x_{-\theta} - x_0\|^2 \geq 2$. This gives

$$\|T_{\mu_\theta} - T_{\mu_0}\|_{L^2(\rho, \mathbb{R}^d)}^2 \geq 2|D_\theta| = 2\frac{|\theta|}{2\pi}|\mathcal{X}| = |\theta|. \quad \square$$

A much stronger negative result comes from Andoni, Naor and Neiman (Andoni et al., 2018, Theorem 7) showing that one cannot construct a bi-Hölder embedding of $(\mathcal{P}_2(\mathbb{R}^d), W_2)$, $d \geq 3$, into a Hilbert space:

Theorem (Andoni, Naor, Neiman). *$(\mathcal{P}_2(\mathbb{R}^3), W_2)$ does not admit a uniform, coarse or quasimetric embedding into any Banach space of nontrivial type.*

This theorem implies in particular that one cannot hope to prove that $\mu \mapsto T_\mu$ is bi-Hölder on the whole set $\mathcal{P}_2(\mathbb{R}^d)$ of probability measures with finite second moment.

Positive results. Existing quantitative stability results can be summed up under the two following statements. A first result due to Ambrosio and reported in (Gigli, 2011), shows a local $1/2$ -Hölder behaviour near probability densities μ whose associated Brenier map T_μ is Lipschitz continuous. We have already noticed in Part I that such setting induces a strong convexity estimate for the Kantorovich functional when evaluated *near* ψ_μ (Proposition 1.20). This strong convexity estimate directly entails the result of (Gigli, 2011). We quote here a variant of this statement, from (Mérigot et al., 2020):

Theorem (Ambrosio). *Let ρ be a probability density over a compact set \mathcal{X} . Let $\mathcal{Y} \subset \mathbb{R}^d$ be a compact set and $\mu, \nu \in \mathcal{P}(\mathcal{Y})$. Assume that the Brenier map T_μ from ρ to μ is L -Lipschitz. Then,*

$$\|T_\mu - T_\nu\|_{L^2(\rho, \mathbb{R}^d)} \leq 2\sqrt{\text{diam}(\mathcal{X})LW_1(\mu, \nu)^{1/2}}.$$

Proof. Assuming that $T_\mu = \nabla\phi_\mu$ is L -Lipschitz continuous is equivalent to assume that $\psi_\mu = (\phi_\mu)^*$ is $\frac{1}{L}$ -strongly convex (Remark 1.16). We know in these conditions from Proposition 1.20 that \mathcal{K}_ρ satisfies the following strong convexity estimate at ψ_ν near ψ_μ :

$$\frac{1}{2L} \|T_\mu - T_\nu\|_{L^2(\rho, \mathbb{R}^d)}^2 \leq \mathcal{K}_\rho(\psi_\mu) - \mathcal{K}_\rho(\psi_\nu) + \langle \psi_\mu - \psi_\nu | \nu \rangle.$$

We also know by Lemma 1.8 that $-\mu$ belongs to the subdifferential of \mathcal{K}_ρ at ψ_μ , that is

$$0 \leq \mathcal{K}_\rho(\psi_\nu) - \mathcal{K}_\rho(\psi_\mu) + \langle \psi_\nu - \psi_\mu | \mu \rangle.$$

Summing the last two inequalities thus yields

$$\frac{1}{2L} \|T_\mu - T_\nu\|_{L^2(\rho, \mathbb{R}^d)}^2 \leq \langle \psi_\mu - \psi_\nu | \nu - \mu \rangle.$$

Using that ψ_μ and ψ_ν are the convex conjugates of ϕ_μ and ϕ_ν respectively, both defined on \mathcal{X} , one can show that $\psi_\mu - \psi_\nu$ is $\text{diam}(\mathcal{X})$ -Lipschitz continuous. This ensures, together with the Kantorovich-Rubinstein duality formula (Proposition A.8), the bound

$$\frac{1}{2L} \|T_\mu - T_\nu\|_{L^2(\rho, \mathbb{R}^d)}^2 \leq \text{diam}(\mathcal{X})W_1(\mu, \nu). \quad \square$$

As already noticed after the statement of Proposition 1.20, assuming Lipschitzness of the Brenier map is rather strong: it requires at least that the support of μ is connected – so that the previous theorem cannot be applied when both μ and ν are finitely supported – and in order to prove that T_μ is Lipschitz, one has to invoke the regularity theory for optimal transport maps, which requires very strong assumptions on μ . A more recent result, due to Berman (Berman, 2020), proves quantitative stability of the map $\mu \mapsto T_\mu$ under milder assumptions on the target probability measures. Berman assumes that the source measure ρ is the restriction of the Lebesgue measure to a compact convex set \mathcal{X} with unit volume. Under this assumption, he proves a stability result on the inverse transport maps when the target measure is required to remain in a fixed compact set (Berman, 2020, Proposition 3.2). This result implies the following quantitative stability of the Brenier maps (first reported in (Mérigot et al., 2020)):

Theorem (Berman). *Let \mathcal{X} be a compact convex subset of \mathbb{R}^d , let $\rho \in \mathcal{P}_{a.c.}(\mathcal{X})$ with density bounded from above and below by positive constants. Let \mathcal{Y} be a bounded connected open subset of \mathbb{R}^d with a Lipschitz boundary. Then there exists a constant $C_{d,\rho,\mathcal{X},\mathcal{Y}}$ depending only on d, ρ, \mathcal{X} and \mathcal{Y} such that for any $\mu, \nu \in \mathcal{P}(\mathcal{Y})$,*

$$\|T_\mu - T_\nu\|_{L^2(\rho, \mathbb{R}^d)} \leq C_{d,\rho,\mathcal{X},\mathcal{Y}} W_1(\mu, \nu)^{\frac{1}{2^{(d-1)}(d+2)}}.$$

Proof. Let $\mu, \nu \in \mathcal{P}(\mathcal{Y})$. Proposition 3.2 of (Berman, 2020) asserts that there exists $C_{\rho, \mathcal{X}, \mathcal{Y}}$ depending only on ρ , \mathcal{X} and \mathcal{Y} such that

$$\|\nabla\psi_\mu - \nabla\psi_\nu\|_{L^2(\mathcal{Y}, \mathbb{R}^d)}^2 \leq C_{\rho, \mathcal{X}, \mathcal{Y}} \left(\int_{\mathcal{Y}} (\psi_\nu - \psi_\mu) d(\mu - \nu) \right)^{\frac{1}{2d-1}}.$$

Because $\psi_\mu = \phi_\mu^*$ and $\psi_\nu = \phi_\nu^*$, $\partial\psi_\mu$ and $\partial\psi_\nu$ are valued in \mathcal{X} (Remarks 1.6 and 1.7). The difference $\psi_\nu - \psi_\mu$ is thus $\text{diam}(\mathcal{X})$ -Lipschitz continuous, so that using the Kantorovich-Rubinstein formula (Proposition A.8) together with the Poincaré-Wirtinger inequality on the compact set \mathcal{Y} in the last inequality, there exists $C_{\rho, \mathcal{X}, \mathcal{Y}}$ depending only on ρ , \mathcal{X} and \mathcal{Y} such that

$$\|\psi_\mu - \psi_\nu\|_{L^2(\mathcal{Y})}^2 \leq C_{\rho, \mathcal{X}, \mathcal{Y}} W_1(\mu, \nu)^{\frac{1}{2d-1}}, \quad (5.1)$$

where we assumed for simplicity that $\int_{\mathcal{Y}} \psi_\mu dy = \int_{\mathcal{Y}} \psi_\nu dy$. Again, because $\psi_\nu - \psi_\mu$ is $\text{diam}(\mathcal{X})$ -Lipschitz continuous on the Lipschitz domain \mathcal{Y} , one can show that there exists a constant $C_{d, \mathcal{X}, \mathcal{Y}} > 0$ that depends on the dimension d and the domains \mathcal{X} and \mathcal{Y} such that

$$\|\psi_\mu - \psi_\nu\|_{L^\infty(\mathcal{Y})} \leq C_{d, \mathcal{X}, \mathcal{Y}} \|\psi_\mu - \psi_\nu\|_{L^2(\mathcal{Y})}^{\frac{2}{d+2}}. \quad (5.2)$$

Then using that $\phi_\mu = \psi_\mu^*$ and $\phi_\nu = \psi_\nu^*$, one easily has

$$\|\phi_\mu - \phi_\nu\|_{L^\infty(\mathcal{X})} \leq \|\psi_\mu - \psi_\nu\|_{L^\infty(\mathcal{Y})}. \quad (5.3)$$

We finally quote Theorem 22 of (Chazal et al., 2017), that ensures that there exists a constant $C_{\mathcal{X}}$ depending only on \mathcal{X} such that for any f and g convex functions on \mathcal{X} ,

$$\|\nabla f - \nabla g\|_{L^2(\mathcal{X}, \mathbb{R}^d)} \leq C_{\mathcal{X}} \|f - g\|_{L^\infty(\mathcal{X})}^{1/2} (\|\nabla f\|_{L^\infty(\mathcal{X}, \mathbb{R}^d)}^{1/2} + \|\nabla g\|_{L^\infty(\mathcal{X}, \mathbb{R}^d)}^{1/2}).$$

Setting $f = \phi_\mu$ and $g = \phi_\nu$ in this inequality and using that $\nabla\phi_\mu, \nabla\phi_\nu$ take their image in \mathcal{Y} then yields that there exists $C_{\mathcal{X}, \mathcal{Y}}$ depending only on \mathcal{X} and \mathcal{Y} such that

$$\|\nabla\phi_\mu - \nabla\phi_\nu\|_{L^2(\mathcal{X}, \mathbb{R}^d)} \leq C_{\mathcal{X}, \mathcal{Y}} \|\phi_\mu - \phi_\nu\|_{L^\infty(\mathcal{X})}^{1/2}. \quad (5.4)$$

Therefore, using that ρ is bounded away from zero and infinity on \mathcal{X} and that $\nabla\phi_\mu = T_\mu$ and $\nabla\phi_\nu = T_\nu$, one has with inequalities (5.1), (5.2), (5.3) and (5.4) the existence of a constant $C_{d, \rho, \mathcal{X}, \mathcal{Y}}$ depending only on d , ρ , \mathcal{X} and \mathcal{Y} such that

$$\|T_\mu - T_\nu\|_{L^2(\rho, \mathbb{R})} \leq C_{d, \rho, \mathcal{X}, \mathcal{Y}} W_1(\mu, \nu)^{\frac{1}{2d-1(d+2)}}. \quad \square$$

Unlike in Ambrosio's theorem, the Hölder behavior deduced from Berman's stability result does not depend on the regularity of the transport map T_μ . On the other hand, the Hölder exponent depends exponentially on the ambient dimension d . As we will see below, this is not optimal.

5.1.3 Contributions

In this chapter, we prove quantitative stability results for quadratic optimal transport maps between a probability density ρ and target measure μ . We do not assume that μ is compactly supported. Introducing $M_p(\mu) = \int_{\mathbb{R}^d} \|x\|^p d\mu(x)$ the p -th moment of $\mu \in \mathcal{P}_2(\mathbb{R}^d)$, we prove in particular the following theorem. We denote by C_{a_1, \dots, a_n} a non-negative constant which depends on a_1, \dots, a_n .

Theorem (Theorems 5.14, 5.12, Corollary 5.8). *Let \mathcal{X} be a compact convex set and let ρ be a probability density on \mathcal{X} , bounded from above and below by positive constants. Let $p > d$ and $p \geq 4$. Assume that $\mu, \nu \in \mathcal{P}_2(\mathbb{R}^d)$ have bounded p -th moment, i.e. $\max(M_p(\mu), M_p(\nu)) \leq M_p < +\infty$. Then*

$$\|T_\mu - T_\nu\|_{L^2(\rho, \mathbb{R}^d)} \leq C_{d,p,\mathcal{X},\rho,M_p} W_1(\mu, \nu)^{\frac{p}{6p+16d}},$$

$$\|\phi_\mu - \phi_\nu\|_{L^2(\rho)} \leq C_{d,p,\mathcal{X},\rho,M_p} W_1(\mu, \nu)^{1/2}.$$

If μ, ν are supported on a compact set \mathcal{Y} , we have an improved Hölder exponent for the Brenier map:

$$\|T_\mu - T_\nu\|_{L^2(\rho, \mathbb{R}^d)} \leq C_{d,\mathcal{X},\mathcal{Y},\rho} W_1(\mu, \nu)^{\frac{1}{6}}.$$

Remark 5.3 (Comparison between W_1 and W_2). We note that since $W_1 \leq W_2$, the estimates in all the previous theorems indeed imply a bi-Hölder behaviour of the map $\mu \mapsto T_\mu$ on subsets of $\mathcal{P}_2(\mathbb{R}^d)$ with respect to both Wasserstein distances W_1 and W_2 .

Remark 5.4 (Linearized Optimal Transport). As an example application, we will show in Part III (Section 8.3) that our bi-Hölder embedding result can be used to analyze the behaviour of clustering algorithms in the linearized optimal transport framework.

Remark 5.5 (Constants). The constants appearing in the above theorem may all be tracked down and all feature the product of three terms that depend respectively on the dimension d , the diameter and perimeter of \mathcal{X} , and the bounds $m_\rho, M_\rho > 0$ on ρ that are such that $m_\rho \leq \rho \leq M_\rho$. If μ, ν are supported on a compact set \mathcal{Y} , the constants also feature a factor that only depends on the smallest positive real $R_\mathcal{Y}$ such that $\mathcal{Y} \subset B(0, R_\mathcal{Y})$. For instance in such compact setting, the constant controlling the $L^2(\rho)$ distance between ϕ_μ and ϕ_ν reads:

$$C_{d,p,\mathcal{X},\rho,M_p} = C_{d,p,\mathcal{X},\rho,\mathcal{Y}} = e(d+1)2^d \frac{M_\rho^2}{m_\rho^2} \text{diam}(\mathcal{X})^2 R_\mathcal{Y}.$$

In the non-compact setting, a factor involving M_p appears, as well as a factor involving the Poincaré constant of order p of \mathcal{X} and the p -th power of the ratio $\frac{R_\mathcal{X}}{r_\mathcal{X}}$, where $r_\mathcal{X}, R_\mathcal{X} > 0$ are the largest and smallest reals such that $B(0, r_\mathcal{X}) \subset \mathcal{X} \subset B(0, R_\mathcal{X})$ (assuming without any loss of generality \mathcal{X} contains the origin).

A large class of probability measures verifies the moment assumption, such as sub-Gaussian or sub-exponential measures (see Remark 5.9). A preliminary version of this theorem was announced in (Mérigot et al., 2020) (not reported in this thesis), with a different proof strategy, relying on the study of the case where both μ, ν are supported on the same finite set. The proof in (Mérigot et al., 2020) led to a worse Hölder exponent in the compact case, and couldn't deal with non-compactly supported measures. We do not know whether the Hölder exponents in this theorem are optimal.

Outline. To prove these stability estimates, we use the fact that the dual potentials solve a convex minimization problem involving the Kantorovich functional $\mathcal{K}(\psi) = \int \psi^* d\rho$ studied in Part I. We use the strong convexity estimate for the Kantorovich functional derived in Chapter 3, which holds under the assumption that the Brenier potentials are bounded, to give a stability estimate concerning the dual and Brenier potentials (§5.2). The stability of Brenier maps is then obtained (§5.3), relying in particular on a Gagliardo-Nirenberg type inequality for the difference of convex functions (§5.4), which might be of independent interest.

5.2 Stability of potentials

A direct consequence of the strong convexity estimate of the Kantorovich functional derived in Chapter 3 (Theorem 3.1) is a quantitative stability result on the dual potential ψ_μ with respect to the target measure μ . This estimate on dual potentials is readily transferred to the Brenier (primal) potentials thanks to Proposition 1.30 of Chapter 1, for which we recall a statement here:

Proposition (Proposition 1.30). *Let ρ be a probability density over a compact convex set \mathcal{X} , and let ϕ^0, ϕ^1 be convex functions on \mathcal{X} . Denote ψ^k the convex conjugate of ϕ^k and μ^k the image of ρ under $\nabla\phi^k$. Then for any $p > 0$,*

$$\|\phi^1 - \phi^0\|_{L^p(\rho)} \leq \|\psi^1 - \psi^0\|_{L^p(\mu^0 + \mu^1)}.$$

In particular,

$$\frac{1}{2}\text{Var}_\rho(\phi^1 - \phi^0) \leq \text{Var}_{\frac{1}{2}(\mu^0 + \mu^1)}(\psi^1 - \psi^0).$$

The stability estimates resulting from Theorem 3.1 and this proposition are expressed in Corollary 5.6 in terms of variance for the potentials and 1-Wasserstein distance for the target measures. Assuming that one of the target measures is absolutely continuous with respect to the other, these estimates can also be expressed in term of χ^2 or Kagan's divergence of the target measures. The χ^2 divergence reduces to the χ^2 test-statistic used for goodness of fit testing when the compared measures are finitely and commonly supported and one of them is observed empirically. Note that such divergence can be interpreted as the *square of a divergence*, noting for instance that the total variation distance is only $\frac{1}{2}$ -Hölder stable with respect to it (Peyré and Cuturi, 2019).

Corollary 5.6 (Stability of potentials). *Let ρ be a probability density over a compact convex set \mathcal{X} , satisfying $0 < m_\rho \leq \rho \leq M_\rho$ and let $\mu^0, \mu^1 \in \mathcal{P}_2(\mathbb{R}^d)$. For $k \in \{0, 1\}$, denote $\phi^k = \phi_{\mu^k}$ the Brenier potential between ρ and μ^k . Assume that*

$$\forall k \in \{0, 1\}, \quad -\infty < m_\phi \leq \min_{\mathcal{X}} \phi^k \leq \max_{\mathcal{X}} \phi^k \leq M_\phi < +\infty.$$

Denote ψ^0 and ψ^1 the convex conjugates of ϕ^0 and ϕ^1 . Then,

$$\text{Var}_\rho(\phi^1 - \phi^0) \leq 2\text{Var}_{\frac{1}{2}(\mu^0 + \mu^1)}(\psi^1 - \psi^0) \leq C_{d,\rho} \text{diam}(\mathcal{X})(M_\phi - m_\phi)W_1(\mu^0, \mu^1)$$

with $C_{d,\rho} = e(d+1)2^d \frac{M_\rho^2}{m_\rho^2}$. Assuming additionally that μ^1 is absolutely continuous w.r.t. μ^0 , then

$$\text{Var}_\rho(\phi^1 - \phi^0) \leq 2\text{Var}_{\frac{1}{2}(\mu^0 + \mu^1)}(\psi^1 - \psi^0) \leq C_{d,\rho}^2 (M_\phi - m_\phi)^2 D_{\chi^2}(\mu^1 | \mu^0)$$

where $D_{\chi^2}(\mu^1 | \mu^0)$ stands for the χ^2 or Kagan's divergence from μ^1 to μ^0 .

Proof. Proposition 1.30 combined with Theorem 3.1 give the inequalities

$$\begin{aligned} \text{Var}_\rho(\phi^1 - \phi^0) &\leq 2\text{Var}_{\frac{1}{2}(\mu^0 + \mu^1)}(\psi^1 - \psi^0) \\ &\leq e(d+1)2^d \frac{M_\rho^2}{m_\rho^2} (M_\phi - m_\phi) (\mathcal{K}_\rho(\psi^1) - \mathcal{K}_\rho(\psi^0) + \langle \psi^1 - \psi^0 | \mu^0 \rangle), \end{aligned}$$

where $\mathcal{K}_\rho : \psi \mapsto \langle \psi^* | \rho \rangle$ denotes the Kantorovich functional associated to ρ and studied in Part I (see Section 1.1.3). From Lemma 1.8, we know that $-\mu^1 = -(\nabla \phi^1)_{\#} \rho$ satisfies

$$\mathcal{K}_\rho(\psi^1) + \langle \psi^0 - \psi^1 | -\mu^1 \rangle \leq \mathcal{K}_\rho(\psi^0).$$

Combining the last two inequalities thus yields

$$\text{Var}_\rho(\phi^1 - \phi^0) \leq C_{d,\rho}(M_\phi - m_\phi) \langle \psi^0 - \psi^1 | \mu^1 - \mu^0 \rangle,$$

where $C_{d,\rho} = e(d+1)2^d \frac{M_\rho^2}{m_\rho^2}$.

The first estimate follows from Kantorovich-Rubinstein duality result: for any $x \in \mathbb{R}^d$ and $k \in \{0, 1\}$, for any $g^k \in \partial \psi^k(x)$, one has $g^k \in \mathcal{X}$ (see Remarks 1.6 and 1.7) so that $\psi^1 - \psi^0$ is $\text{diam}(\mathcal{X})$ -Lipschitz continuous. Kantorovich-Rubinstein duality formula (Proposition A.8) then ensures

$$\langle \psi^0 - \psi^1 | \mu^1 - \mu^0 \rangle \leq \text{diam}(\mathcal{X}) W_1(\mu^0, \mu^1).$$

The second estimate follows from the fact that, if μ^1 is absolutely continuous with respect to μ^0 , then we have for any constant $c \in \mathbb{R}$:

$$\begin{aligned} \langle \psi^0 - \psi^1 | \mu^1 - \mu^0 \rangle &= \langle \psi^0 - \psi^1 - c | \mu^1 - \mu^0 \rangle \\ &= \int_{\mathbb{R}^d} (\psi^0 - \psi^1 - c) \left(\frac{d\mu^1}{d\mu^0} - 1 \right) d\mu^0 \\ &\leq \left(\int_{\mathbb{R}^d} (\psi^0 - \psi^1 - c)^2 d\mu^0 \right)^{1/2} \left(\int_{\mathbb{R}^d} \left(\frac{d\mu^1}{d\mu^0} - 1 \right)^2 d\mu^0 \right)^{1/2} \\ &= \| \psi^0 - \psi^1 - c \|_{L^2(\mu^0)} D_{\chi^2}(\mu^1 | \mu^0)^{1/2} \\ &\leq \sqrt{2} \| \psi^0 - \psi^1 - c \|_{L^2(\frac{1}{2}(\mu^0 + \mu^1))} D_{\chi^2}(\mu^1 | \mu^0)^{1/2}. \end{aligned}$$

The second estimate comes after minimizing with respect to c in the last inequality:

$$\langle \psi^0 - \psi^1 | \mu^1 - \mu^0 \rangle \leq \sqrt{2} \text{Var}_{\frac{1}{2}(\mu^0 + \mu^1)}(\psi^1 - \psi^0)^{1/2} D_{\chi^2}(\mu^1 | \mu^0)^{1/2}. \quad \square$$

All the stability estimates that have been established so far involve the oscillation of the Brenier potentials $M_\phi - m_\phi$. It is then natural to wonder under what assumption on a measure $\mu \in \mathcal{P}_2(\mathbb{R}^d)$ can we control this oscillation. The next proposition, found in (Berman and Berndtsson, 2013), shows that a sufficient condition is that μ admits a finite moment of order $p > d$. This assumption seems nearly tight : Remark 5.10 below shows that there exists a measure μ such that $M_p(\mu) < +\infty$ with $p < d$, whose associated Brenier potential is unbounded.

Proposition 5.7 (Proposition 2.22 in (Berman and Berndtsson, 2013)). *Let ρ be a probability density over a compact convex set \mathcal{X} , satisfying $0 < m_\rho \leq \rho \leq M_\rho$ and let $\mu \in \mathcal{P}_2(\mathbb{R}^d)$. Denote ϕ the Brenier potential for the quadratic optimal transport between ρ and μ . Assume that there exists $p > d$ and $M_p < +\infty$ such that*

$$M_p(\mu) = \int_{\mathbb{R}^d} \|y\|^p d\mu(y) \leq M_p.$$

Then ϕ is Hölder continuous and verifies for all $x, x' \in \mathcal{X}$:

$$|\phi(x) - \phi(x')| \leq C_{d,p,\mathcal{X}} \left(\frac{M_p}{m_\rho} \right)^{1/p} \|x - x'\|^{1-\frac{d}{p}}.$$

In particular, there exists $m_\phi, M_\phi \in \mathbb{R}$ that can be chosen such that for any $x \in \mathcal{X}$, $m_\phi \leq \phi(x) \leq M_\phi$ and such that

$$M_\phi - m_\phi \leq C_{d,p,\mathcal{X}} \left(\frac{M_p}{m_\rho} \right)^{1/p} \text{diam}(\mathcal{X})^{1-\frac{d}{p}}.$$

Before proving Proposition 5.7, let us mention the following consequence when combined with Corollary 5.6.

Corollary 5.8 (Stability with enough moments). *Let ρ be a probability density over a compact convex set \mathcal{X} , satisfying $0 < m_\rho \leq \rho \leq M_\rho$. For any $\mu \in \mathcal{P}_2(\mathbb{R}^d)$, denote ϕ_μ the Brenier potential for the optimal transport between ρ and μ . Let $p > d$. Then the restriction of the mapping $\mu \mapsto \phi_\mu$ to the set of probability measures with bounded p -th moment is $1/2$ -Hölder with respect to the W_1 distance. More precisely, if $\max(M_p(\mu^0), M_p(\mu^1)) \leq M_p < +\infty$, then*

$$\|\phi_{\mu^1} - \phi_{\mu^0}\|_{L^2(\rho)} \leq C_{d,p,\mathcal{X},\rho,M_p} W_1(\mu^0, \mu^1)^{1/2}.$$

Remark 5.9. A large class of probability distributions admit a finite moment of order $p > d$. For instance, sub-exponential measures, which encompass most of the commonly used heavy-tailed distributions fall into this class. We say that a measure $\mu \in \mathcal{P}(\mathbb{R}^d)$ is sub-exponential with variance proxy σ^2 for $\sigma > 0$ if it has zero mean and if for all $r > 0$,

$$\mu(\{x \in \mathbb{R}^d \mid \|x\| \geq r\}) \leq 2e^{-2r/\sigma}.$$

We refer to Proposition 2.7.1 in (Vershynin, 2018) for equivalent characterization. The moments of such a measure are all bounded, and more precisely,

$$M_p(\mu) \leq 2p! \left(\frac{\sigma}{2} \right)^p.$$

We report the proof of Proposition 5.7 from (Berman and Berndtsson, 2013) for completeness.

Proof of Proposition 5.7. The gradient $\nabla\phi$ corresponds to the optimal transport map between ρ and μ . Using that μ is the image of ρ under $\nabla\phi$, the moment assumption gives,

$$\|\nabla\phi\|_{L^p(\mathcal{X})}^p = \int_{\mathcal{X}} \|\nabla\phi(x)\|^p dx \leq \frac{1}{m_\rho} \int_{\mathcal{X}} \|\nabla\phi(x)\|^p d\rho(x) \leq \frac{M_p}{m_\rho}.$$

We can add a constant to ϕ so that $\int_{\mathcal{X}} \phi(x) dx = 0$ without changing its modulus of continuity. The Poincaré-Wirtinger inequality then ensures that $\|\phi\|_{L^p(\mathcal{X})} \leq C_{p,\mathcal{X}} \|\nabla\phi\|_{L^p(\mathcal{X})}$. In particular, the potential ϕ belongs to the Sobolev space $W^{1,p}(\mathcal{X})$. Morrey's inequality (Theorem 11.34 and Theorem 12.15 in (Leoni, 2009)) ensures that ϕ is $(1 - \frac{d}{p})$ -Hölder and that there exists a constant depending only on d, p and \mathcal{X} such that

$$\forall x \neq x' \in \mathcal{X}, \quad \frac{|\phi(x) - \phi(x')|}{\|x - x'\|^{1-\frac{d}{p}}} \leq C_{d,p,\mathcal{X}} \|\phi\|_{W^{1,p}(\mathcal{X})} \leq C_{d,p,\mathcal{X}} \left(\frac{M_p}{m_\rho} \right)^{1/p}. \quad \square$$

Remark 5.10 (Morrey's inequality for convex functions). Since the Brenier potentials ϕ are convex, one may wonder whether Morrey's inequality and the resulting Sobolev embedding can be improved when restricted to the class of convex functions. However, one can show that for $\mathcal{X} = [0, 1]^d$ and $p < d$, for $\alpha \in \left(0, \frac{d}{p} - 1\right)$, the potential

$$\phi : \begin{cases} \mathcal{X} \rightarrow \mathbb{R} \\ (x_1, \dots, x_d) \mapsto (x_1 + \dots + x_d)^{-\alpha} \end{cases}$$

is convex, belongs to $W^{1,p}(\mathcal{X})$, but obviously neither Hölder continuous nor even bounded. In other words, assuming that $M_p(\mu) < +\infty$ for $p < d$ does not guarantee that the Brenier potential from ρ to μ is α -Hölder, or even bounded.

5.3 Stability of optimal transport maps

In this section, we derive quantitative stability estimates on optimal transport maps with respect to the target measures from the stability estimates on Brenier potentials given in the preceding section. This derivation relies on a Gagliardo–Nirenberg type inequality on the difference of convex functions, which is reported here but will be proven in Section 5.4.

Proposition 5.11. *Let K be a compact domain of \mathbb{R}^d with rectifiable boundary and let $u, v : K \rightarrow \mathbb{R}$ be two Lipschitz functions on K that are convex on any segment included in K . Then there exists a constant C_d depending only on d such that*

$$\|\nabla u - \nabla v\|_{L^2(K, \mathbb{R}^d)}^2 \leq C_d \mathcal{H}^{d-1}(\partial K)^{\frac{2}{3}} (\|\nabla u\|_{L^\infty(K, \mathbb{R}^d)} + \|\nabla v\|_{L^\infty(K, \mathbb{R}^d)})^{\frac{4}{3}} \|u - v\|_{L^2(K)}^{\frac{2}{3}},$$

where \mathcal{H}^{d-1} denotes the $(d-1)$ -dimensional Hausdorff measure.

With this proposition at hand, the stability result for Brenier potentials can readily be transferred to stability of the corresponding optimal transport maps – that is, to their gradient – at least when the target measures are compactly supported. Indeed, Proposition 5.11 together with Corollary 5.6 directly imply:

Theorem 5.12 (Stability of the Brenier map, compact case). *Let \mathcal{X}, \mathcal{Y} be compact subsets of \mathbb{R}^d with \mathcal{X} convex, let ρ be a probability density over \mathcal{X} bounded from above and below by positive constants and let $\mu^0, \mu^1 \in \mathcal{P}(\mathcal{Y})$. Denoting T_{μ^k} the Brenier map from ρ to μ^k , we have*

$$W_2(\mu^0, \mu^1) \leq \|T_{\mu^0} - T_{\mu^1}\|_{L^2(\rho, \mathbb{R}^d)} \leq C_{d, \rho, \mathcal{X}, \mathcal{Y}} W_1(\mu^0, \mu^1)^{\frac{1}{6}}.$$

In particular, the embedding $\mu \in \mathcal{P}_2(\mathcal{Y}) \rightarrow T_\mu \in L^2(\rho, \mathbb{R}^d)$ is bi-Hölder continuous.

Remark 5.13 (bi-Hölder embedding via potentials). The previous theorem and Proposition 5.11 together with Corollary 5.6 also ensure the following bi-Hölder behavior for the Brenier potentials (with zero mean against ρ on \mathcal{X}):

$$\forall \mu^0, \mu^1 \in \mathcal{P}(\mathcal{Y}), \quad W_2(\mu^0, \mu^1)^3 \lesssim \|\phi^1 - \phi^0\|_{L^2(\rho)} \lesssim W_1(\mu^0, \mu^1)^{\frac{1}{2}},$$

where the \lesssim notation hides multiplicative constants depending on $d, \rho, \mathcal{X}, \mathcal{Y}$.

We now phrase a similar stability result for probability measures whose Brenier potential is Hölder continuous and that admit a bounded fourth order moment. This includes a large class of probability measures, as noticed in Proposition 5.7 and Remark 5.9.

Theorem 5.14 (Stability of the Brenier map). *Let ρ be a probability density over a compact convex set $\mathcal{X} \subset \mathbb{R}^d$, satisfying $0 < m_\rho \leq \rho \leq M_\rho$. Let $\mu^0, \mu^1 \in \mathcal{P}_2(\mathbb{R}^d)$ and denote ϕ^0, ϕ^1 the Brenier potentials for the quadratic optimal transport between ρ and μ^0, μ^1 respectively. Assume that there exists $M_\alpha > 0$ and $\alpha \in (0, 1)$ such that for all $x, x' \in \mathcal{X}$ and $k \in \{0, 1\}$,*

$$\left| \phi^k(x) - \phi^k(x') \right| \leq M_\alpha \|x - x'\|^\alpha.$$

Assume that there exists $0 < M < +\infty$ such that for $k \in \{0, 1\}$, $M_4(\mu^k) \leq M$. Then

$$W_2(\mu^0, \mu^1) \leq \|\nabla\phi^1 - \nabla\phi^0\|_{L^2(\rho, \mathbb{R}^d)} \leq C_{d, \rho, \mathcal{X}, \alpha, M_\alpha, M} W_1(\mu^0, \mu^1)^{\frac{1}{2(11-8\alpha)}}. \quad (5.5)$$

Remark 5.15. The assumption $M_4(\mu^k) < +\infty$ comes from a use of the Cauchy-Schwarz inequality in the proof of Theorem 5.14. However, one could use Hölder's inequality instead, under different moment assumption and show that for any $q \geq 1$, assuming that $M_{2q}(\mu^k) \leq M_{2q} < +\infty$ for $k \in \{0, 1\}$, one has

$$\|\nabla\phi^1 - \nabla\phi^0\|_{L^2(\rho)} \leq C_{d, \rho, \mathcal{X}, M_\alpha, \alpha, M_{2q}} W_1(\mu^0, \mu^1)^{\frac{q-1}{2(q(7-4\alpha)-3)}}.$$

Since the exponent is an increasing function of q , a stronger stability can be obtained at the cost of stronger moment assumptions.

Theorem 5.14 and Proposition 5.7 directly imply the following.

Corollary 5.16 (Stability with enough moments). *Let ρ be a probability density over a compact convex set $\mathcal{X} \subset \mathbb{R}^d$, satisfying $0 < m_\rho \leq \rho \leq M_\rho$. For $\mu \in \mathcal{P}_2(\mathbb{R}^d)$, denote $\nabla\phi_\mu$ the optimal transport map for the quadratic optimal transport between ρ and μ . Let $p \in \mathbb{R}$ and assume $p \geq 4$ and $p > d$. Then, the map $\mu \mapsto T_\mu$ is Hölder when restricted to the set of probability measures with bounded p -th moment. More precisely, if $\max(M_p(\mu^0), M_p(\mu^1)) \leq M_p < +\infty$, then*

$$W_2(\mu^0, \mu^1) \leq \|\nabla\phi_{\mu^1} - \nabla\phi_{\mu^0}\|_{L^2(\rho, \mathbb{R}^d)} \leq C_{d, p, \mathcal{X}, \rho, M_p} W_1(\mu^0, \mu^1)^{\frac{p}{6p+16d}}.$$

To prove Theorem 5.14, we first show that whenever a Brenier potential defined on the compact and convex set \mathcal{X} is Hölder continuous, it is possible to control its Lipschitz constant on erosions of \mathcal{X} . We recall that for $\eta > 0$, the η -erosion of \mathcal{X} , denoted $\mathcal{X}_{-\eta}$, corresponds to the set of points of \mathcal{X} that are at least at a distance η from $\partial\mathcal{X}$. The proof of this proposition is inspired by Proposition 3.3 in (Klartag, 2014).

Proposition 5.17 (Lipschitz behavior on erosion). *Let ρ be a probability density over a compact convex set $\mathcal{X} \subset \mathbb{R}^d$, satisfying $0 < m_\rho \leq \rho \leq M_\rho$. Let $\mu \in \mathcal{P}_2(\mathbb{R}^d)$ and denote ϕ the Brenier potential for the quadratic optimal transport between ρ and μ . Assume that there exists $M_\alpha > 0$ and $\alpha \in (0, 1)$ such that for all $x, x' \in \mathcal{X}$,*

$$|\phi(x) - \phi(x')| \leq M_\alpha \|x - x'\|^\alpha.$$

Then, ϕ is R -Lipschitz on the erosion $\mathcal{X}_{-\eta_R}$ with $\eta_R = \left(\frac{M_\alpha}{R}\right)^{\frac{1}{1-\alpha}}$.

Proof. Let $x \in \mathcal{X}$ be such that $d(x, \partial\mathcal{X}) \geq \eta_R$, and let $g \in \partial\phi(x)$. We will show that $\|g\| \leq R$, thus implying the statement. Denoting $\psi = (\phi)^*$, the Fenchel-Young equality and inequality ensures that

$$\begin{cases} \psi(g) = \langle g|x \rangle - \phi(x), \\ \psi(g) \geq \langle g|x' \rangle - \phi(x') \quad \text{for all } x' \in \mathcal{X}. \end{cases}$$

Putting these equations together, we get that for any $x' \in \mathcal{X}$,

$$\langle g|x' - x \rangle \leq \phi(x') - \phi(x) \leq M_\alpha \|x' - x\|^\alpha, \quad (5.6)$$

where we used the Hölder continuity assumption on ϕ . We now choose x' to be the unique point in the intersection between the ray $x + \mathbb{R}^+g$ and $\partial\mathcal{X}$, so that $\langle g|x' - x \rangle = \|x - x'\| \|g\|$ and in (5.6),

$$\|g\| \leq \frac{M_\alpha}{\|x - x'\|^{1-\alpha}}.$$

Now using $\|x' - x\| \geq d(x, \partial\mathcal{X}) \geq \eta_R$ in this last inequality yields $\|g\| \leq R$. □

Proposition 5.17 allows to control the Lipschitz constant of the restriction ϕ^k to $\mathcal{X}_{-\eta}$ assuming that ϕ^k is α -Hölder continuous. Combining it with the inequality of Proposition 5.11, we get a stability estimate for the restriction of the transport map to $\mathcal{X}_{-\eta}$. To conclude the proof of the theorem, we will rely on an upper bound on the volume of the symmetric difference between \mathcal{X} and its erosion $\mathcal{X}_{-\eta}$ given in the next proposition.

Proposition 5.18 (Volume of boundary slices). *Let $\mathcal{X} \subset \mathbb{R}^d$ be a compact convex set containing the origin, and denote $r_{\mathcal{X}} > 0$ and $R_{\mathcal{X}} > 0$ the largest and smallest radii such that $B(0, r_{\mathcal{X}}) \subseteq \mathcal{X} \subseteq B(0, R_{\mathcal{X}})$. Then, for all $\eta \geq 0$,*

$$\text{vol}^d(\mathcal{X} \setminus \mathcal{X}_{-\eta}) \leq 2S_{d-1}(R_{\mathcal{X}} + r_{\mathcal{X}})^{d-1} \frac{R_{\mathcal{X}}}{r_{\mathcal{X}}} \eta,$$

where S_{d-1} denotes the surface area of the $(d-1)$ -dimensional unit sphere.

We quote a lemma extracted from (Matheron, 1978) that allows to control the volume of the difference between a convex \mathcal{X} and its η -erosion $\mathcal{X}_{-\eta}$ using the volume of η -dilation of \mathcal{X} , denoted $\mathcal{X}_{+\eta} = \{x \in \mathbb{R}^d \mid d(x, \mathcal{X}) \leq \eta\}$.

Lemma 5.19 ((Matheron, 1978), Lemma 1). *For all $\eta \leq r_{\mathcal{X}}$, $\text{vol}^d(\mathcal{X} \setminus \mathcal{X}_{-\eta}) \leq \text{vol}^d(\mathcal{X}_{+\eta} \setminus \mathcal{X})$.*

This lemma, together with Steiner's formula already implies that $\text{vol}^d(\mathcal{X} \setminus \mathcal{X}_{-\eta})$ grows linearly in η for small values of η . We provide a direct proof below.

Proof of Proposition 5.18. This result is proven using the radial function of \mathcal{X} , $\Lambda_{\mathcal{X}}(x) = \max\{\lambda \geq 0 \mid \lambda x \in \mathcal{X}\}$. Since $x \in \mathcal{S}^{d-1} \mapsto \Lambda_{\mathcal{X}}(x)x$ is a radial parametrization of $\partial\mathcal{X}$, we have:

$$\text{vol}^d(\mathcal{X}) = \int_{\mathcal{X}} 1 dx = \int_{\mathcal{S}^{d-1}} \int_0^{\Lambda_{\mathcal{X}}(u)} r^{d-1} dr du = \frac{1}{d} \int_{\mathcal{S}^{d-1}} \Lambda_{\mathcal{X}}(u)^d du.$$

Combined with Lemma 5.19, this implies that for any $0 \leq \eta \leq r_{\mathcal{X}}$,

$$\begin{aligned} \text{vol}^d(\mathcal{X} \setminus \mathcal{X}_{-\eta}) &\leq \text{vol}^d(\mathcal{X}_{+\eta} \setminus \mathcal{X}) = \frac{1}{d} \int_{\mathcal{S}^{d-1}} \left(\Lambda_{\mathcal{X}_{+\eta}}(u)^d - \Lambda_{\mathcal{X}}(u)^d \right) du \\ &= \frac{1}{d} \int_{\mathcal{S}^{d-1}} (\Lambda_{\mathcal{X}_{+\eta}}(u) - \Lambda_{\mathcal{X}}(u)) \left(\sum_{k=0}^{d-1} \Lambda_{\mathcal{X}_{+\eta}}(u)^{d-1-k} \Lambda_{\mathcal{X}}(u)^k \right) du \\ &\leq \frac{1}{d} \int_{\mathcal{S}^{d-1}} (\Lambda_{\mathcal{X}_{+\eta}}(u) - \Lambda_{\mathcal{X}}(u)) d \cdot (R_{\mathcal{X}} + r_{\mathcal{X}})^{d-1} du. \end{aligned}$$

Using the inclusions $B(0, r_{\mathcal{X}}) \subseteq \mathcal{X} \subseteq B(0, R_{\mathcal{X}})$, one can prove that for any $\eta > 0$ and for any unit vector u ,

$$0 \leq \Lambda_{\mathcal{X}_{+\eta}}(u) - \Lambda_{\mathcal{X}}(u) \leq \frac{(r_{\mathcal{X}}^2 + R_{\mathcal{X}}^2)^{1/2}}{r_{\mathcal{X}}} \eta \leq \frac{2R_{\mathcal{X}}}{r_{\mathcal{X}}} \eta.$$

This can be seen from the *worst case* where \mathcal{X} is an *ice cream cone* made from the convex hull of $B(0, r_{\mathcal{X}})$ and a point at distance $R_{\mathcal{X}}$ of the origin. This finally gives, for $\eta \in [0, r_{\mathcal{X}}]$,

$$\text{vol}^d(\mathcal{X} \setminus \mathcal{X}_{-\eta}) \leq \int_{\mathcal{S}^{d-1}} \frac{2R_{\mathcal{X}}}{r_{\mathcal{X}}} \eta (R_{\mathcal{X}} + r_{\mathcal{X}})^{d-1} du = 2S_{d-1} (R_{\mathcal{X}} + r_{\mathcal{X}})^{d-1} \frac{R_{\mathcal{X}}}{r_{\mathcal{X}}} \eta.$$

One can easily check that in the case $\eta \geq r_{\mathcal{X}}$ the inequality also holds. \square

Proof of Theorem 5.14. In the following, the \lesssim notation hides multiplicative constants that might depend on $d, \rho, \mathcal{X}, \alpha, M_{\alpha}, M$. We get the left inequality of (5.5) by recalling that

$$\mathbb{W}_2(\mu^0, \mu^1)^2 = \min_{\gamma \in \Pi(\mu^0, \mu^1)} \int_{\mathbb{R}^d \times \mathbb{R}^d} \|x - y\|^2 d\gamma(x, y),$$

and by noticing that the optimal transport maps $\nabla\phi^0, \nabla\phi^1$ between ρ and μ^0, μ^1 yield an admissible coupling $\gamma^{0,1} := (\nabla\phi^0, \nabla\phi^1)_{\#} \rho \in \Pi(\mu^0, \mu^1)$, which leads to:

$$\mathbb{W}_2(\mu^0, \mu^1)^2 \leq \int_{\mathbb{R}^d \times \mathbb{R}^d} \|x - y\|^2 d\gamma^{0,1}(x, y) = \int_{\mathcal{X}} \|\nabla\phi^1 - \nabla\phi^0\|^2 d\rho.$$

We now prove the right inequality of (5.5). We recall that $\eta_R = \left(\frac{M_{\alpha}}{R}\right)^{\frac{1}{1-\alpha}}$. Then, denoting ρ_R the restriction of ρ to $\mathcal{X}_{-\eta_R}$ and $\rho_R^{\perp} = \rho - \rho_R$,

$$\|\nabla\phi^1 - \nabla\phi^0\|_{L^2(\rho, \mathbb{R}^d)}^2 = \|\nabla\phi^1 - \nabla\phi^0\|_{L^2(\rho_R, \mathbb{R}^d)}^2 + \|\nabla\phi^1 - \nabla\phi^0\|_{L^2(\rho_R^{\perp}, \mathbb{R}^d)}^2.$$

On $\mathcal{X}_{-\eta_R}$, Proposition 5.17 ensures that $\|\nabla\phi^k\| \leq R$ for $k \in \{0, 1\}$. This fact thus ensures with Proposition 5.11 that for any $c \in \mathbb{R}$:

$$\|\nabla\phi^1 - \nabla\phi^0\|_{L^2(\rho_R, \mathbb{R}^d)}^2 \lesssim R^{4/3} \|\phi^1 - \phi^0 - c\|_{L^2(\rho_R)}^2.$$

Note that we used the inequality $\mathcal{H}^{d-1}(\partial\mathcal{X}_{-\eta_R}) \leq \mathcal{H}^{d-1}(\partial\mathcal{X})$ obtained from the inclusion of the convex set $\mathcal{X}_{-\eta_R}$ into \mathcal{X} , where the convexity of $\mathcal{X}_{-\eta_R}$ is visible from $\mathcal{X}_{-\eta_R} = \bigcap_{\|e\|=\eta_R} (\mathcal{X} - e)$. Minimizing over c in the last inequality thus ensures

$$\|\nabla\phi^1 - \nabla\phi^0\|_{L^2(\rho_R, \mathbb{R}^d)}^2 \lesssim R^{4/3} \text{Var}_{\rho}(\phi^1 - \phi^0)^{1/3} \lesssim R^{4/3} \mathbb{W}_1(\mu^0, \mu^1)^{1/3}, \quad (5.7)$$

where we used Corollary 5.6 to get the second inequality. On the other hand, notice that

$$\|\nabla\phi^1 - \nabla\phi^0\|_{L^2(\rho_R^\perp, \mathbb{R}^d)}^2 \leq 2\|\nabla\phi^1\|_{L^2(\rho_R^\perp, \mathbb{R}^d)}^2 + 2\|\nabla\phi^0\|_{L^2(\rho_R^\perp, \mathbb{R}^d)}^2.$$

By the Cauchy-Schwartz inequality we have for $k \in \{0, 1\}$

$$\begin{aligned} \|\nabla\phi^k\|_{L^2(\rho_R^\perp, \mathbb{R}^d)}^2 &= \int_{\mathcal{X} \setminus \mathcal{X}_{-\eta_R}} \|\nabla\phi^k\|^2 d\rho \\ &\leq \left(\int_{\mathcal{X} \setminus \mathcal{X}_{-\eta_R}} \|\nabla\phi^k\|^4 d\rho \right)^{1/2} \left(\int_{\mathcal{X} \setminus \mathcal{X}_{-\eta_R}} 1^2 d\rho \right)^{1/2} \\ &\lesssim M_4(\mu^k)^{1/2} \text{vol}^d(\mathcal{X} \setminus \mathcal{X}_{-\eta_R})^{1/2}. \end{aligned}$$

Proposition 5.18 ensures that for any $R \geq 0$, we have

$$\text{vol}^d(\mathcal{X} \setminus \mathcal{X}_{-\eta_R}) \lesssim \eta_R = \left(\frac{M_\alpha}{R} \right)^{1/(1-\alpha)}.$$

This gives thus the estimation

$$\|\nabla\phi^1 - \nabla\phi^0\|_{L^2(\rho_R^\perp, \mathbb{R}^d)}^2 \lesssim R^{-1/2(1-\alpha)} \quad (5.8)$$

Estimations (5.7) and (5.8) thus give for $R \geq 0$

$$\|\nabla\phi^1 - \nabla\phi^0\|_{L^2(\rho, \mathbb{R}^d)}^2 \lesssim R^{4/3} \mathbb{W}_1(\mu^0, \mu^1)^{1/3} + R^{-1/2(1-\alpha)}. \quad (5.9)$$

Solving for $R^{4/3} \mathbb{W}_1(\mu^0, \mu^1)^{1/3} = R^{-1/2(1-\alpha)}$ yields $R = \mathbb{W}_1(\mu^0, \mu^1)^{\frac{-2(1-\alpha)}{11-8\alpha}}$. Injecting this value of R in (5.9) yields the desired estimate. \square

We finally prove that if the target measures μ^0, μ^1 are supported on a compact set $\mathcal{Y} \subset \mathbb{R}^d$, if they are absolutely continuous and if their densities are bounded away from zero and infinity, then the Hölder exponents can be slightly improved.

Corollary 5.20. *Let \mathcal{X}, \mathcal{Y} be compact subsets of \mathbb{R}^d , and assume that \mathcal{X} is convex and that \mathcal{Y} has a rectifiable boundary. Let ρ be a probability density over \mathcal{X} satisfying $0 < m_\rho \leq \rho \leq M_\rho < +\infty$ and let μ^0, μ^1 be probability densities over \mathcal{Y} satisfying*

$$\forall k \in \{0, 1\}, \quad 0 < c_\mu \leq \mu^k \leq C_\mu < +\infty.$$

Then, if ϕ^k (resp. T^k) is the Brenier potential (resp. Brenier map) from ρ to μ^k , we have

$$\begin{aligned} \mathbb{W}_2(\mu^0, \mu^1)^6 &\lesssim \text{Var}_\rho(\phi^1 - \phi^0) \leq 2\text{Var}_{\frac{1}{2}(\mu^0 + \mu^1)}(\psi^1 - \psi^0) \lesssim \mathbb{W}_2(\mu^0, \mu^1)^{\frac{6}{5}}, \\ \mathbb{W}_2(\mu^0, \mu^1) &\leq \|T^1 - T^0\|_{L^2(\rho, \mathbb{R}^d)} \lesssim \mathbb{W}_2(\mu^0, \mu^1)^{\frac{1}{5}}, \end{aligned}$$

where the \lesssim notation hides multiplicative constants depending on $d, \rho, \mathcal{X}, \mathcal{Y}, c_\mu$ and C_μ .

This corollary will be a consequence of the following lemma from (Maury et al., 2010), which we will use as a replacement of the Kantorovich-Rubinstein inequality.

Lemma 5.21 (Lemma 3.5 in (Maury et al., 2010)). *Assume that μ^0 and μ^1 are absolutely continuous measures on the compact \mathcal{Y} , whose densities are bounded by a common constant C_μ . Then, for any function $f \in H^1(\mathcal{Y})$, we have the following inequality:*

$$\int_{\mathcal{Y}} f d(\mu^1 - \mu^0) \leq \sqrt{C_\mu} \|\nabla f\|_{L^2(\mathcal{Y})} W_2(\mu^0, \mu^1).$$

Proof of Corollary 5.20. Because \mathcal{Y} is compact, the Brenier potentials ϕ^0, ϕ^1 are $R_{\mathcal{Y}}$ -Lipschitz continuous for any $R_{\mathcal{Y}} \in \mathbb{R}_+$ such that $\mathcal{Y} \subset B(0, R_{\mathcal{Y}})$. One can thus find $m_\phi, M_\phi \in \mathbb{R}$ such that for $k \in \{0, 1\}$, $m_\phi \leq \phi^k \leq M_\phi$ on \mathcal{X} and $M_\phi - m_\phi \leq R_{\mathcal{Y}} \text{diam}(\mathcal{X})$. Setting $\psi^0 = (\phi^0)^*$, $\psi^1 = (\phi^1)^*$, we thus have from Theorem 3.1:

$$\text{Var}_{\frac{1}{2}(\mu^0 + \mu^1)}(\psi^1 - \psi^0) \lesssim \langle \psi^0 - \psi^1 | \mu^1 - \mu^0 \rangle. \quad (5.10)$$

For $c \in \mathbb{R}$ such that $\|\psi^1 - \psi^0 - c\|_{L^2(\frac{1}{2}(\mu^0 + \mu^1))}^2 = \text{Var}_{\frac{1}{2}(\mu^0 + \mu^1)}(\psi^1 - \psi^0)$, estimation (5.10) and Lemma 5.21 ensure that:

$$\|\psi^1 - \psi^0 - c\|_{L^2(\frac{1}{2}(\mu^0 + \mu^1))}^2 \lesssim \|\nabla \psi^1 - \nabla \psi^0\|_{L^2(\mathcal{Y})} W_2(\mu^0, \mu^1). \quad (5.11)$$

But Proposition 5.11 applied to the convex and Lipschitz functions $\psi^0 + c, \psi^1$ ensures that

$$\|\nabla \psi^1 - \nabla \psi^0\|_{L^2(\mathcal{Y})} \lesssim \|\psi^1 - \psi^0 - c\|_{L^2(\frac{1}{2}(\mu^0 + \mu^1))}^{1/3}.$$

Injecting this estimation into (5.11) yields

$$\|\psi^1 - \psi^0 - c\|_{L^2(\frac{1}{2}(\mu^0 + \mu^1))}^2 \lesssim W_2(\mu^0, \mu^1)^{6/5}.$$

This gives thus with Proposition 1.30

$$\text{Var}_\rho(\phi^1 - \phi^0) \leq 2 \text{Var}_{\frac{1}{2}(\mu^0 + \mu^1)}(\psi^1 - \psi^0) \lesssim W_2(\mu^0, \mu^1)^{6/5}.$$

Finally, a last use of Proposition 5.11 also ensures that under these assumptions on the targets μ^0, μ^1 we have

$$W_2(\mu^0, \mu^1) \leq \|\nabla \phi^1 - \nabla \phi^0\|_{L^2(\rho, \mathbb{R}^d)} \lesssim \text{Var}_\rho(\phi^1 - \phi^0)^{\frac{1}{6}} \lesssim W_2(\mu^0, \mu^1)^{\frac{1}{5}}. \quad \square$$

5.4 Gagliardo–Nirenberg type inequality for difference of convex functions

We prove here Proposition 5.11, a sort of reverse Poincaré inequality which allows to control the L^2 distance $\|\nabla u - \nabla v\|_{L^2(K, \mathbb{R}^d)}$ between the gradients of Lipschitz convex functions u, v using the L^2 distance between these functions $\|u - v\|_{L^2(K)}$. This proposition is a refinement of Theorem 3.5 in (Chazal et al., 2010), in which the upper bound involved the uniform distance $\|u - v\|_\infty$. Proposition 5.11 is first proven in dimension $d = 1$ and on a segment (Lemma 5.24) and then generalized to higher dimensions using arguments from integral geometry (Lemma 5.25).

Remark 5.22 (Relation to the Gagliardo–Nirenberg inequality). Although the estimate of Proposition 5.11 resembles the Gagliardo–Nirenberg inequality, it cannot be deduced from it. More precisely, we note that without convexity of u and v , the inequality in (5.11) does not hold. One can see this by taking $u = 0$ and $v_n(x) = 1/n \sin(nx)$ on $K = [0, 1]$.

Remark 5.23 (Optimality of exponents). The inequality proposed in Proposition 5.11 is sharp in term of the exponents of the norms $\|\nabla u\|_{L^\infty(K)} + \|\nabla v\|_{L^\infty(K)}$ and $\|u - v\|_{L^2(K)}$ in the right-hand side. In the case $d = 1$, let $L > 0, \varepsilon > 0$ and define on $K = [0, 1]$, $u(x) = L|x - \frac{1}{2}|$ and $v = \max(u, \varepsilon)$. Then u, v are convex and L -Lipschitz and we have:

$$\|u - v\|_{L^2([0,1])}^2 = \frac{2\varepsilon^3}{3L} \quad \text{and} \quad \|u' - v'\|_{L^2([0,1])}^2 = 2L\varepsilon.$$

so that $\|u' - v'\|_{L^2([0,1])}^2 = 12^{1/3}L^{4/3} \|u - v\|_{L^2([0,1])}^{2/3}$.

Lemma 5.24. *Let $I \subset \mathbb{R}$ be a compact segment and let $u, v : I \rightarrow \mathbb{R}$ be two convex functions with uniformly bounded gradients on I . Then*

$$\|u' - v'\|_{L^2(I)}^2 \leq 8(\|u'\|_{L^\infty(I)} + \|v'\|_{L^\infty(I)})^{4/3} \|u - v\|_{L^2(I)}^{2/3}. \quad (5.12)$$

Proof. We first assume that $I = [0, 1]$. Using a simple approximation, we may assume that u, v are \mathcal{C}^2 on I to get the following integration by part:

$$\|u' - v'\|_{L^2([0,1])}^2 = [(u - v)(u' - v')]_0^1 - \int_{[0,1]} (u - v)(u'' - v'').$$

The convexity hypothesis then allows to get a L^∞ estimate. Indeed, $[(u - v)(u' - v')]_0^1 \leq 2(\|u'\|_{L^\infty} + \|v'\|_{L^\infty}) \|u - v\|_{L^\infty}$, and by convexity

$$\begin{aligned} \left| \int_{[0,1]} (u - v)(u'' - v'') \right| &\leq \|u - v\|_{L^\infty} \left(\int_{[0,1]} |u''| + \int_{[0,1]} |v''| \right) \\ &= \|u - v\|_{L^\infty} \left(\int_{[0,1]} u'' + \int_{[0,1]} v'' \right) \\ &\leq 2(\|u'\|_{L^\infty} + \|v'\|_{L^\infty}) \|u - v\|_{L^\infty}. \end{aligned}$$

This gives

$$\|u' - v'\|_{L^2([0,1])}^2 \leq 4(\|u'\|_{L^\infty} + \|v'\|_{L^\infty}) \|u - v\|_{L^\infty}. \quad (5.13)$$

We now bound the L^∞ norm of $u - v$ with its L^2 norm using that the Lipschitz constant of $u - v$ is less than $L = \|u'\|_{L^\infty} + \|v'\|_{L^\infty}$. Let $\varepsilon = \|u - v\|_{L^\infty}$ and let $x^* \in [0, 1]$ where the maximum of $|u - v|$ is attained. Since $\text{Lip}(u - v) \leq L$, one gets $|u(x) - v(x)| \geq \frac{\varepsilon}{2}$ on the interval $I_* = I \cap [x^* - \frac{\varepsilon}{2L}, x^* + \frac{\varepsilon}{2L}]$. The length of I_* is at least $\min(\frac{\varepsilon}{2L}, 1)$, so that

$$\|u - v\|_{L^2([0,1])}^2 \geq \frac{1}{4} \min\left(\frac{\varepsilon}{2L}, 1\right) \varepsilon^2. \quad (5.14)$$

Assume first that $\varepsilon \leq 2L$. Then, equation (5.14) gives $\varepsilon^3 = \|u - v\|_{L^\infty}^3 \leq 8L \|u - v\|_{L^2([0,1])}^2$, thus implying

$$\|u - v\|_{L^\infty} \leq 2(\|u'\|_{L^\infty} + \|v'\|_{L^\infty})^{1/3} \|u - v\|_{L^2([0,1])}^{2/3}.$$

This gives, with equation (5.13):

$$\|u' - v'\|_{L^2([0,1])}^2 \leq 8(\|u'\|_{L^\infty} + \|v'\|_{L^\infty})^{4/3} \|u - v\|_{L^2([0,1])}^{2/3}. \quad (5.15)$$

On the other hand, if $\varepsilon \geq 2L$, then $\|u - v\|_{L^2([0,1])} \geq \frac{\varepsilon}{2}$ by equation (5.14), so that

$$8(\|u'\|_{L^\infty} + \|v'\|_{L^\infty})^{4/3} \|u - v\|_{L^2([0,1])}^{2/3} \geq 8L^{4/3} \left(\frac{\varepsilon}{2}\right)^{2/3} \geq L^{4/3+2/3} = L^2,$$

which allows to conclude using $L^2 \geq \|u' - v'\|_{L^2([0,1])}^2$. We get inequality (5.12) for a general interval $I = [a, b]$ by an affine change of variable. \square

The one-dimensional result from Lemma 5.24 is generalized to higher dimensions thanks to two formulas from integral geometry that allow to rewrite the L^2 norms of the scalar-field $u - v$ and vector-field $\nabla u - \nabla v$ over set $K \subset \mathbb{R}^d$ using integrals over lines intersecting K .

Integral geometry. Denote V_d the volume of the unit d -ball and S_{d-1} the area of the unit $(d-1)$ -sphere. Let \mathcal{L}^d be the set of oriented affine lines ℓ in \mathbb{R}^d , identified to the submanifold of \mathbb{R}^{2d} consisting of pairs of directions and offsets $(e, p) \in \mathbb{R}^d \times \mathbb{R}^d$, with $e \in \mathcal{S}^{d-1}$ and p in the hyperplane $\{e\}^\perp$, and endowed with the induced Riemannian metric. The volume measure $d\mathcal{L}^d$ is invariant under rigid motions. Denoting \mathcal{H}^k the k -dimensional Hausdorff measure, the usual Crofton formula – see for instance the first paragraph of Chapter 5 in (Hug and Weil, 2020) – states that for any \mathcal{H}^{d-1} -rectifiable subset S of \mathbb{R}^d ,

$$\mathcal{H}^{d-1}(S) = \frac{1}{4V_{d-1}} \int_{\ell \in \mathcal{L}^d} \#(\ell \cap S) d\mathcal{L}^d(\ell), \quad (5.16)$$

where $\#X$ is the cardinality of the set X . We denote \mathcal{L}_e^d the set of oriented lines with a fixed direction $e \in \mathcal{S}^{d-1}$, endowed with the $(d-1)$ -dimensional Lebesgue measure $d\mathcal{L}_e^d$ on $\{e\}^\perp$, so that for any $\phi : \mathcal{L}^d \rightarrow \mathbb{R}$,

$$\int_{\ell \in \mathcal{L}^d} \phi(\ell) d\mathcal{L}^d(\ell) = \int_{e \in \mathcal{S}^{d-1}} \int_{p \in \{e\}^\perp} \phi((e, p)) d\mathcal{L}_e^d(p) d\mathcal{S}^{d-1}(e).$$

We will also use the following formula, which easily follows from Fubini's theorem: if K is a measurable subset of \mathbb{R}^d , then for any fixed direction $e \in \mathcal{S}^{d-1}$,

$$\mathcal{H}^d(K) = \int_{\ell \in \mathcal{L}_e^d} \mathcal{H}^1(\ell \cap K) d\mathcal{L}_e^d(\ell). \quad (5.17)$$

We begin with an elementary lemma.

Lemma 5.25. *Let K be a compact subset of \mathbb{R}^d and $f \in L^2(K)$. Then,*

$$\|f\|_{L^2(K)}^2 = \frac{1}{S_{d-1}} \int_{\ell \in \mathcal{L}^d} \int_{y \in \ell \cap K} f(y)^2 dy d\mathcal{L}^d(\ell). \quad (5.18)$$

Similarly, for any vector field $F \in L^2(K, \mathbb{R}^d)$, one has

$$\|F\|_{L^2(K, \mathbb{R}^d)}^2 = C_d \int_{\ell \in \mathcal{L}^d} \int_{y \in \ell \cap K} \langle F(y) | e(\ell) \rangle^2 dy d\mathcal{L}^d(\ell), \quad (5.19)$$

where for $\ell \in \mathcal{L}^d$, $e(\ell) \in \mathcal{S}^{d-1}$ is the oriented direction of ℓ , and C_d depends only on d .

Proof. Piecewise constant functions (resp. vector fields) are dense in $L^2(K)$ (resp. $L^2(K, \mathbb{R}^d)$). Using this fact and the continuity of equations (5.18), (5.19), it is therefore enough to prove these equations when f and F are of the form $f = \chi_{K'}$ and $F = x\chi_{K'}$ for some fixed measurable set $K' \subset K$ and $x \in \mathcal{S}^{d-1}$, where $\chi_{K'}(u) = 1$ if $u \in K'$ and 0 else. We have for $f = \chi_{K'}$, using formula (5.17):

$$\begin{aligned} S_{d-1} \|f\|_{L^2(K)}^2 &= S_{d-1} \mathcal{H}^d(K') = \int_{e \in \mathcal{S}^{d-1}} \int_{\ell \in \mathcal{L}_e^d} \mathcal{H}^1(\ell \cap K') d\mathcal{L}_e^d(\ell) de \\ &= \int_{e \in \mathcal{S}^{d-1}} \int_{\ell \in \mathcal{L}_e^d} \int_{y \in \ell \cap K'} dy d\mathcal{L}_e^d(\ell) de \\ &= \int_{\ell \in \mathcal{L}^d} \int_{y \in \ell \cap K} f(y)^2 dy d\mathcal{L}^d(\ell), \end{aligned}$$

which proves equation (5.18). Now for $F = x\chi_{K'}$, we get for $e \in \mathcal{S}^{d-1}$:

$$\begin{aligned} \langle x|e \rangle^2 \|F\|_{L^2(K, \mathbb{R}^d)}^2 &= \langle x|e \rangle^2 \mathcal{H}^d(K') = \langle x|e \rangle^2 \int_{\ell \in \mathcal{L}_e^d} \mathcal{H}^1(\ell \cap K') d\mathcal{L}_e^d(\ell) \\ &= \int_{\ell \in \mathcal{L}_e^d} \int_{y \in \ell \cap K} \langle F(y)|e \rangle^2 dy d\mathcal{L}_e^d(\ell). \end{aligned}$$

Hence we get:

$$\left(\int_{e \in \mathcal{S}^{d-1}} \langle x|e \rangle^2 de \right) \|F\|_{L^2(K, \mathbb{R}^d)}^2 = \int_{\ell \in \mathcal{L}^d} \int_{y \in \ell \cap K} \langle F(y)|e(\ell) \rangle^2 dy d\mathcal{L}^d(\ell).$$

The first integral does not depend on x , thus establishing the result. \square

We are now ready to prove the Gagliardo-Nirenberg type inequality of Proposition 5.11.

Proof of Proposition 5.11. We apply formula (5.19) from Lemma 5.25 to $(\nabla u - \nabla v)$:

$$\|\nabla u - \nabla v\|_{L^2(K, \mathbb{R}^d)}^2 = C_d \int_{\ell \in \mathcal{L}^d} \int_{y \in \ell \cap K} \langle (\nabla u - \nabla v)(y)|e(\ell) \rangle^2 dy d\mathcal{L}^d(\ell).$$

For any $\ell \in \mathcal{L}^d$, denote $u_\ell = u|_{\ell \cap K}$, $v_\ell = v|_{\ell \cap K}$, and notice that the last equation reads:

$$\|\nabla u - \nabla v\|_{L^2(K, \mathbb{R}^d)}^2 = C_d \int_{\ell \in \mathcal{L}^d} \|u'_\ell - v'_\ell\|_{L^2(\ell \cap K)}^2 d\mathcal{L}^d(\ell).$$

Given any oriented line $\ell \in \mathcal{L}^d$, denote $n_\ell \in \mathbb{N} \cup \{+\infty\}$ the number of connected components of $\ell \cap K$. Then, $n_\ell \leq \#(\ell \cap \partial K)$ so that by Crofton's formula,

$$\int_{\ell \in \mathcal{L}^d} n_\ell d\mathcal{L}^d(\ell) \leq \int_{\ell \in \mathcal{L}^d} \#(\ell \cap \partial K) d\mathcal{L}^d(\ell) < +\infty.$$

This implies that for almost every $\ell \in \mathcal{L}^d$, the set $\ell \cap K$ may be decomposed as a finite union of n_ℓ segments, i.e. $\ell \cap K = \bigcup_{i=1}^{n_\ell} I_\ell^i$. This gives

$$\|u'_\ell - v'_\ell\|_{L^2(\ell \cap K)}^2 = \sum_{i=1}^{n_\ell} \|u'_\ell - v'_\ell\|_{L^2(I_\ell^i)}^2; \|u_\ell - v_\ell\|_{L^2(\ell \cap K)}^2 = \sum_{i=1}^{n_\ell} \|u_\ell - v_\ell\|_{L^2(I_\ell^i)}^2.$$

Lemma 5.24 combined with Jensen's inequality then ensure that we have for almost every $l \in \mathcal{L}^d$:

$$\begin{aligned} \|u'_\ell - v'_\ell\|_{L^2(\ell \cap K)}^2 &\leq 8(\|u'_\ell\|_{L^\infty(\ell \cap K)} + \|v'_\ell\|_{L^\infty(\ell \cap K)})^{4/3} \sum_{i=1}^{n_\ell} \|u_\ell - v_\ell\|_{L^2(I_\ell^i)}^{2/3} \\ &\leq 8(\|u'_\ell\|_{L^\infty(\ell \cap K)} + \|v'_\ell\|_{L^\infty(\ell \cap K)})^{4/3} n_\ell^{2/3} \|u_\ell - v_\ell\|_{L^2(\ell \cap K)}^{2/3}. \end{aligned}$$

This leads to the inequality

$$\|\nabla u - \nabla v\|_{L^2(K)}^2 \leq 8C_d(2L)^{4/3} \int_{\ell \in \mathcal{L}^d} n_\ell^{2/3} \|u_\ell - v_\ell\|_{L^2(\ell \cap K)}^{2/3} d\mathcal{L}^d(\ell),$$

where $L = \max(\|\nabla u\|_{L^\infty(K)}, \|\nabla v\|_{L^\infty(K)})$. But Hölder's inequality together with formula (5.18) give

$$\int_{\ell \in \mathcal{L}^d} n_\ell^{2/3} \|u_\ell - v_\ell\|_{L^2(\ell \cap K)}^{2/3} d\mathcal{L}^d(\ell) \leq \left(\int_{\ell \in \mathcal{L}^d} n_\ell d\mathcal{L}^d(\ell) \right)^{2/3} S_{d-1}^{1/3} \|u - v\|_{L^2(K)}^{2/3}.$$

The conclusion comes after using again that $n_\ell \leq \#(\ell \cap \partial K)$ and Crofton's formula (5.16)

$$\int_{\ell \in \mathcal{L}^d} n_\ell d\mathcal{L}^d(\ell) \leq \int_{\ell \in \mathcal{L}^d} \#(\ell \cap \partial K) d\mathcal{L}^d(\ell) = 4V_{d-1} \mathcal{H}^{d-1}(\partial K). \quad \square$$

Quantitative stability of Wasserstein barycenters with respect to the marginals

Abstract

This chapter, extracted from (Carlier et al., 2022), derives quantitative stability estimates for Wasserstein barycenters with respect to their marginals. Wasserstein barycenters define averages of probability measures in a geometrically meaningful way. Their use is increasingly popular in applied fields, such as image, geometry or language processing. In these fields however, the probability measures of interest are often not accessible in their entirety and the practitioner may have to deal with statistical or computational approximations instead. In this chapter, we quantify the effect of such approximations on the corresponding barycenters. We show that the Wasserstein barycenter depends in a Hölder-continuous way on its marginals under relatively mild assumptions. Our proof relies on the strong convexity estimates for the Kantorovich functional presented in Part I and a new result quantifying the modulus of continuity of the push-forward operation under a (not necessarily smooth) optimal transport map.

6.1 Introduction

Wasserstein barycenters are Fréchet means in Wasserstein spaces: they define averages of families of probability measures that are consistent with the optimal transport geometry and generalize to more than two measures the fundamental notion of displacement interpolation due to McCann (1997). As such, they average out probability measures in a geometrically meaningful way and appear as a relevant tool to interpolate or summarize measure data. Such notion of barycenter have indeed found many successful applications, for instance in image processing (Rabin et al., 2011), geometry processing (Solomon et al., 2015), language processing (Dognin et al., 2019; Colombo et al., 2021; Lian et al., 2020), statistics (Srivastava et al., 2018) or machine learning (Cuturi and Doucet, 2014; Ho et al., 2017). We refer the readers to existing surveys (Peyré and Cuturi, 2019; Panaretos and Zemel, 2020) for further applications. In such applications however, the probability measures of interest are often not accessible in their

entirety. They may be accessible for instance only through noisy samples in a statistical context, or they may be approximated in order to use existing computational methods that estimate Wasserstein barycenters (see e.g. (Carlier, Guillaume et al., 2015; Benamou et al., 2015; Cuturi and Doucet, 2014; Altschuler and Boix-Adsera, 2021)) while paying an affordable computational cost. This means that in addition to the computational error induced by the algorithm used to calculate the barycenter, the practitioner may be subject to an extra statistical or approximation error that corresponds to the approximation of the marginal measures of interest. While works focusing on the computation of Wasserstein barycenters may now come with guarantees on the first type of error (see e.g. (Altschuler and Boix-Adsera, 2021)), very little is known on the second type of error, which corresponds broadly speaking to a stability error since it quantifies the effect of a perturbation of the marginals on the corresponding barycenters. In this chapter, we focus on this kind of error and show that the Wasserstein barycenter depends in an Hölder-continuous way on its marginal measures under regularity assumptions on (some of) the latter.

Outline. In the remaining of this section, we define Wasserstein barycenters and the setting we focus on. We then show that mild regularity assumptions are necessary in order to hope for any stability result. Next, we announce the dual formulation of the Wasserstein barycenter problem in our context, that is necessary to present our main assumption. This assumption and our main result are then stated and we give some immediate but useful consequences of this result. We conclude this section with the principal elements of proof to our main result, that are made of two stability estimates: a stability estimate for the dual solutions to the Wasserstein barycenter problem (justified in §6.2) and a stability estimate for the push-forward operation (justified in §6.3). The proof of the dual formulation is postponed to the end of this chapter (§6.4).

6.1.1 Wasserstein barycenters

Introduced in (Agueh and Carlier, 2011) for finite families of probability measures supported over a Euclidean space, the definition of Wasserstein barycenters have been extended to infinite families of probability measures in (Bigot and Klein, 2018; Pass, 2013), possibly supported over a Riemannian manifold in (Kim and Pass, 2017; Le Gouic and Loubes, 2017). In this work, we focus on families of probability measures supported over a compact Euclidean domain. Let $\Omega = B(0, R_\Omega) \subset \mathbb{R}^d$ be the ball of \mathbb{R}^d centered at zero and of radius $R_\Omega > 0$ and denote $\mathcal{P}(\Omega)$ the set of Borel probability measures over Ω . We endow $\mathcal{P}(\Omega)$ with the 2-Wasserstein distance W_2 defined for any $\rho, \mu \in \mathcal{P}(\Omega)$ by

$$W_2(\rho, \mu) = \left(\min_{\gamma \in \Gamma(\rho, \mu)} \int_{\Omega \times \Omega} \|x - y\|^2 d\gamma(x, y) \right)^{1/2},$$

where the minimum is taken over the set $\Gamma(\rho, \mu)$ of transport plans between ρ and μ (see e.g. Chapter A for more details about Wasserstein distances). We equip $\mathcal{P}(\Omega)$ with the the topology induced by W_2 (i.e. the weak topology) and denote $\mathcal{P}(\mathcal{P}(\Omega))$ the set of corresponding Borel probability measures over $\mathcal{P}(\Omega)$. A Wasserstein barycenter of $\mathbb{P} \in \mathcal{P}(\mathcal{P}(\Omega))$ is then defined as a minimizer $\mu_{\mathbb{P}}$ of

$$\min \left\{ \frac{1}{2} \int_{\mathcal{P}(\Omega)} W_2^2(\rho, \mu) d\mathbb{P}(\rho), \quad \mu \in \mathcal{P}(\Omega) \right\} = (\text{P})_{\mathbb{P}}.$$

Such a minimizer always exists, and it is uniquely defined whenever $\mathbb{P}(\mathcal{P}_{a.c.}(\Omega) > 0)$, where $\mathcal{P}_{a.c.}(\Omega)$ denotes the set of probability measures over Ω that are absolutely continuous with respect to the Lebesgue measure (Kim and Pass, 2017; Le Gouic and Loubes, 2017).

6.1.2 Stability of Wasserstein barycenters

As mentioned above, the population of interest $\mathbb{P} \in \mathcal{P}(\mathcal{P}(\Omega))$ may not always be accessible in practice, and one may have to deal with another measure $\mathbb{Q} \in \mathcal{P}(\mathcal{P}(\Omega))$ instead. The stability question that then comes up is the following: can we bound a distance between minimizers $\mu_{\mathbb{P}}$ of $(\mathbb{P})_{\mathbb{P}}$ and $\mu_{\mathbb{Q}}$ of $(\mathbb{P})_{\mathbb{Q}}$ in terms of a distance between \mathbb{P} and \mathbb{Q} ? While the above-defined 2-Wasserstein distance gives a natural metric to compare $\mu_{\mathbb{P}}$ and $\mu_{\mathbb{Q}}$, there remains to choose a metric in order to compare \mathbb{P} and \mathbb{Q} . For this, we will use the following 1-Wasserstein distance over $\mathcal{P}(\mathcal{P}(\Omega))$, defined for any \mathbb{P}, \mathbb{Q} in $\mathcal{P}(\mathcal{P}(\Omega))$ by

$$\mathcal{W}_1(\mathbb{P}, \mathbb{Q}) = \min_{\gamma \in \Gamma(\mathbb{P}, \mathbb{Q})} \int_{\mathcal{P}(\Omega) \times \mathcal{P}(\Omega)} W_2(\rho, \tilde{\rho}) d\gamma(\rho, \tilde{\rho}).$$

This choice of distance is justified by the fact that Wasserstein distances are naturally defined for probability measures on the compact metric space $(\mathcal{P}(\Omega), W_2)$ and that they allow to compare measures that have incomparable support. The 1-Wasserstein distance being the weakest of the Wasserstein distances, our bounds are ensured to be the sharpest in terms of this optimal transport geometry. We are thus interested in bounding $W_2(\mu_{\mathbb{P}}, \mu_{\mathbb{Q}})$ in terms of $\mathcal{W}_1(\mathbb{P}, \mathbb{Q})$ for $\mathbb{P}, \mathbb{Q} \in \mathcal{P}(\mathcal{P}(\Omega))$.

Consistency of Wasserstein barycenters. Before looking for any quantitative stability result, one may first wonder if the barycenters depend at least in a continuous way on their marginals. This question, framed under the notion of *consistency* of Wasserstein barycenters, has been answered positively in (Bigot and Klein, 2018; Boissard et al., 2015) in some specific settings and in (Le Gouic and Loubes, 2017) in the most general setting. Theorem 3 of (Le Gouic and Loubes, 2017) ensures in particular the following:

Theorem (Le Gouic, Loubes). *Let $\mathbb{P} \in \mathcal{P}(\mathcal{P}(\Omega))$ and a sequence $(\mathbb{P}_n)_{n \geq 1} \in \mathcal{P}(\mathcal{P}(\Omega))$ be such that*

$$\mathcal{W}_1(\mathbb{P}_n, \mathbb{P}) \xrightarrow{n \rightarrow +\infty} 0.$$

For all $n \geq 1$, denote $\mu_{\mathbb{P}_n}$ a barycenter of \mathbb{P}_n . Then the sequence $(\mu_{\mathbb{P}_n})_{n \geq 1}$ is precompact in $(\mathcal{P}(\Omega), W_2)$ and any limit is a barycenter of \mathbb{P} .

This result ensures the continuity of Wasserstein barycenters with respect to the marginal measures, at least in our setting, so that we can now legitimately look for bounds that quantify this continuity.

Quantitative stability in dimension $d = 1$. In dimension $d = 1$, the derivation of quantitative stability bounds for Wasserstein barycenters is straightforward. Indeed, in such setting W_2 is Hilbertian, which ensures a Lipschitz behavior of the barycenters with respect to their marginals. More precisely, denoting F_{ρ}^{-1} the *quantile function* of a measure $\rho \in \mathcal{P}(\Omega)$ (i.e. the generalized inverse of its cumulative distribution function




Figure 6.1: Let $\rho_1 = \frac{1}{2}(\delta_{(0;1)} + \delta_{(0;-1)})$. For $\varepsilon > 0$ and $x_\varepsilon = (1; \varepsilon/2) \in \mathbb{R}^2$, let $\rho_2^\varepsilon = \frac{1}{2}(\delta_{x_\varepsilon} + \delta_{-x_\varepsilon})$. Introduce $\mathbb{P}_\varepsilon = \frac{1}{2}(\delta_{\rho_1} + \delta_{\rho_2^\varepsilon})$. Then for $\varepsilon \leq \frac{1}{2}$, $W_2(\mu_{\mathbb{P}_\varepsilon}, \mu_{\mathbb{P}_{-\varepsilon}}) = 1$ while $W_1(\mathbb{P}_\varepsilon, \mathbb{P}_{-\varepsilon}) \leq \varepsilon$.

F_ρ), one has for any measures $\rho, \mu \in \mathcal{P}(\Omega)$, $W_2(\rho, \mu) = \|F_\rho^{-1} - F_\mu^{-1}\|_{L^2([0,1])}$. This leads for any $\mathbb{P} \in \mathcal{P}(\mathcal{P}(\Omega))$ to a simple formula for the unique barycenter:

$$\mu_{\mathbb{P}} = \left(\int_{\mathcal{P}(\Omega)} F_\rho^{-1} d\mathbb{P}(\rho) \right)_{\#} \lambda_{[0,1]},$$

where $\lambda_{[0,1]}$ denotes the Lebesgue measure over $[0, 1]$. Using this fact and the triangle inequality, one immediately obtains the following Lipschitz stability result, that actually holds for any families of measures in the set $\mathcal{P}_2(\mathbb{R})$ of probability measures supported over \mathbb{R} that admit a finite second moment:

Proposition. *Let $\mathbb{P}, \mathbb{Q} \in \mathcal{P}(\mathcal{P}_2(\mathbb{R}))$ and denote $\mu_{\mathbb{P}}, \mu_{\mathbb{Q}}$ their respective barycenters. Then*

$$W_2(\mu_{\mathbb{P}}, \mu_{\mathbb{Q}}) \leq W_1(\mathbb{P}, \mathbb{Q}).$$

This fact was exploited in (Bigot et al., 2018) to characterize the statistical rate of convergence of empirical Wasserstein barycenters towards their population counterpart in an asymptotic setting for probability measures supported over the real line.

Quantitative stability in dimension $d \geq 2$. In dimension $d \geq 2$, the derivation of any quantitative stability bound turns out to be much more difficult. This may first come from the fact that without any assumption on \mathbb{P} and \mathbb{Q} , the barycenters $\mu_{\mathbb{P}}$ and $\mu_{\mathbb{Q}}$ may not be uniquely defined, which makes hopeless the derivation of any stability result. Even when uniqueness of the barycenters is ensured, one can easily build examples where no quantitative stability bound holds, see for instance the setting illustrated in Figure 6.1. This example relies on barycenters with only discrete marginals, and recovers in the limit $\varepsilon = 0$ the pathological case where the barycenter is not uniquely defined. One may circumvent this issue by ensuring, even in the limit $\varepsilon = 0$, uniqueness of the barycenter. As mentioned above, this may be done by imposing that some of the marginal measures are absolutely continuous. Nevertheless, even under such assumption on the marginals, one can find an example where the barycenter achieves an Hölder behavior with respect to its marginal, but with an arbitrary low Hölder exponent, see Figure 6.2. These negative results show that, even in dimension $d = 2$, some regularity assumptions on the marginals \mathbb{P}, \mathbb{Q} are necessary in order to hope to derive stability estimates for their barycenters.




Figure 6.2: Let $\rho_1 = \frac{1}{2}(\delta_{(0;1)} + \delta_{(0;-1)})$. For $a \in (0, 1)$ and $\varepsilon > 0$, let $c_\varepsilon = [1 - \frac{a}{2}; 1 + \frac{a}{2}] \times [-\frac{a}{2} + \varepsilon; \frac{a}{2} + \varepsilon]$ and ρ_2^ε the probability measure with density $\rho_2^\varepsilon(x, y) = \frac{\alpha}{2^{1-2\alpha}a^{1+2\alpha}} (|y - \varepsilon|^{2\alpha-1} \mathbb{1}_{c_\varepsilon}(x, y) + |y + \varepsilon|^{2\alpha-1} \mathbb{1}_{-c_\varepsilon}(x, y))$ for some $\alpha > 0$. Introduce $\mathbb{P}_\varepsilon = \frac{1}{2}(\delta_{\rho_1} + \delta_{\rho_2^\varepsilon})$. Then for $\varepsilon \leq \frac{a}{2}$, $W_2(\mu_{\mathbb{P}_0}, \mu_{\mathbb{P}_\varepsilon}) \sim \varepsilon^\alpha$ while $W_1(\mathbb{P}_0, \mathbb{P}_\varepsilon) \leq \varepsilon$.

Previous works. Consistently with the above remarks, previous works having dealt with the stability of Wasserstein barycenters have either worked under stringent assumptions on the marginal measures or regularized the barycenter problem in order to ensure more regular solutions. In (Ahidar-Coutrix et al., 2020; Le Gouic et al., 2022) for instance, the question of the rate of convergence of the empirical barycenter in a Wasserstein space toward its population counterpart has been answered at the cost of assumptions that require in particular to have guarantees on the regularity of the (unknown) population barycenter (see sub-section 6.1.5 for more details). In (Bigot et al., 2019b; Carlier et al., 2021), a regularization of the barycenter functional has been considered and stability bounds and central limit theorems were deduced for the solutions to these regularized problem. In this work, we do not regularize the barycenter functional and work under less restrictive assumptions on the marginal measures than previous works having dealt with the stability of Wasserstein barycenters. In order to state these assumptions, we first need to introduce the dual problem to $(P)_{\mathbb{P}}$.

6.1.3 Dual formulation

Building from (Agueh and Carlier, 2011), we show that $(P)_{\mathbb{P}}$ admits the following dual formulation with strong duality (the proof is deferred to the last section of this chapter, Section 6.4):

Proposition 6.1 (Dual formulation). *For any $\mathbb{P} \in \mathcal{P}(\mathcal{P}(\Omega))$, problem $(P)_{\mathbb{P}}$ satisfies*

$$(P)_{\mathbb{P}} = \frac{1}{2} \int_{\mathcal{P}(\Omega)} M_2(\rho) d\mathbb{P}(\rho) - (D)_{\mathbb{P}},$$

where $M_2(\rho) = \langle \|\cdot\|^2 | \rho \rangle$ is the second moment of ρ and where $(D)_{\mathbb{P}}$ corresponds to the dual value

$$(D)_{\mathbb{P}} = \min \left\{ \int_{\mathcal{P}(\Omega)} \langle \psi_\rho^* | \rho \rangle d\mathbb{P}(\rho), \quad (\psi_\rho)_\rho \in L^\infty(\mathbb{P}; W^{1,\infty}(\Omega)), \quad \int_{\mathcal{P}(\Omega)} \psi_\rho(\cdot) d\mathbb{P}(\rho) = \frac{\|\cdot\|^2}{2} \right\}.$$

In the expression above, $\psi_\rho^*(\cdot) = \sup_{y \in \Omega} \{\langle \cdot | y \rangle - \psi_\rho(y)\}$ corresponds to the convex conjugate of ψ_ρ and $L^\infty(\mathbb{P}; W^{1,\infty}(\Omega))$ denotes the set of essentially bounded \mathbb{P} -measurable mappings from $\mathcal{P}(\Omega)$ to the Sobolev space $W^{1,\infty}(\Omega)$ of bounded Lipschitz continuous functions from Ω to \mathbb{R} .

Remark 6.2. Note that in the above minimization problem, $(\psi_\rho)_\rho$ is to be understood as the following mapping, defined \mathbb{P} -almost everywhere:

$$(\psi_\rho)_\rho : \begin{cases} \mathcal{P}(\Omega) & \rightarrow W^{1,\infty}(\Omega), \\ \rho & \mapsto \psi_\rho. \end{cases}$$

Remark 6.3. By Kantorovich duality (see Chapter 1 or (Villani, 2008)), for $\mathbb{P} \in \mathcal{P}(\mathcal{P}(\Omega))$, the collection of functions $(\psi_\rho)_\rho$ solving (D) $_{\mathbb{P}}$ give solutions to the (dual) quadratic optimal transport problems between \mathbb{P} -a.e. $\rho \in \mathcal{P}(\Omega)$ and any barycenter $\mu_{\mathbb{P}} \in \arg \min (\mathbf{P})_{\mathbb{P}}$:

$$\begin{aligned} \frac{1}{2}W_2^2(\rho, \mu_{\mathbb{P}}) &= \frac{1}{2}M_2(\rho) + \frac{1}{2}M_2(\mu_{\mathbb{P}}) - (\langle \psi_\rho^* | \rho \rangle + \langle \psi_\rho | \mu_{\mathbb{P}} \rangle) \\ &= \frac{1}{2}M_2(\rho) + \frac{1}{2}M_2(\mu_{\mathbb{P}}) - \left(\min_{\psi \in \mathcal{C}^0(\Omega)} \langle \psi^* | \rho \rangle + \langle \psi | \mu_{\mathbb{P}} \rangle \right). \end{aligned} \quad (6.1)$$

As such, we know from Chapter 1 that $\psi_\rho = \psi_\rho^{**}$ for \mathbb{P} -a.e. ρ , so that this function – that we call later on a (Kantorovich) potential – is convex and Lipschitz continuous with Lipschitz constant smaller than R_Ω . When $\mathbb{P}(\mathcal{P}_{a.c.}(\Omega)) > 0$ and $\rho \in \text{spt}(\mathbb{P}) \cap \mathcal{P}_{a.c.}(\Omega)$, the convex function ψ_ρ^* corresponds to a Brenier potential ((Brenier, 1991) or Theorem 1.12) and its gradients achieves the optimal transport from ρ to the unique barycenter $\mu_{\mathbb{P}}$:

$$(\nabla \psi_\rho^*)_{\#} \rho = \mu_{\mathbb{P}}, \quad \text{and} \quad W_2^2(\rho, \mu_{\mathbb{P}}) = \|\nabla \psi_\rho^* - \text{id}\|_{L^2(\rho, \mathbb{R}^d)}^2.$$

6.1.4 Contributions

As studied in Part I, the minimization problem that appears in (6.1) is convex but it is not, in general, globally strongly-convex. Our main result relies on the assumption that one marginal distribution, say \mathbb{P} , gives positive mass to a set of absolutely continuous measures that are such that problem (6.1) presents a form of local strong-convexity. In particular, for $\rho \in \mathcal{P}(\Omega)$, following Part I and denoting $\mathcal{K}_\rho : \psi \mapsto \langle \psi^* | \rho \rangle$ the associated *Kantorovich functional* appearing in the minimization problem (6.1) (which is convex and whose gradient reads $\nabla \mathcal{K}_\rho(\psi) = -(\nabla \psi^*)_{\#} \rho$ by Lemma 1.11), we will make the following assumption:

Assumption 6.4. The measure $\mathbb{P} \in \mathcal{P}(\mathcal{P}(\Omega))$ is such that there exists constants $\alpha_{\mathbb{P}} > 0$, $c_{\mathbb{P}}, \text{per}_{\mathbb{P}}, m_{\mathbb{P}}, M_{\mathbb{P}} \in (0, +\infty)$ and a measurable set $S_{\mathbb{P}} \subset \mathcal{P}(\Omega)$ verifying $\mathbb{P}(S_{\mathbb{P}}) = \alpha_{\mathbb{P}}$ such that for all $\rho \in S_{\mathbb{P}}$,

- (i) $\rho \in \mathcal{P}_{a.c.}(\Omega)$,
- (ii) $m_{\mathbb{P}} \leq \rho|_{\text{spt}(\rho)} \leq M_{\mathbb{P}}$,
- (iii) $\text{spt}(\rho)$ has a \mathcal{H}^{d-1} -rectifiable boundary and $\mathcal{H}^{d-1}(\partial \text{spt}(\rho)) \leq \text{per}_{\mathbb{P}}$,
- (iv) $\forall \psi, \tilde{\psi} \in \mathcal{C}^0(\Omega), \quad c_{\mathbb{P}} \text{Var}_\rho(\tilde{\psi}^* - \psi^*) \leq \mathcal{K}_\rho(\tilde{\psi}) - \mathcal{K}_\rho(\psi) - \langle \psi - \tilde{\psi} | (\nabla \psi^*)_{\#} \rho \rangle$,

where $\text{spt}(\rho)$ denotes the support of ρ , $\partial \text{spt}(\rho)$ denotes the topological boundary of this support and \mathcal{H}^{d-1} denotes the $(d-1)$ -dimensional Hausdorff measure.

For a population of marginals $\mathbb{P} \in \mathcal{P}(\mathcal{P}(\Omega))$ satisfying Assumption 6.4, we prove that the Wasserstein barycenters depend in a Hölder-continuous way on their marginals near \mathbb{P} :

Theorem 6.5. *Let $\mathbb{P}, \mathbb{Q} \in \mathcal{P}(\mathcal{P}(\Omega))$ and assume that \mathbb{P} satisfies Assumption 6.4. Let $\mu_{\mathbb{P}}$ be the barycenter of \mathbb{P} and $\mu_{\mathbb{Q}}$ be a barycenter of \mathbb{Q} . Then*

$$W_2(\mu_{\mathbb{P}}, \mu_{\mathbb{Q}}) \lesssim \frac{1}{\alpha_{\mathbb{P}}^{1/4}} \mathcal{W}_1(\mathbb{P}, \mathbb{Q})^{1/6},$$

where \lesssim hides the multiplicative constant

$$C_{d, R_{\Omega}, m_{\mathbb{P}}, M_{\mathbb{P}}, \text{per}_{\mathbb{P}}, c_{\mathbb{P}}} = C_d (1 + M_{\mathbb{P}})^{1/4} (1 + R_{\Omega})^{d/4 + 1} \left(1 + \frac{M_{\mathbb{P}}^{1/2} \text{per}_{\mathbb{P}}^{1/3}}{c_{\mathbb{P}}^{1/6} m_{\mathbb{P}}^{1/6}}\right)$$

and C_d is a constant that depends only on d .

Before discussing consequences of this result, we make some comments on our main Assumption 6.4. This assumption mainly corresponds to the assumption that the population \mathbb{P} gives positive mass to a set of marginals ρ that are such that the Kantorovich functionals $\mathcal{K}_{\rho} : \psi \mapsto \langle \psi^* | \rho \rangle$ associated to them satisfies a strong convexity estimate in the sense adopted in Part I. More prosaically, the conditions (i), (ii) and (iii) speak for themselves, and conditions under which (iv) holds are given in Theorem 3.1 or Corollary 1.31. In particular, if a measure $\rho \in \mathcal{P}(\Omega)$ already satisfies (i), (ii) and (iii), then property (iv) will be satisfied at the moment ρ satisfies a Poincaré-Wirtinger inequality and its support is convex or is made of a finite connected union of convex sets. On a more technical side, we note that the Borel measurability of a set $S_{\mathbb{P}} \subset \mathcal{P}(\Omega)$ as defined in Assumption 6.4 needs to be checked depending on the application. Obviously, measurability holds when the number of marginals is finite (\mathbb{P} is discrete).

6.1.5 Consequences of Theorem 6.5

Statistical estimation of barycenter with a finite number of marginals. For a probability measure $\rho \in \mathcal{P}(\Omega)$ and an i.i.d. sequence $(x_j)_{j=1, \dots, n}$ sampled from ρ , it is well-known that the empirical measure $\hat{\rho}^n = \frac{1}{n} \sum_{j=1}^n \delta_{x_j}$ converges weakly to ρ almost-surely as $n \rightarrow \infty$ (Varadarajan, 1958). By Theorem 1 of (Fournier and Guillin, 2015), the rate of this convergence can be controlled in Wasserstein distance: there exists a constant C_d depending only on d such that

$$\mathbb{E}W_2^2(\hat{\rho}^n, \rho) \leq C_d R_{\Omega}^2 \begin{cases} n^{-1/2} & \text{if } d < 4, \\ n^{-1/2} \log(n) & \text{if } d = 4, \\ n^{-2/d} & \text{else,} \end{cases}$$

where the expectation is taken with respect to $(x_j)_{j=1, \dots, n} \sim \rho^{\otimes n}$. Theorem 6.5 together with a double use of Jensen's inequality allows to translate these rates to the statistical estimation of a Wasserstein barycenter with a finite number of marginals:

Corollary 6.6. *Let $\mathbb{P}_m = \sum_{i=1}^m \alpha_i \delta_{\rho_i} \in \mathcal{P}(\mathcal{P}(\Omega))$ satisfying Assumption 6.4. For all $i \in \{1, \dots, m\}$, denote $\hat{\rho}_i^n = \frac{1}{n} \sum_{j=1}^n \delta_{x_{i,j}}$ an empirical measure built from an i.i.d. sequence $(x_{i,j})_{1 \leq j \leq n}$ sampled from ρ_i . Then the barycenters $\mu_{\mathbb{P}_m}$ of \mathbb{P}_m and $\mu_{\hat{\mathbb{P}}_m^n}$ of $\hat{\mathbb{P}}_m^n = \sum_{i=1}^m \alpha_i \delta_{\hat{\rho}_i^n}$ verify*

$$\mathbb{E}W_2^2(\mu_{\hat{\mathbb{P}}_m^n}, \mu_{\mathbb{P}_m}) \lesssim \frac{1}{\alpha_{\mathbb{P}_m}^{1/2}} \begin{cases} n^{-1/12} & \text{if } d < 4, \\ n^{-1/12} \log(n)^{1/6} & \text{if } d = 4, \\ n^{-1/(3d)} & \text{else,} \end{cases}$$

where \lesssim hides a multiplicative constant depending on $d, R_{\Omega}, m_{\mathbb{P}}, M_{\mathbb{P}}, \text{per}_{\mathbb{P}}, c_{\mathbb{P}}$.

Convergence rate of empirical barycenters in Wasserstein spaces. Another statistical question occurs in the setting where the population of marginals $\mathbb{P} \in \mathcal{P}(\mathcal{P}(\Omega))$ is only known through samples $(\rho_i)_{1 \leq i \leq m} \sim \mathbb{P}^{\otimes m}$. Introducing the plug-in estimator $\mathbb{P}_m = \frac{1}{m} \sum_{i=1}^m \delta_{\rho_i}$, it is natural to wonder how well $\mu_{\mathbb{P}_m}$ approaches $\mu_{\mathbb{P}}$ in terms of m . This question, asked in the more general framework of barycenters in Alexandrov spaces, has been the object of recent research (Ahidar-Coutrix et al., 2020; Le Gouic et al., 2022). In Wasserstein spaces, the authors of (Le Gouic et al., 2022) show in particular that $\mathbb{E}W_2(\mu_{\mathbb{P}}, \mu_{\mathbb{P}_m})$ may converge at the parametric rate $m^{-1/2}$ under the assumption that \mathbb{P} admits a barycenter $\mu_{\mathbb{P}}$ that it is such that there exists a bi-Lipschitz optimal transport map between any $\rho \in \text{spt}(\mathbb{P})$ and $\mu_{\mathbb{P}}$, and that the Lipschitz constants of these maps and their inverses do not differ by a value more than 1. Under similar assumptions, the authors of (Chewi et al., 2020) derive a strong-convexity estimate of the barycenter functional at its minimum which helps them derive rates of convergence of gradient descent algorithms for the (stochastic) estimation of barycenters. Such assumptions however require to have strong guarantees on the regularity of a barycenter of \mathbb{P} . These guarantees can be obtained when restricted to specific families of probability measures (e.g. Gaussian measures), but are difficult to get in more general cases. For instance, barycenters of measures with convex support may not have a convex support (Santambrogio and Wang, 2016), which hampers a straightforward use of Caffarelli's regularity theory. In contrast, our stability result entails that for barycenters $\mu_{\mathbb{P}}$ of \mathbb{P} and $\mu_{\mathbb{P}_m}$ of \mathbb{P}_m ,

$$\mathbb{E}W_2(\mu_{\mathbb{P}}, \mu_{\mathbb{P}_m}) \lesssim \frac{1}{\alpha_{\mathbb{P}}^{1/4}} \mathbb{E}W_1(\mathbb{P}, \mathbb{P}_m)^{1/6},$$

whenever \mathbb{P} satisfies Assumption 6.4. This implies that any rate of convergence of $\mathbb{E}W_1(\mathbb{P}, \mathbb{P}_m)$ w.r.t. m is readily transferred to $\mathbb{E}W_2(\mu_{\mathbb{P}}, \mu_{\mathbb{P}_m})$, up to an exponent. However, $\mathcal{P}(\Omega)$ is an infinite dimensional space and there is no general convergence rate for $\mathbb{E}W_1(\mathbb{P}, \mathbb{P}_m)$. Nonetheless, assuming some structure on the population \mathbb{P} may help to derive bounds. One may use for instance the notion of upper Wasserstein dimension of \mathbb{P} introduced in (Weed and Bach, 2019) (Definition 4), defined from quantities depending on the covering numbers of (subsets of) the support of \mathbb{P} . Assuming that this dimension is strictly upper bounded by $s > 0$, the authors of (Weed and Bach, 2019) show that

$$\mathbb{E}W_1(\mathbb{P}, \mathbb{P}_m) \lesssim m^{-1/s},$$

where \lesssim hides a multiplicative constant that depends on R_{Ω} and s .

Error induced by a discretization of the marginals. Let $\rho \in \mathcal{P}(\Omega)$ and let $h > 0$ be a discretization parameter. Denoting $(x_i^h)_{1 \leq i \leq N_h}$ an h -net of Ω and $(V_i^h)_{1 \leq i \leq N_h}$ the corresponding Voronoi tessellation of Ω , it is easy to verify that the discretization $\rho^h = \sum_{i=1}^{N_h} \rho(V_i^h) \delta_{x_i^h}$ verifies

$$W_2(\rho, \rho^h) \leq h.$$

Such kind of discretization, with controlled error bound, may be useful in practice for computational purposes. The stability result of Theorem 6.5 allows to translate the error bound made when discretizing the marginals to the corresponding barycenter:

Corollary 6.7. *Let $\mathbb{P}_m = \sum_{i=1}^m \alpha_i \delta_{\rho_i} \in \mathcal{P}(\mathcal{P}(\Omega))$ satisfying Assumption 6.4. Let $h > 0$ and for all $i \in \{1, \dots, m\}$, denote $\rho_i^h = \sum_{j=1}^{N_h} \rho_i(V_j^h) \delta_{x_j^h}$ a discretization of ρ_i built from*

the h -net $(x_j^h)_{1 \leq j \leq N_h}$. Then the barycenters $\mu_{\mathbb{P}_m}$ of \mathbb{P}_m and $\mu_{\mathbb{P}_m^h}$ of $\mathbb{P}_m^h = \sum_{i=1}^m \alpha_i \delta_{\rho_i^h}$ verify

$$W_2(\mu_{\mathbb{P}_m^h}, \mu_{\mathbb{P}_m}) \lesssim \frac{1}{\alpha_{\mathbb{P}_m}^{1/4}} h^{1/6},$$

where \lesssim hides a multiplicative constant depending on $d, R_\Omega, m_{\mathbb{P}}, M_{\mathbb{P}}, \text{per}_{\mathbb{P}}, c_{\mathbb{P}}$.

6.1.6 Main elements of proof for Theorem 6.5

The derivation of Theorem 6.5 is decomposed into two separate sub-problems. Introduce $(\psi_\rho)_\rho, (\tilde{\psi}_{\tilde{\rho}})_{\tilde{\rho}}$ solutions to the dual problems $(D)_{\mathbb{P}}, (D)_{\mathbb{Q}}$ with respective populations \mathbb{P}, \mathbb{Q} . For an optimal $\gamma \in \Gamma(\mathbb{P}, \mathbb{Q})$ such that $\mathcal{W}_1(\mathbb{P}, \mathbb{Q}) = \int_{\mathcal{P}(\Omega) \times \mathcal{P}(\Omega)} W_2(\rho, \tilde{\rho}) d\gamma(\rho, \tilde{\rho})$, recalling that $\mathbb{P}(S_{\mathbb{P}}) = \alpha_{\mathbb{P}}$, one can notice the following:

$$\begin{aligned} W_2(\mu_{\mathbb{P}}, \mu_{\mathbb{Q}}) &= \frac{1}{\alpha_{\mathbb{P}}} \int_{S_{\mathbb{P}} \times \mathcal{P}(\Omega)} W_2(\mu_{\mathbb{P}}, \mu_{\mathbb{Q}}) d\gamma(\rho, \tilde{\rho}) \\ &\leq \frac{1}{\alpha_{\mathbb{P}}} \int_{S_{\mathbb{P}} \times \mathcal{P}(\Omega)} \left(W_2(\mu_{\mathbb{P}}, (\nabla \tilde{\psi}_{\tilde{\rho}}^*)_{\#} \rho) + W_2((\nabla \tilde{\psi}_{\tilde{\rho}}^*)_{\#} \rho, \mu_{\mathbb{Q}}) \right) d\gamma(\rho, \tilde{\rho}). \end{aligned} \quad (6.2)$$

Because $\rho \in S_{\mathbb{P}}$ is absolutely continuous, we may replace $\mu_{\mathbb{P}}$ by $(\nabla \psi_\rho^*)_{\#} \rho$ in the first term of (6.2), using Remark 6.3. Bounding the first term of (6.2) will thus amount to quantify how $\nabla \psi_\rho^*$ deviates from $\nabla \tilde{\psi}_{\tilde{\rho}}^*$ in terms of $\mathcal{W}_1(\mathbb{P}, \mathbb{Q})$. This corresponds to getting a quantitative estimate on the stability of the solutions to $(D)_{\mathbb{P}}$, for which the local-strong convexity assumption made for measures in $S_{\mathbb{P}}$ almost readily gives the following bound, proven in Section 6.2:

Proposition 6.8. *Let $\mathbb{P}, \mathbb{Q} \in \mathcal{P}(\mathcal{P}(\Omega))$ and assume that \mathbb{P} satisfies Assumption 6.4. Let $\gamma \in \Gamma(\mathbb{P}, \mathbb{Q})$ be such that $\mathcal{W}_1(\mathbb{P}, \mathbb{Q}) = \int_{\mathcal{P}(\Omega) \times \mathcal{P}(\Omega)} W_2(\rho, \tilde{\rho}) d\gamma(\rho, \tilde{\rho})$ and let $(\tilde{\psi}_{\tilde{\rho}})_{\tilde{\rho}}$ be a solution to $(D)_{\mathbb{Q}}$. Then the barycenter $\mu_{\mathbb{P}}$ of \mathbb{P} satisfies*

$$\frac{1}{\alpha_{\mathbb{P}}} \int_{S_{\mathbb{P}} \times \mathcal{P}(\Omega)} W_2(\mu_{\mathbb{P}}, (\nabla \tilde{\psi}_{\tilde{\rho}}^*)_{\#} \rho) d\gamma(\rho, \tilde{\rho}) \lesssim \left(\frac{\mathcal{W}_1(\mathbb{P}, \mathbb{Q})}{\alpha_{\mathbb{P}}} \right)^{1/6},$$

where \lesssim hides the multiplicative constant $C_{d, R_\Omega, m_{\mathbb{P}}, M_{\mathbb{P}}, \text{per}_{\mathbb{P}}, c_{\mathbb{P}}} = \left(\frac{C_d R_\Omega^5 M_{\mathbb{P}}^3 \text{per}_{\mathbb{P}}^2}{m_{\mathbb{P}} c_{\mathbb{P}}} \right)^{1/6}$, and C_d is a constant that depends only on d .

In the second term of (6.2), we may replace *formally* $\mu_{\mathbb{Q}}$ by $(\nabla \tilde{\psi}_{\tilde{\rho}}^*)_{\#} \tilde{\rho}$. Bounding the second term of (6.2) thus amounts to finding a stability estimate for the push-forward operation under the mapping $\nabla \tilde{\psi}_{\tilde{\rho}}^*$. Proposition 6.12 of Section 6.3 gives such estimates and allows to get the following bound.

Proposition 6.9. *Let $\mathbb{P}, \mathbb{Q} \in \mathcal{P}(\mathcal{P}(\Omega))$ and assume that \mathbb{P} satisfies Assumption 6.4. Let $\gamma \in \Gamma(\mathbb{P}, \mathbb{Q})$ be such that $\mathcal{W}_1(\mathbb{P}, \mathbb{Q}) = \int_{\mathcal{P}(\Omega) \times \mathcal{P}(\Omega)} W_2(\rho, \tilde{\rho}) d\gamma(\rho, \tilde{\rho})$ and let $(\tilde{\psi}_{\tilde{\rho}})_{\tilde{\rho}}$ be a solution to $(D)_{\mathbb{Q}}$. Then any barycenter $\mu_{\mathbb{Q}}$ of \mathbb{Q} satisfies*

$$\frac{1}{\alpha_{\mathbb{P}}} \int_{S_{\mathbb{P}} \times \mathcal{P}(\Omega)} W_2((\nabla \tilde{\psi}_{\tilde{\rho}}^*)_{\#} \rho, \mu_{\mathbb{Q}}) d\gamma(\rho, \tilde{\rho}) \lesssim \left(\frac{\mathcal{W}_1(\mathbb{P}, \mathbb{Q})}{\alpha_{\mathbb{P}}} \right)^{1/4},$$

where \lesssim hides the multiplicative constant $C_{d, R_\Omega, M_{\mathbb{P}}} = C_d (1 + M_{\mathbb{P}})^{1/4} (1 + R_\Omega)^{\frac{d+1}{4}}$, and C_d is a constant that depends only on d .

The proof of Theorem 6.5 is then immediate from bound (6.2) and Propositions 6.8 and 6.9.

6.2 Stability of potentials

6.2.1 Lipschitz behavior of the primal and dual values

A first immediate fact is the Lipschitz behavior of the primal and dual values with respect to the marginals \mathbb{P}, \mathbb{Q} , that holds without any assumptions on their regularity.

Proposition 6.10. *For any $\mathbb{P}, \mathbb{Q} \in \mathcal{P}(\mathcal{P}(\Omega))$, the following bounds hold:*

$$\begin{aligned} |(P)_{\mathbb{Q}} - (P)_{\mathbb{P}}| &\leq 3R_{\Omega} \mathcal{W}_1(\mathbb{P}, \mathbb{Q}), \\ |(D)_{\mathbb{Q}} - (D)_{\mathbb{P}}| &\leq 4R_{\Omega} \mathcal{W}_1(\mathbb{P}, \mathbb{Q}). \end{aligned}$$

Proof. Let $\gamma \in \Gamma(\mathbb{P}, \mathbb{Q})$ be such that $\mathcal{W}_1(\mathbb{P}, \mathbb{Q}) = \int_{\mathcal{P}(\Omega) \times \mathcal{P}(\Omega)} W_2(\rho, \tilde{\rho}) d\gamma(\rho, \tilde{\rho})$. By definition of $(P)_{\mathbb{P}}$ and using the triangle inequality we have:

$$\begin{aligned} (P)_{\mathbb{P}} &\leq \frac{1}{2} \int_{\mathcal{P}(\Omega)} W_2^2(\rho, \mu_{\mathbb{Q}}) d\mathbb{P}(\rho) \\ &= \frac{1}{2} \int_{\mathcal{P}(\Omega) \times \mathcal{P}(\Omega)} W_2^2(\rho, \mu_{\mathbb{Q}}) d\gamma(\rho, \tilde{\rho}) \\ &\leq \frac{1}{2} \int_{\mathcal{P}(\Omega) \times \mathcal{P}(\Omega)} (W_2(\rho, \tilde{\rho}) + W_2(\tilde{\rho}, \mu_{\mathbb{Q}}))^2 d\gamma(\rho, \tilde{\rho}) \\ &= \frac{1}{2} \int_{\mathcal{P}(\Omega) \times \mathcal{P}(\Omega)} (W_2^2(\rho, \tilde{\rho}) + 2W_2(\rho, \tilde{\rho})W_2(\tilde{\rho}, \mu_{\mathbb{Q}}) + W_2^2(\tilde{\rho}, \mu_{\mathbb{Q}})) d\gamma(\rho, \tilde{\rho}). \end{aligned}$$

Now using that Ω is compact and included in a ball centered at the origin and of radius $R_{\Omega} > 0$, we have for any $\rho, \tilde{\rho} \in \mathcal{P}(\Omega)$ the upper bound $W_2(\rho, \tilde{\rho}) \leq 2R_{\Omega}$. We thus get the bound

$$\begin{aligned} (P)_{\mathbb{P}} &\leq \frac{1}{2} \int_{\mathcal{P}(\Omega) \times \mathcal{P}(\Omega)} (6R_{\Omega}W_2(\rho, \tilde{\rho}) + W_2^2(\tilde{\rho}, \mu_{\mathbb{Q}})) d\gamma(\rho, \tilde{\rho}) \\ &= 3R_{\Omega} \int_{\mathcal{P}(\Omega) \times \mathcal{P}(\Omega)} W_2(\rho, \tilde{\rho}) d\gamma(\rho, \tilde{\rho}) + \frac{1}{2} \int_{\mathcal{P}(\Omega)} W_2^2(\tilde{\rho}, \mu_{\mathbb{Q}}) d\mathbb{Q}(\tilde{\rho}) \\ &= 3R_{\Omega} \mathcal{W}_1(\mathbb{P}, \mathbb{Q}) + (P)_{\mathbb{Q}}. \end{aligned}$$

The first inequality of the statement is then deduced by symmetry. From this bound on the primal values, we can deduce the following bound on the dual values:

$$\begin{aligned} (D)_{\mathbb{P}} &= \frac{1}{2} \int_{\mathcal{P}(\Omega)} M_2(\rho) d\mathbb{P}(\rho) - (P)_{\mathbb{P}}, \\ &\leq \frac{1}{2} \int_{\mathcal{P}(\Omega)} M_2(\rho) d\mathbb{P}(\rho) - (P)_{\mathbb{Q}} + 3R_{\Omega} \mathcal{W}_1(\mathbb{P}, \mathbb{Q}) \end{aligned}$$

Then, using the triangle inequality:

$$\begin{aligned} \int_{\mathcal{P}(\Omega)} M_2(\rho) d\mathbb{P}(\rho) &= \int_{\mathcal{P}(\Omega) \times \mathcal{P}(\Omega)} M_2(\rho) d\gamma(\rho, \tilde{\rho}) \\ &\leq \int_{\mathcal{P}(\Omega) \times \mathcal{P}(\Omega)} (|M_2(\rho) - M_2(\tilde{\rho})| + M_2(\tilde{\rho})) d\gamma(\rho, \tilde{\rho}) \end{aligned}$$

Notice that for all $\rho, \tilde{\rho} \in \mathcal{P}(\Omega)$, one has

$$\begin{aligned} |M_2(\rho) - M_2(\tilde{\rho})| &= |W_2^2(\rho, \delta_0) - W_2^2(\tilde{\rho}, \delta_0)| \\ &\leq 2R_\Omega |W_2(\rho, \delta_0) - W_2(\tilde{\rho}, \delta_0)| \\ &\leq 2R_\Omega W_2(\rho, \tilde{\rho}). \end{aligned}$$

We thus have $\int_{\mathcal{P}(\Omega)} M_2(\rho) d\mathbb{P}(\rho) \leq 2R_\Omega \mathcal{W}_1(\mathbb{P}, \mathbb{Q}) + \int_{\mathcal{P}(\Omega)} M_2(\tilde{\rho}) d\mathbb{Q}(\tilde{\rho})$, from which we deduce:

$$(\mathbf{D})_{\mathbb{P}} \leq \frac{1}{2} \int_{\mathcal{P}(\Omega)} M_2(\tilde{\rho}) d\mathbb{Q}(\tilde{\rho}) - (\mathbf{P})_{\mathbb{Q}} + 4R_\Omega \mathcal{W}_1(\mathbb{P}, \mathbb{Q}) = (\mathbf{D})_{\mathbb{Q}} + 4R_\Omega \mathcal{W}_1(\mathbb{P}, \mathbb{Q}).$$

The second inequality of the statement then follows by symmetry. \square

6.2.2 Stability of the dual solutions

Denote $(\psi_\rho)_\rho$ and $(\tilde{\psi}_{\tilde{\rho}})_{\tilde{\rho}}$ solutions to the dual problems $(\mathbf{D})_{\mathbb{P}}, (\mathbf{D})_{\mathbb{Q}}$ with respective populations \mathbb{P}, \mathbb{Q} . These families of potentials verify

$$\begin{aligned} (\mathbf{D})_{\mathbb{P}} &= \int_{\mathcal{P}(\Omega)} \mathcal{K}_\rho(\psi_\rho) d\mathbb{P}(\rho) \quad \text{and} \quad \int_{\mathcal{P}(\Omega)} \psi_\rho(\cdot) d\mathbb{P}(\rho) = \frac{\|\cdot\|^2}{2}, \\ (\mathbf{D})_{\mathbb{Q}} &= \int_{\mathcal{P}(\Omega)} \mathcal{K}_{\tilde{\rho}}(\tilde{\psi}_{\tilde{\rho}}) d\mathbb{Q}(\tilde{\rho}) \quad \text{and} \quad \int_{\mathcal{P}(\Omega)} \tilde{\psi}_{\tilde{\rho}}(\cdot) d\mathbb{Q}(\tilde{\rho}) = \frac{\|\cdot\|^2}{2}. \end{aligned}$$

We can then show that $(\tilde{\psi}_{\tilde{\rho}})_{\tilde{\rho}}$ is almost a minimizer for $(\mathbf{D})_{\mathbb{P}}$ whenever $\mathcal{W}_1(\mathbb{P}, \mathbb{Q})$ is small:

Proposition 6.11. *Let $\mathbb{P}, \mathbb{Q} \in \mathcal{P}(\mathcal{P}(\Omega))$ and $(\psi_\rho)_\rho, (\tilde{\psi}_{\tilde{\rho}})_{\tilde{\rho}}$ respective solutions of the dual problems $(\mathbf{D})_{\mathbb{P}}, (\mathbf{D})_{\mathbb{Q}}$. Let $\gamma \in \Gamma(\mathbb{P}, \mathbb{Q})$ be such that $\mathcal{W}_1(\mathbb{P}, \mathbb{Q}) = \int_{\mathcal{P}(\Omega) \times \mathcal{P}(\Omega)} W_2(\rho, \tilde{\rho}) d\gamma(\rho, \tilde{\rho})$. Then*

$$\left| \int_{\mathcal{P}(\Omega) \times \mathcal{P}(\Omega)} (\mathcal{K}_\rho(\psi_\rho) - \mathcal{K}_\rho(\tilde{\psi}_{\tilde{\rho}})) d\gamma(\rho, \tilde{\rho}) \right| \leq 5R_\Omega \mathcal{W}_1(\mathbb{P}, \mathbb{Q}).$$

Proof. Using the triangle inequality, we have

$$\begin{aligned} \left| \int_{\mathcal{P}(\Omega) \times \mathcal{P}(\Omega)} (\mathcal{K}_\rho(\psi_\rho) - \mathcal{K}_\rho(\tilde{\psi}_{\tilde{\rho}})) d\gamma(\rho, \tilde{\rho}) \right| &\leq \left| \int_{\mathcal{P}(\Omega) \times \mathcal{P}(\Omega)} (\mathcal{K}_\rho(\psi_\rho) - \mathcal{K}_{\tilde{\rho}}(\tilde{\psi}_{\tilde{\rho}})) d\gamma(\rho, \tilde{\rho}) \right| \\ &\quad + \left| \int_{\mathcal{P}(\Omega) \times \mathcal{P}(\Omega)} (\mathcal{K}_{\tilde{\rho}}(\tilde{\psi}_{\tilde{\rho}}) - \mathcal{K}_\rho(\tilde{\psi}_{\tilde{\rho}})) d\gamma(\rho, \tilde{\rho}) \right|. \end{aligned} \quad (6.3)$$

Using that $\gamma \in \Gamma(\mathbb{P}, \mathbb{Q})$, the first term of sum (6.3) is equal to

$$\left| \int_{\mathcal{P}(\Omega)} \mathcal{K}_\rho(\psi_\rho) d\mathbb{P}(\rho) - \int_{\mathcal{P}(\Omega)} \mathcal{K}_{\tilde{\rho}}(\tilde{\psi}_{\tilde{\rho}}) d\mathbb{Q}(\tilde{\rho}) \right| = |(\mathbf{D})_{\mathbb{P}} - (\mathbf{D})_{\mathbb{Q}}|,$$

which can be upper bounded by $4R_\Omega \mathcal{W}_1(\mathbb{P}, \mathbb{Q})$ using Proposition 6.10. The second term in (6.3) can then be upper bounded by $R_\Omega \mathcal{W}_1(\mathbb{P}, \mathbb{Q})$ using the Kantorovich-Rubinstein duality result (Proposition A.8), which ensures that for any $\rho, \tilde{\rho} \in \mathcal{P}(\Omega)$, since $\tilde{\psi}_{\tilde{\rho}}^*$ is R_Ω -Lipschitz,

$$\mathcal{K}_\rho(\tilde{\psi}_{\tilde{\rho}}) - \mathcal{K}_{\tilde{\rho}}(\tilde{\psi}_{\tilde{\rho}}) = \langle \tilde{\psi}_{\tilde{\rho}}^* | \rho - \tilde{\rho} \rangle \leq R_\Omega W_1(\rho, \tilde{\rho}) \leq R_\Omega W_2(\rho, \tilde{\rho}). \quad \square$$

6.2.3 Proof of Proposition 6.8

Proof of Proposition 6.8. Let $\rho \in \text{spt}(\mathbb{P})$ and $\pi \in \Gamma(\rho, \mu_{\mathbb{P}})$ be an optimal coupling for the quadratic optimal transport between ρ and $\mu_{\mathbb{P}}$. Then by (6.1) we have:

$$\int_{\Omega \times \Omega} \langle x|y \rangle d\pi(x, y) = \mathcal{K}_{\rho}(\psi_{\rho}) + \langle \psi_{\rho} | \mu_{\mathbb{P}} \rangle. \quad (6.4)$$

On the other hand, the Fenchel-Young inequality ensures that for any $\tilde{\rho} \in \text{spt}(\mathbb{Q})$, for any pair of points $x, y \in \Omega$,

$$\langle x|y \rangle \leq \tilde{\psi}_{\tilde{\rho}}^*(x) + \tilde{\psi}_{\tilde{\rho}}(y).$$

This ensures

$$\int_{\Omega \times \Omega} \langle x|y \rangle d\pi(x, y) \leq \mathcal{K}_{\rho}(\tilde{\psi}_{\tilde{\rho}}) + \langle \tilde{\psi}_{\tilde{\rho}} | \mu_{\mathbb{P}} \rangle.$$

Injecting (6.4) into this last inequality then yields

$$\langle \psi_{\rho} - \tilde{\psi}_{\tilde{\rho}} | \mu_{\mathbb{P}} \rangle \leq \mathcal{K}_{\rho}(\tilde{\psi}_{\tilde{\rho}}) - \mathcal{K}_{\rho}(\psi_{\rho}). \quad (6.5)$$

When ρ belongs to $S_{\mathbb{P}}$, Assumption 6.4 allows us to improve the previous bound:

$$\langle \psi_{\rho} - \tilde{\psi}_{\tilde{\rho}} | \mu_{\mathbb{P}} \rangle + c_{\mathbb{P}} \text{Var}_{\rho}(\tilde{\psi}_{\tilde{\rho}}^* - \psi_{\rho}^*) \leq \mathcal{K}_{\rho}(\tilde{\psi}_{\tilde{\rho}}) - \mathcal{K}_{\rho}(\psi_{\rho}). \quad (6.6)$$

Hence, weighting (6.5) and (6.6) by $d\gamma(\rho, \tilde{\rho})$ and summing over $(\rho, \tilde{\rho}) \in \mathcal{P}(\Omega) \times \mathcal{P}(\Omega)$, we get

$$\begin{aligned} \left\langle \int_{\mathcal{P}(\Omega)} \psi_{\rho} d\mathbb{P}(\rho) - \int_{\mathcal{P}(\Omega)} \tilde{\psi}_{\tilde{\rho}} d\mathbb{Q}(\tilde{\rho}) | \mu_{\mathbb{P}} \right\rangle + c_{\mathbb{P}} \int_{S_{\mathbb{P}} \times \mathcal{P}(\Omega)} \text{Var}_{\rho}(\tilde{\psi}_{\tilde{\rho}}^* - \psi_{\rho}^*) d\gamma(\rho, \tilde{\rho}) \\ \leq \int_{\mathcal{P}(\Omega) \times \mathcal{P}(\Omega)} (\mathcal{K}_{\rho}(\tilde{\psi}_{\tilde{\rho}}) - \mathcal{K}_{\rho}(\psi_{\rho})) d\gamma(\rho, \tilde{\rho}). \end{aligned}$$

Using that $\int_{\mathcal{P}(\Omega)} \psi_{\rho}(\cdot) d\mathbb{P}(\rho) = \int_{\mathcal{P}(\Omega)} \tilde{\psi}_{\tilde{\rho}}(\cdot) d\mathbb{Q}(\tilde{\rho}) = \frac{\|\cdot\|^2}{2}$ and Proposition 6.11, this leads to

$$c_{\mathbb{P}} \int_{S_{\mathbb{P}} \times \mathcal{P}(\Omega)} \text{Var}_{\rho}(\tilde{\psi}_{\tilde{\rho}}^* - \psi_{\rho}^*) d\gamma(\rho, \tilde{\rho}) \leq 5R_{\Omega} \mathcal{W}_1(\mathbb{P}, \mathbb{Q}). \quad (6.7)$$

We now recall the Galiardo-Nirenberg type inequality of Proposition 5.11, that ensures that there exists a constant C_d depending only on d such that for any compact domain K of \mathbb{R}^d with \mathcal{H}^{d-1} -rectifiable boundary and $u, v : K \rightarrow \mathbb{R}$ two Lipschitz functions on K that are convex on any segment included in K ,

$$\|\nabla u - \nabla v\|_{L^2(K)}^6 \leq C_d \mathcal{H}^{d-1}(\partial K)^2 (\|\nabla u\|_{L^{\infty}(K)} + \|\nabla v\|_{L^{\infty}(K)})^4 \|u - v\|_{L^2(K)}^2.$$

This inequality can be used to bound from below the left-hand side term of equation (6.7). Indeed, measures belonging to $S_{\mathbb{P}}$ are assumed to have a support which has a \mathcal{H}^{d-1} -rectifiable topological boundary. This ensures the following for any $\rho \in S_{\mathbb{P}}$ and $\tilde{\rho} \in \mathcal{P}(\Omega)$:

$$\left\| \nabla \tilde{\psi}_{\tilde{\rho}}^* - \nabla \psi_{\rho}^* \right\|_{L^2(\rho)}^6 \leq C_d \text{per}_{\mathbb{P}}^2 R_{\Omega}^4 \frac{M_{\mathbb{P}}^3}{m_{\mathbb{P}}} \text{Var}_{\rho}(\tilde{\psi}_{\tilde{\rho}}^* - \psi_{\rho}^*), \quad (6.8)$$

where C_d is a constant that only depends on the dimension d . Therefore, using that for any $\rho \in \mathcal{S}_{\mathbb{P}}$ and $\tilde{\rho} \in \mathcal{P}(\Omega)$,

$$\mathbb{W}_2(\mu_{\mathbb{P}}, (\nabla \tilde{\psi}_{\tilde{\rho}}^*)_{\#}\rho) = \mathbb{W}_2((\nabla \psi_{\rho}^*)_{\#}\rho, (\nabla \tilde{\psi}_{\tilde{\rho}}^*)_{\#}\rho) \leq \left\| \nabla \tilde{\psi}_{\tilde{\rho}}^* - \nabla \psi_{\rho}^* \right\|_{L^2(\rho)},$$

the combination of bounds (6.7) and (6.8) with Jensen's inequality yields:

$$\frac{1}{\alpha_{\mathbb{P}}} \int_{\mathcal{S}_{\mathbb{P}} \times \mathcal{P}(\Omega)} \mathbb{W}_2(\mu_{\mathbb{P}}, (\nabla \tilde{\psi}_{\tilde{\rho}}^*)_{\#}\rho) d\gamma(\rho, \tilde{\rho}) \leq \left(\frac{C_d R_{\Omega}^5 M_{\mathbb{P}}^3 \text{per}_{\mathbb{P}}^2}{m_{\mathbb{P}} c_{\mathbb{P}}} \right)^{1/6} \left(\frac{\mathcal{W}_1(\mathbb{P}, \mathbb{Q})}{\alpha_{\mathbb{P}}} \right)^{1/6}. \quad \square$$

6.3 Stability of push-forwards

We prove the following proposition, which is more general than Proposition 6.9, and which entails the latter by a direct use of Jensen's inequality. In this statement, $p_1 : (x, y) \mapsto x$ and $p_2 : (x, y) \mapsto y$ are the projections onto the first and second coordinates respectively.

Proposition 6.12. *Let $\rho, \tilde{\rho} \in \mathcal{P}(\Omega)$ and assume that ρ is absolutely continuous with density bounded above by $M_{\rho} \in (0, +\infty)$. Let $\phi \in \mathcal{C}^0(\Omega)$ be a convex R_{Ω} -Lipschitz function. Let $\tilde{\gamma} \in \mathcal{P}(\Omega \times \Omega)$ be such that $(p_1)_{\#}\tilde{\gamma} = \tilde{\rho}$ and assume that $\tilde{\gamma}$ is concentrated on*

$$\partial\phi = \{(x, y) | \phi(x) + \phi^*(y) = \langle x | y \rangle\}.$$

Then

$$\mathbb{W}_2((\nabla\phi)_{\#}\rho, (p_2)_{\#}\tilde{\gamma}) \leq C_{d, R_{\Omega}, M_{\rho}} \mathbb{W}_2(\rho, \tilde{\rho})^{1/4},$$

where $C_{d, R_{\Omega}, M_{\rho}} = C_d(1 + M_{\rho})(1 + R_{\Omega})^{d+1}$, with C_d a constant that depends only on d .

Remark 6.13. Whenever ϕ is differentiable $\tilde{\rho}$ -almost-everywhere, Proposition 6.12 ensures the following stability result for the push-forward operation by $\nabla\phi$:

$$\mathbb{W}_2((\nabla\phi)_{\#}\rho, (\nabla\phi)_{\#}\tilde{\rho}) \leq C_{d, R_{\Omega}, L, M_{\rho}} \mathbb{W}_2(\rho, \tilde{\rho})^{1/4}.$$

We will rely on the following lemma, whose proof is deferred to the end of this section. Note that this proof heavily relies on the proof of Lemma 3.2 of (Carlier et al., 2021).

Lemma 6.14. *Let ϕ be a convex Lipschitz function over \mathbb{R}^d . Then for any $x \in \mathbb{R}^d$ and $r > 0$,*

$$\text{diam}(\partial\phi(B(x, r))) \leq \frac{12}{\omega_d r^d} \|\nabla\phi\|_{L^1(B(x, 4r))},$$

where ω_d denotes the volume of the unit ball of \mathbb{R}^d .

We are now ready to prove Proposition 6.12.

Proof of Proposition 6.12. Denote $\tilde{\gamma} = \tilde{\gamma}_x \otimes \tilde{\rho}$ the disintegration of $\tilde{\gamma}$ with respect to $\tilde{\rho}$, i.e. the collection of measures $(\tilde{\gamma}_x)_{x \in \Omega}$ that satisfy for any function $\xi : \Omega \times \Omega \rightarrow \mathbb{R}$,

$$\int_{\Omega \times \Omega} \xi(x, y) d\tilde{\gamma}(x, y) = \int_{\Omega} \int_{\Omega} \xi(x, y) d\tilde{\gamma}_x(y) d\tilde{\rho}(x).$$

Notice that for $\tilde{\rho}$ -almost-every $x \in \Omega$, if $y \in \text{spt}(\tilde{\gamma}_x)$, then $y \in \partial\phi(x)$ by assumption on $\tilde{\gamma}$. Introduce $S : \mathbb{R}^d \rightarrow \mathbb{R}^d$ the optimal transport map from ρ to $\tilde{\rho}$ and the measure $\gamma = \tilde{\gamma}_{S(x)} \otimes \rho$, that satisfies for any function $\xi : \Omega \times \Omega \rightarrow \mathbb{R}$,

$$\int_{\Omega \times \Omega} \xi(x, y) d\gamma(x, y) = \int_{\Omega} \int_{\Omega} \xi(x, y) d\tilde{\gamma}_{S(x)}(y) d\rho(x).$$

The measure γ is a coupling between ρ and $(p_2)_{\#}\tilde{\gamma}$, which implies that $(\nabla\phi, \text{id})_{\#}\gamma$ is a coupling between $(\nabla\phi)_{\#}\rho$ and $(p_2)_{\#}\tilde{\gamma}$. We therefore have the bound:

$$\begin{aligned} W_2^2((\nabla\phi)_{\#}\rho, (p_2)_{\#}\tilde{\gamma}) &\leq \int_{\Omega \times \Omega} \|\nabla\phi(x) - y\|^2 d\gamma(x, y) \\ &= \int_{\Omega} \int_{\Omega} \|\nabla\phi(x) - y\|^2 d\tilde{\gamma}_{S(x)}(y) d\rho(x) \\ &= \int_{x \in \Omega} \int_{y \in \partial\phi(S(x))} \|\nabla\phi(x) - y\|^2 d\tilde{\gamma}_{S(x)}(y) d\rho(x), \end{aligned}$$

where to get to the last line, we used that for ρ -almost-every $x \in \Omega$, if $y \in \text{spt}(\tilde{\gamma}_{S(x)})$ then $y \in \partial\phi(S(x))$. For $\eta \in (0, (2R_{\Omega})^{3/4}]$, we will find an upper bound on the right-hand side by splitting the integral over Ω into two integrals: one on Ω_{η} and one on Ω_{η}^c , where

$$\Omega_{\eta} = \{x \in \text{spt}(\rho) \mid \|S(x) - x\|^2 \geq \eta^2\}, \quad \Omega_{\eta}^c = \text{spt}(\rho) \setminus \Omega_{\eta}.$$

Upper bound on Ω_{η} . The optimal transport map S from ρ to $\tilde{\rho}$ satisfies

$$\|S - \text{id}\|_{L^2(\rho)} = W_2(\rho, \tilde{\rho}).$$

Then by Markov's inequality, $\rho(\Omega_{\eta}) \leq \frac{W_2^2(\rho, \tilde{\rho})}{\eta^2}$, so that

$$\begin{aligned} \int_{x \in \Omega_{\eta}} \int_{y \in \partial\phi(S(x))} \|\nabla\phi(x) - y\|^2 d\tilde{\gamma}_{S(x)}(y) d\rho(x) &\leq \int_{x \in \Omega_{\eta}} 4R_{\Omega}^2 d\rho(x) \\ &\leq \frac{4R_{\Omega}^2}{\eta^2} W_2^2(\rho, \tilde{\rho}). \end{aligned} \quad (6.9)$$

Upper bound on Ω_{η}^c . By definition of Ω_{η} , for any $x \in \Omega_{\eta}^c$, $\|S(x) - x\| \leq \eta$, i.e. $S(x) \in B(x, \eta)$. Then for any such x , $\partial\phi(S(x)) \subset \partial\phi(B(x, \eta))$. Therefore for any $g \in \partial\phi(x)$ and $y \in \partial\phi(S(x))$, one has

$$\|g - y\| \leq \text{diam}(\partial\phi(B(x, \eta))),$$

which leads to

$$\int_{x \in \Omega_{\eta}^c} \int_{y \in \partial\phi(S(x))} \|\nabla\phi(x) - y\|^2 d\tilde{\gamma}_{S(x)}(y) d\rho(x) \leq \int_{x \in \Omega_{\eta}^c} \text{diam}(\partial\phi(B(x, \eta)))^2 d\rho(x). \quad (6.10)$$

For $\alpha > 0$, introduce $\mathcal{X}_{\alpha} = \{x \in \Omega_{\eta}^c \mid \text{diam}(\partial\phi(B(x, \eta))) \geq \eta^{\alpha}\}$. We will quantify $\rho(\mathcal{X}_{\alpha})$ by finding an upper bound on its covering number. Let $(\mathcal{X}_{\alpha})_{4\eta}^{\text{pack}} \subset \mathcal{X}_{\alpha}$ be a maximal (4η) -packing of \mathcal{X}_{α} , i.e. a finite subset of \mathcal{X}_{α} such that for any $x, y \in (\mathcal{X}_{\alpha})_{4\eta}^{\text{pack}}$, $x \neq y$, $B(x, 4\eta) \cap B(y, 4\eta) = \emptyset$ and such that for any $z \in \mathcal{X}_{\alpha}$, there exists $x \in (\mathcal{X}_{\alpha})_{4\eta}^{\text{pack}}$

such that $B(z, 4\eta) \cap B(x, 4\eta) \neq \emptyset$. We denote $N_{4\eta}^{pack}(\mathcal{X}_\alpha)$ the cardinal of $(\mathcal{X}_\alpha)_{4\eta}^{pack}$. By definition of \mathcal{X}_α , for any $x \in (\mathcal{X}_\alpha)_{4\eta}^{pack} \subset \mathcal{X}_\alpha$, one has

$$\eta^\alpha \leq \text{diam}(\partial\phi(B(x, \eta))).$$

Lemma 6.14 then ensures that for any $c \in \mathbb{R}^d$,

$$\eta^\alpha \leq \frac{12}{\omega_d \eta^d} \|\nabla\phi - c\|_{L^1(B(x, 4\eta))}. \quad (6.11)$$

Choosing $c = \frac{1}{|B(x, 4\eta)|} \int_{B(x, 4\eta)} \nabla\phi(u) du$, the Poincaré-Wirtinger inequality on $B(x, 4\eta)$ ensures

$$\|\nabla\phi - c\|_{L^1(B(x, 4\eta))} \leq 4\eta \int_{B(x, 4\eta)} \|D^2\phi(u)\|_{1,1} du.$$

Using that for any positive semi-definite $d \times d$ matrix M , $\|M\|_{1,1} \leq \text{dtr}(M)$, we then have

$$\|\nabla\phi - c\|_{L^1(B(x, 4\eta))} \leq 4\eta d \int_{B(x, 4\eta)} \Delta\phi(u) du,$$

where Δ stands for the Laplace operator. Injecting this last bound into (6.11) yields

$$\eta^\alpha \lesssim \frac{1}{\eta^{d-1}} \int_{B(x, 4\eta)} \Delta\phi(u) du,$$

where \lesssim hides multiplicative constants depending on d, M_ρ and R_Ω . Summing the last bound over $x \in (\mathcal{X}_\alpha)_{4\eta}^{pack}$ then yields

$$\begin{aligned} N_{4\eta}^{pack}(\mathcal{X}_\alpha) &\lesssim \eta^{1-d-\alpha} \sum_{x \in (\mathcal{X}_\alpha)_{4\eta}^{pack}} \int_{B(x, 4\eta)} \Delta\phi(u) du \\ &\lesssim \eta^{1-d-\alpha} \int_{\Omega+B(0, 4\eta)} \Delta\phi(u) du \\ &\lesssim \eta^{1-d-\alpha} \int_{\partial(\Omega+B(0, 4\eta))} \langle \nabla\phi(u) | n_u \rangle du, \end{aligned}$$

where n_u is the outward pointing unit normal at $u \in \partial(\Omega + B(0, 4\eta))$. Using that $\|\partial\phi\| \leq R_\Omega$, one then has

$$N_{4\eta}^{pack}(\mathcal{X}_\alpha) \lesssim \eta^{1-d-\alpha}.$$

We can now upper bound the (8η) -covering number of \mathcal{X}_α , denoted $N_{8\eta}^{cov}(\mathcal{X}_\alpha)$, with its (4η) -packing number:

$$N_{8\eta}^{cov}(\mathcal{X}_\alpha) \leq N_{4\eta}^{pack}(\mathcal{X}_\alpha).$$

Hence for any minimal (8η) -covering of \mathcal{X}_α denoted by $(\mathcal{X}_\alpha)_{8\eta}^{cov}$, one has

$$\rho(\mathcal{X}_\alpha) \leq \rho \left(\bigcup_{x \in (\mathcal{X}_\alpha)_{8\eta}^{cov}} B(x, 8\eta) \right) \leq N_{8\eta}^{cov}(\mathcal{X}_\alpha) M_\rho \omega_d (8\eta)^d \lesssim \eta^{1-\alpha}.$$

We can finally write

$$\begin{aligned}
\int_{x \in \Omega_\eta^c} \text{diam}(\partial\phi(B(x, \eta)))^2 d\rho(x) &= \int_{x \in \Omega_\eta^c \cap \mathcal{X}_\alpha} \text{diam}(\partial\phi(B(x, \eta)))^2 d\rho(x) \\
&\quad + \int_{x \in \Omega_\eta^c \cap \mathcal{X}_\alpha^c} \text{diam}(\partial\phi(B(x, \eta)))^2 d\rho(x) \\
&\leq 4(R_\Omega)^2 \rho(\mathcal{X}_\alpha) + \int_{\Omega_\eta^c \cap \mathcal{X}_\alpha^c} \eta^{2\alpha} d\rho(x) \\
&\lesssim \eta^{1-\alpha} + \eta^{2\alpha}.
\end{aligned}$$

Choosing $\alpha = \frac{1}{3}$ then yields in (6.10):

$$\int_{x \in \Omega_\eta^c} \int_{y \in \partial\phi(S(x))} \|\nabla\phi(x) - y\|^2 d\tilde{\gamma}_{S(x)}(y) d\rho(x) \lesssim \eta^{2/3}.$$

Combining this last bound with (6.9) leads to

$$W_2^2((\nabla\phi)_\# \rho, (p_2)_\# \tilde{\gamma}) \lesssim \frac{1}{\eta^2} W_2^2(\rho, \tilde{\rho}) + \eta^{2/3}.$$

Setting $\eta = W_2(\rho, \tilde{\rho})^{3/4}$ then gives us

$$W_2^2((\nabla\phi)_\# \rho, (p_2)_\# \tilde{\gamma}) \lesssim W_2(\rho, \tilde{\rho})^{1/2}. \quad \square$$

We now prove Lemma 6.14, which is key to the proof of Proposition 6.12.

Proof of Lemma 6.14. Let $x \in \mathbb{R}^d$ and $r > 0$. One has by definition:

$$\begin{aligned}
\text{diam}(\partial\phi(B(x, r))) &= \max_{y, y' \in B(x, r)} \max_{g \in \partial\phi(y), g' \in \partial\phi(y')} \|g - g'\| \\
&\leq \max_{y, y' \in B(x, r)} \max_{g \in \partial\phi(y), g' \in \partial\phi(y')} \|g\| + \|g'\| \\
&= 2 \max_{y \in B(x, r)} \max_{g \in \partial\phi(y)} \|g\| \\
&= 2 \|\partial\phi\|_{L^\infty(B(x, r))}.
\end{aligned}$$

But for any $y, y' \in \mathbb{R}^d$ and $g \in \partial\phi(y)$, the convexity of ϕ entails

$$\langle g | y' - y \rangle \leq |\phi(y') - \phi(y)|.$$

Therefore, choosing $y \in B(x, r)$ and $g \in \partial\phi(y)$ such that $\|\partial\phi\|_{L^\infty(B(x, r))} = \|g\|$, one has for $y' = y + r \frac{g}{\|g\|} \in B(y, r) \subset B(x, 2r)$ the following bound:

$$r \|g\| \leq |\phi(y') - \phi(y)| \leq \text{osc}_{B(x, 2r)}(\phi),$$

where $\text{osc}_K(f) = \sup_{u, v \in K} |f(u) - f(v)|$. We thus have shown

$$\text{diam}(\partial\phi(B(x, r))) \leq \frac{2}{r} \text{osc}_{B(x, 2r)}(\phi). \quad (6.12)$$

We conclude exactly as in the proof of Lemma 3.2 of (Carlier et al., 2021), that we report here only for completeness: let $y_0 \in \arg \min_{B(x, 2r)} \phi$, $y_1 \in \arg \max_{B(x, 2r)} \phi$, $g_1 \in \partial\phi(y_1)$. Then by convexity of ϕ , for any $y \in \mathbb{R}^d$ and $g \in \partial\phi(y)$ one has

$$\phi(y_1) + \langle g_1 | y - y_1 \rangle \leq \phi(y) \leq \phi(y_0) + \langle g | y - y_0 \rangle.$$

It follows that

$$\|g\| \geq \frac{\text{osc}_{B(x,2r)}(\phi) + \langle g_1 | y - y_1 \rangle}{\|y - y_0\|}.$$

Introducing $W_r(y_1, g_1) = \{y \in B(y_1, 2r) \mid \langle g_1 | y - y_1 \rangle \geq 0\} \subset B(x, 4r)$, one then has

$$\begin{aligned} \|\nabla\phi\|_{L^1(B(x,4r))} &\geq \int_{W_r(y_1, g_1)} \|\nabla\phi\| \, dy \\ &\geq \int_{W_r(y_1, g_1)} \frac{\text{osc}_{B(x,2r)}(\phi)}{\|y - y_0\|} \, dy \\ &\geq \text{osc}_{B(x,2r)}(\phi) \int_{W_r(y_1, g_1)} \frac{1}{\|y - y_1\| + \|y_1 - y_0\|} \, dy \\ &\geq \frac{\text{osc}_{B(x,2r)}(\phi)}{6r} \int_{B(y_1 + r \frac{g_1}{\|g_1\|}, r)} \, dy \\ &= \omega_d \frac{r^{d-1}}{6} \text{osc}_{B(x,2r)}(\phi), \end{aligned}$$

where ω_d denotes the volume of the unit ball of \mathbb{R}^d and where we used the fact that $B(y_1 + r \frac{g_1}{\|g_1\|}, r) \subset W_r(y_1, g_1)$. Plugging this last bound into (6.12) finally yields

$$\text{diam}(\partial\phi(B(x, r))) \leq \frac{12}{\omega_d r^d} \|\nabla\phi\|_{L^1(B(x,4r))}. \quad \square$$

6.4 Dual formulation

In this final section, we justify the dual formulation we claimed in Proposition 6.1. Note that we use a very similar approach to the one we used to prove Kantorovich duality for quadratic optimal transport in Theorem 1.2 of Chapter 1.

Proof of Proposition 6.1. Instead of showing directly the formulation of Proposition 6.1, we will rather show

$$(\mathbf{P})_{\mathbb{P}} = \max \left\{ \int_{\mathcal{P}(\Omega)} \langle \phi_{\rho}^c | \rho \rangle \, d\mathbb{P}(\rho), \quad (\phi_{\rho})_{\rho} \in L^{\infty}(\mathbb{P}; W^{1,\infty}(\Omega)), \quad \int_{\mathcal{P}(\Omega)} \phi_{\rho}(\cdot) \, d\mathbb{P}(\rho) = 0 \right\},$$

where for any $\rho \in \mathcal{P}(\Omega)$, ϕ_{ρ}^c denotes the following c -transform of ϕ_{ρ} :

$$\forall x \in \Omega, \quad \phi_{\rho}^c(x) = \inf_{y \in \Omega} \frac{1}{2} \|x - y\|^2 - \phi_{\rho}(y).$$

Such a formulation entails the result of Proposition 6.1 by the change of variable

$$(\psi_{\rho})_{\rho} = \left(\frac{\|\cdot\|^2}{2} - \phi_{\rho} \right)_{\rho} \in L^{\infty}(\mathbb{P}; W^{1,\infty}(\Omega)).$$

Duality. Let's first show that $(\mathbf{P})_{\mathbb{P}}$ is equal to the value of the following supremum

$$\widetilde{(\mathbf{P})}_{\mathbb{P}} := \sup \left\{ \int_{\mathcal{P}(\Omega)} \langle \phi_{\rho}^c | \rho \rangle \, d\mathbb{P}(\rho), \quad (\phi_{\rho})_{\rho} \in L^1(\mathbb{P}; C^0(\Omega)), \quad \int_{\mathcal{P}(\Omega)} \phi_{\rho}(\cdot) \, d\mathbb{P}(\rho) = 0 \right\},$$

where $L^1(\mathbb{P}; \mathcal{C}^0(\Omega))$ denotes the set of \mathbb{P} -measurable and Bochner integrable mappings from $\mathcal{P}(\Omega)$ to the space $(\mathcal{C}^0(\Omega), \|\cdot\|_\infty)$ of continuous function from Ω to \mathbb{R} equipped with the supremum norm. Introduce the functional $H : \mathcal{C}^0(\Omega) \rightarrow \mathbb{R}$ defined for all $\varphi \in \mathcal{C}^0(\Omega)$ by

$$H(\varphi) = \inf \left\{ - \int_{\mathcal{P}(\Omega)} \langle \phi_\rho^c | \rho \rangle d\mathbb{P}(\rho), \quad (\phi_\rho)_\rho \in L^1(\mathbb{P}; \mathcal{C}^0(\Omega)), \quad \int_{\mathcal{P}(\Omega)} \phi_\rho(\cdot) d\mathbb{P}(\rho) = \varphi(\cdot) \right\}.$$

Notice then that $\widetilde{(\mathbf{P})}_{\mathbb{P}} = -H(0)$. On the other hand, notice that H has the following convex conjugate: for $\nu \in \mathcal{P}(\Omega)$,

$$\begin{aligned} H^*(\nu) &= \sup \{ \langle \varphi | \nu \rangle - H(\varphi), \quad \varphi \in \mathcal{C}^0(\Omega) \} \\ &= \sup \left\{ \langle \varphi | \nu \rangle + \int_{\mathcal{P}(\Omega)} \langle \phi_\rho^c | \rho \rangle d\mathbb{P}(\rho), \quad \varphi \in \mathcal{C}^0(\Omega), (\phi_\rho)_\rho \in L^1(\mathbb{P}; \mathcal{C}^0(\Omega)), \int_{\mathcal{P}(\Omega)} \phi_\rho(\cdot) d\mathbb{P}(\rho) = \varphi(\cdot) \right\} \\ &= \sup \left\{ \int_{\mathcal{P}(\Omega)} \langle \phi_\rho | \nu \rangle d\mathbb{P}(\rho) + \int_{\mathcal{P}(\Omega)} \langle \phi_\rho^c | \rho \rangle d\mathbb{P}(\rho), \quad (\phi_\rho)_\rho \in L^1(\mathbb{P}; \mathcal{C}^0(\Omega)) \right\} \\ &= \int_{\mathcal{P}(\Omega)} \left(\sup_{\phi_\rho \in \mathcal{C}^0(\Omega)} \langle \phi_\rho | \nu \rangle + \langle \phi_\rho^c | \rho \rangle \right) d\mathbb{P}(\rho) \\ &= \int_{\mathcal{P}(\Omega)} \frac{1}{2} W_2^2(\nu, \rho) d\mathbb{P}(\rho), \end{aligned}$$

where we used the Kantorovich duality formula (see Theorems 1.2 or A.5 or for instance (Villani, 2008)) to get to the last line. By definition of $(\mathbf{P})_{\mathbb{P}}$ we then have

$$(\mathbf{P})_{\mathbb{P}} = \inf_{\nu \in \mathcal{P}(\Omega)} H^*(\nu) = -H^{**}(0).$$

Therefore, showing that $(\mathbf{P})_{\mathbb{P}} = \widetilde{(\mathbf{P})}_{\mathbb{P}}$ corresponds to show that $H(0) = H^{**}(0)$. Since H is convex (by concavity of the c -transform operation), this will follow from the continuity of H at 0 for the supremum-norm over $\mathcal{C}^0(\Omega)$ (Proposition 4.1 of (Ekeland and Témam, 1999)). For this, we can first notice that H never takes the value $-\infty$: for any $\varphi \in \mathcal{C}^0(\Omega)$ and $(\phi_\rho)_\rho \in L^1(\mathbb{P}; \mathcal{C}^0(\Omega))$ such that $\int_{\mathcal{P}(\Omega)} \phi_\rho(\cdot) d\mathbb{P}(\rho) = \varphi(\cdot)$, one has

$$\forall \rho \in \mathcal{P}(\Omega), \quad -\phi_\rho^c(x) = \sup_{y \in \Omega} \phi_\rho(y) - \frac{1}{2} \|x - y\|^2 \geq \phi_\rho(0) - \frac{1}{2} \|x\|^2.$$

It follows that

$$H(\varphi) \geq \varphi(0) - \int_{\mathcal{P}(\Omega)} \frac{M_2(\rho)}{2} d\mathbb{P}(\rho) > -\infty.$$

On the other hand, notice that H is bounded from above in a neighborhood of 0 in $\mathcal{C}^0(\Omega)$: let $\varphi \in \mathcal{C}^0(\Omega)$ be such that $\|\varphi\|_\infty \leq 1$. Then for any $x \in \Omega$,

$$\varphi^c(x) = \inf_{y \in \Omega} \frac{1}{2} \|x - y\|^2 - \varphi(y) \geq \inf_{y \in \Omega} \frac{1}{2} \|x - y\|^2 + \inf_{y \in \Omega} -\varphi(y) \geq 0 + 1,$$

so that $-\varphi^c(x) \leq 1$. Thus

$$H(\varphi) \leq - \int_{\mathcal{P}(\Omega)} \langle (\varphi)^c | \rho \rangle d\mathbb{P}(\rho) \leq 1.$$

A standard convex analysis result (Proposition 2.5 in (Ekeland and T emam, 1999)) then ensures that H is continuous at 0, so that $H(0) = H^{**}(0)$ and $(\mathbb{P})_{\mathbb{P}} = \widetilde{(\mathbb{P})}_{\mathbb{P}}$.

Restriction to $L^\infty(\mathbb{P}; W^{1,\infty}(\Omega))$. We show here that we can run the supremum $\widetilde{(\mathbb{P})}_{\mathbb{P}}$ only over $L^\infty(\mathbb{P}; W^{1,\infty}(\Omega))$ instead of $L^1(\mathbb{P}; \mathcal{C}^0(\Omega))$, that is

$$\widetilde{(\mathbb{P})}_{\mathbb{P}} = \sup \left\{ \int_{\mathcal{P}(\Omega)} \langle \phi_\rho^c | \rho \rangle d\mathbb{P}(\rho), \quad (\phi_\rho)_\rho \in L^\infty(\mathbb{P}; W^{1,\infty}(\Omega)), \quad \int_{\mathcal{P}(\Omega)} \phi_\rho(\cdot) d\mathbb{P}(\rho) = 0 \right\}.$$

Let $(\phi_\rho)_\rho \in L^1(\mathbb{P}; \mathcal{C}^0(\Omega))$ be an admissible solution to $\widetilde{(\mathbb{P})}_{\mathbb{P}}$, i.e. $(\phi_\rho)_\rho$ satisfies

$$\int_{\mathcal{P}(\Omega)} \phi_\rho(\cdot) d\mathbb{P}(\rho) = 0. \quad (6.13)$$

Then we can construct from $(\phi_\rho)_\rho$ another admissible solution $(\tilde{\phi}_\rho)_\rho$ that belongs to $L^\infty(\mathbb{P}; W^{1,\infty}(\Omega))$ and that performs better at $\widetilde{(\mathbb{P})}_{\mathbb{P}}$, i.e. that verifies

$$\int_{\mathcal{P}(\Omega)} \langle \tilde{\phi}_\rho^c | \rho \rangle d\mathbb{P}(\rho) \geq \int_{\mathcal{P}(\Omega)} \langle \phi_\rho^c | \rho \rangle d\mathbb{P}(\rho). \quad (6.14)$$

Indeed, introduce $(\hat{\phi}_\rho)_\rho := (\phi_\rho^{cc})_\rho$. Then for all $\rho \in \mathcal{P}(\Omega)$, $\hat{\phi}_\rho = \phi_\rho^{cc}$ is $2R_\Omega$ -Lipschitz continuous (as a c -transform over $\Omega = B(0, R_\Omega)$) and satisfies $\hat{\phi}_\rho^c = \phi_\rho^c$ and $\hat{\phi}_\rho \geq \phi_\rho$ (as a double c -transform). Using then (6.13), one has that

$$\alpha(\cdot) := \int_{\mathcal{P}(\Omega)} \hat{\phi}_\rho(\cdot) d\mathbb{P}(\rho) \geq 0,$$

where α is also $2R_\Omega$ -Lipschitz. Now denoting $\tilde{\phi}_\rho = \hat{\phi}_\rho - \alpha$ for all $\rho \in \mathcal{P}(\Omega)$, the mapping $(\tilde{\phi}_\rho)_\rho \in L^1(\mathbb{P}; \mathcal{C}^0(\Omega))$ is admissible to $\widetilde{(\mathbb{P})}_{\mathbb{P}}$ by construction and satisfies $\tilde{\phi}_\rho \leq \hat{\phi}_\rho$ for all $\rho \in \mathcal{P}(\Omega)$, so that $\tilde{\phi}_\rho^c \geq \hat{\phi}_\rho^c = \phi_\rho^c$ (using that the c -transform is order-reversing). For each $\rho \in \mathcal{P}(\Omega)$, up to subtracting $\tilde{\phi}_\rho(0)$ to $\tilde{\phi}_\rho$ (this operation leaves $(\tilde{\phi}_\rho)_\rho$ admissible to $\widetilde{(\mathbb{P})}_{\mathbb{P}}$ and does not change its value), one can assume that $\tilde{\phi}_\rho(0) = 0$. Noticing that $\tilde{\phi}_\rho$ is $4R_\Omega$ -Lipschitz by construction, we thus have the bound

$$\|\tilde{\phi}_\rho\|_{W^{1,\infty}(\Omega)} = \|\tilde{\phi}_\rho\|_\infty + \|\nabla \tilde{\phi}_\rho\|_\infty \leq 4R_\Omega(R_\Omega + 1).$$

We thus have built an admissible $(\tilde{\phi}_\rho)_\rho \in L^\infty(\mathbb{P}; W^{1,\infty}(\Omega))$ from an admissible $(\phi_\rho)_\rho \in L^1(\mathbb{P}; \mathcal{C}^0(\Omega))$ that satisfies (6.14), which shows that we can run the supremum $\widetilde{(\mathbb{P})}_{\mathbb{P}}$ only over $L^\infty(\mathbb{P}; W^{1,\infty}(\Omega))$ instead of $L^1(\mathbb{P}; \mathcal{C}^0(\Omega))$

Existence of a maximizer. There now remains to show that the supremum in $\widetilde{(\mathbb{P})}_{\mathbb{P}}$ can be replaced by a maximum. Let $((\phi_\rho^n)_\rho)_{n \geq 1}$ be a maximizing sequence to $\widetilde{(\mathbb{P})}_{\mathbb{P}}$, and assume from what precedes that this sequence belongs to $L^\infty(\mathbb{P}; W^{1,\infty}(\Omega))$ and satisfies for all $n \geq 0$ and $\rho \in \mathcal{P}(\Omega)$, $\|\phi_\rho^n\|_{W^{1,\infty}(\Omega)} \leq 4R_\Omega(1 + R_\Omega)$. Further assume that this sequence verifies for all $n \geq 1$,

$$\int_{\mathcal{P}(\Omega)} \langle (\phi_\rho^n)^c | \rho \rangle d\mathbb{P}(\rho) \geq \widetilde{(\mathbb{P})}_{\mathbb{P}} - \frac{1}{n}. \quad (6.15)$$

For any $n \geq 1$, the mapping $(\rho, x) \mapsto \phi_\rho^n(x)$ is bounded in $L^2(\mathbb{P} \otimes \lambda)$ where λ denotes the Lebesgue measure over Ω . Therefore, by Banach-Alaoglu theorem, the sequence

$((\phi_\rho^n)_\rho)_{n \geq 1}$ (seen as a sequence in $L^2(\mathbb{P} \otimes \lambda)$) admits a weakly converging subsequence in $L^2(\mathbb{P} \otimes \lambda)$, that we do not relabel and for which we denote $(\phi_\rho^\infty)_\rho$ the weak limit in $L^2(\mathbb{P} \otimes \lambda)$. Using Mazur's lemma, we know that there exists a sequence of integers $(N_n)_{n \geq 1}$ and coefficients $((\lambda_{n,k})_{n \leq k \leq N_n})_{n \geq 1} \geq 0$ satisfying for all $n \geq 1$, $\sum_{k=n}^{N_n} \lambda_{n,k} = 1$ such that the sequence $((\bar{\phi}_\rho^n)_\rho)_{n \geq 1}$ defined for all $n \geq 1$ and $\rho \in \mathcal{P}(\Omega)$ by $\bar{\phi}_\rho^n := \sum_{k=n}^{N_n} \lambda_{n,k} \phi_\rho^k$ converges strongly toward $(\phi_\rho^\infty)_\rho$ in $L^2(\mathbb{P} \otimes \lambda)$. By concavity of the c -transform operation and equation (6.15), we then have for all $n \geq 1$ the bound

$$\begin{aligned} \int_{\mathcal{P}(\Omega)} \langle (\bar{\phi}_\rho^n)^c | \rho \rangle d\mathbb{P}(\rho) &\geq \sum_{k=n}^{N_n} \lambda_{n,k} \int_{\mathcal{P}(\Omega)} \langle (\phi_\rho^k)^c | \rho \rangle d\mathbb{P}(\rho) \\ &\geq \sum_{k=n}^{N_n} \lambda_{n,k} \left(\widetilde{(\mathbb{P})}_{\mathbb{P}} - \frac{1}{k} \right) \\ &\geq \widetilde{(\mathbb{P})}_{\mathbb{P}} - \frac{1}{n}. \end{aligned} \quad (6.16)$$

The sequence $((\bar{\phi}_\rho^n)_\rho)_{n \geq 1}$ is therefore also a maximizing sequence of $\widetilde{(\mathbb{P})}_{\mathbb{P}}$ and it also satisfies for any $n \geq 1$ and $\rho \in \mathcal{P}(\Omega)$ the bound

$$\|\bar{\phi}_\rho^n\|_{W^{1,\infty}(\Omega)} \leq 4R_\Omega(1 + R_\Omega). \quad (6.17)$$

Since the sequence $((\bar{\phi}_\rho^n)_\rho)_{n \geq 1}$ strongly converges to $(\phi_\rho^\infty)_\rho$ in $L^2(\mathbb{P} \otimes \lambda)$, one can extract a subsequence (that we do not relabel) such that for \mathbb{P} -almost-every $\rho \in \mathcal{P}(\Omega)$, the sequence $(\bar{\phi}_\rho^n)_{n \geq 1}$ converges to ϕ_ρ^∞ in $L^2(\lambda)$. Using bound (6.17) and Arzelà-Ascoli theorem, we deduce that for \mathbb{P} -almost-every $\rho \in \mathcal{P}(\Omega)$, the sequence $(\bar{\phi}_\rho^n)_{n \geq 1}$ converges uniformly to ϕ_ρ^∞ in $C^0(\Omega)$ and that

$$\|\phi_\rho^\infty\|_{W^{1,\infty}(\Omega)} \leq 4R_\Omega(1 + R_\Omega).$$

In particular, $(\phi_\rho^\infty)_\rho$ belongs to $L^\infty(\mathbb{P}; W^{1,\infty}(\Omega))$ and we have the limit

$$0 = \int_{\mathcal{P}(\Omega)} \bar{\phi}_\rho^n(\cdot) d\mathbb{P}(\rho) \xrightarrow{n \rightarrow \infty} \int_{\mathcal{P}(\Omega)} \phi_\rho^\infty(\cdot) d\mathbb{P}(\rho),$$

so that $(\phi_\rho^\infty)_\rho$ is admissible to $\widetilde{(\mathbb{P})}_{\mathbb{P}}$. Eventually, for \mathbb{P} -almost-every $\rho \in \mathcal{P}(\Omega)$, we have the limit

$$\langle (\bar{\phi}_\rho^n)^c | \rho \rangle \xrightarrow{n \rightarrow \infty} \langle (\phi_\rho^\infty)^c | \rho \rangle, \quad (6.18)$$

so that by Lebesgue's dominated convergence theorem and bound (6.16),

$$\int_{\mathcal{P}(\Omega)} \langle (\phi_\rho^\infty)^c | \rho \rangle d\mathbb{P}(\rho) = \lim_{n \rightarrow +\infty} \int_{\mathcal{P}(\Omega)} \langle (\bar{\phi}_\rho^n)^c | \rho \rangle d\mathbb{P}(\rho) = \widetilde{(\mathbb{P})}_{\mathbb{P}},$$

which proves that $(\phi_\rho^\infty)_\rho \in L^\infty(\mathbb{P}; W^{1,\infty}(\Omega))$ is a maximizer for $\widetilde{(\mathbb{P})}_{\mathbb{P}}$. \square

Quantitative stability of Schrödinger potentials with respect to the temperature in the semi-discrete setting

Abstract

This chapter, extracted from (Delalande, 2022), derives nearly tight and non-asymptotic convergence bounds for solutions of entropic semi-discrete optimal transport. These bounds quantify the stability of the dual solutions of the regularized problem (sometimes called Sinkhorn or Schrödinger potentials) with respect to the regularization parameter, for which we ensure a better than Lipschitz dependence. Such facts may be a first step towards a mathematical justification of ε -scaling heuristics for the numerical resolution of regularized semi-discrete optimal transport. These results also entail a non-asymptotic and tight expansion of the difference between the entropic and the unregularized costs.

7.1 Introduction

Optimal transport and the distances it defines are now widely acknowledged as important tools for machine learning (Canas and Rosasco, 2012; Arjovsky et al., 2017; Genevay et al., 2018; Flamary et al., 2018; Alaux et al., 2019; Gordaliza et al., 2019) and statistics (Ramdas et al., 2015; Cazelles et al., 2018; Bigot et al., 2019a; Weed and Berthet, 2019). In these fields, it is also recognized that the original formulation of the transport problem suffers in general from poor computationability and statistical behavior with respect to the dimension, and that some form of regularization can be helpful. In this state of mind, the entropic regularization of the optimal transport problem that we have already studied in Chapter 4 has proven to be a relevant choice of regularization. Let us recall in this chapter the formulation of this regularized problem and some of its attractive features. For two compact subsets \mathcal{X}, \mathcal{Y} of \mathbb{R}^d , two probability measures $\rho \in \mathcal{P}(\mathcal{X}), \mu \in \mathcal{P}(\mathcal{Y})$, and for $\varepsilon \geq 0$, the quadratic optimal transport problem between ρ and μ with entropic

regularization of parameter ε reads

$$\min_{\gamma \in \Gamma(\rho, \mu)} \int_{\mathcal{X} \times \mathcal{Y}} \|x - y\|^2 d\gamma(x, y) + \varepsilon \text{KL}(\gamma | \rho \otimes \mu), \quad (\text{P}_\varepsilon)$$

where $\Gamma(\rho, \mu)$ denotes the set of couplings between ρ and μ and KL denotes the Kullback-Leibler divergence or relative entropy (up to an additive term):

$$\text{KL}(\gamma | \rho \otimes \mu) = \int_{\mathcal{X} \times \mathcal{Y}} \left(\log \left(\frac{d\gamma}{d\rho \otimes \mu}(x, y) \right) - 1 \right) d\gamma(x, y)$$

if $\gamma \ll \rho \otimes \mu$ and $+\infty$ otherwise. When $\varepsilon = 0$, problem (P_ε) corresponds to the usual quadratic optimal transport problem between ρ and μ , and the value of (P_ε) defines in this case the square of the 2-Wasserstein distance W_2 between ρ and μ (see Chapter A). However, choosing $\varepsilon > 0$ in (P_ε) has several advantages: first, it turns problem (P_ε) in a ε -strongly-convex minimization problem, which enables the use of fast algorithms for its resolution (Cuturi, 2013; Altschuler et al., 2017; Dvurechensky et al., 2018; Peyré and Cuturi, 2019; Schmitzer, 2019). Second, problem (P_ε) enjoys better statistical properties when $\varepsilon > 0$ rather than when $\varepsilon = 0$, with improved sample complexity for its value (Genevay et al., 2019; Mena and Niles-Weed, 2019) and better guarantees when using stochastic optimization algorithms for its resolution (Genevay et al., 2016; Bercu and Bigot, 2020). Thus, with $\varepsilon > 0$, introducing the quantity

$$W_{2,\varepsilon}(\rho, \mu) = \left(\int_{\mathcal{X} \times \mathcal{Y}} \|x - y\|^2 d\gamma^{(\text{P}_\varepsilon)}(x, y) \right)^{1/2},$$

where $\gamma^{(\text{P}_\varepsilon)}$ is the solution to (P_ε) , one may hope that $W_{2,\varepsilon}$ approximates W_2 well when ε is not too large. This fact has been the object of a long line of works, going to very recent developments. The convergence of $W_{2,\varepsilon}$ to W_2 as ε goes to zero is established in general settings (Mikami, 2004; Léonard, 2012; Nutz and Wiesel, 2021; Bernton et al., 2022), and it has been quantified in more specific settings. In the *continuous setting* where both ρ and μ are absolutely continuous, (Adams et al., 2011; Duong et al., 2013; Erbar et al., 2015; Pal, 2019) gave first order asymptotics for $W_{2,\varepsilon}$ in terms of ε and thus showed in this setting an asymptotic linear rate of convergence of $W_{2,\varepsilon}$ to W_2 . These results were recently refined in (Conforti and Tamanini, 2021) where second order asymptotics have been given: the authors of (Conforti and Tamanini, 2021) have shown in particular that if ρ and μ are absolutely continuous with bounded densities, then

$$\begin{aligned} W_{2,\varepsilon}^2(\rho, \mu) + \varepsilon \text{KL}(\gamma^{(\text{P}_\varepsilon)} | \rho \otimes \mu) &= W_2^2(\rho, \mu) - \frac{\varepsilon}{2} (\text{KL}(\rho | \lambda) + \text{KL}(\mu | \lambda)) - \frac{\varepsilon}{2} d \log(\pi \varepsilon) \\ &\quad + \frac{\varepsilon^2}{16} I(\rho, \mu) + o(\varepsilon^2), \end{aligned}$$

where λ denotes the Lebesgue measure on \mathbb{R}^d and $I(\rho, \mu)$ is the integrated Fisher information along the 2-Wasserstein geodesic connecting ρ and μ . In contrast, in the *discrete setting* where both ρ and μ are finitely supported, the rate of convergence of $W_{2,\varepsilon}$ to W_2 was shown to be asymptotically exponential in (Cominetti and Martín, 1994) in the context of the analysis of exponentially penalized finite dimensional linear programs. This result was then refined with a tight non-asymptotic analysis in (Weed, 2018), where it was shown that for ρ and μ discrete, there exists (explicit) positive constants $C_{\rho,\mu}, \tilde{C}_{\rho,\mu}$ depending only on ρ and μ such that for any $\varepsilon > 0$,

$$0 \leq W_{2,\varepsilon}^2(\rho, \mu) - W_2^2(\rho, \mu) \leq C_{\rho,\mu} \exp(-\tilde{C}_{\rho,\mu}/\varepsilon).$$

Notably, such non-asymptotic result may allow to choose ε in terms of the data in order to compute the unregularized cost W_2 to a wanted precision from an entropic approximation $W_{2,\varepsilon}$.

Very different regimes are thus observed between the continuous setting (with an asymptotic linear convergence rate) and the discrete setting (with a non-asymptotic exponential convergence rate). However, very few was known – until recently (Altschuler et al., 2022) – on the intermediate setting of *semi-discrete* optimal transport, where ρ is absolutely continuous and μ is finitely supported, that is of particular importance in some applications. In statistics, it corresponds to the case where one wants to compare an empirical sample to a given probability measures, and it can serve to extend notions of quantiles and ranks to multivariate measures (Chernozhukov et al., 2017). In numerical analysis, the semi-discrete setting gives a natural framework to approximate the solution of the optimal transport problem between a probability density ρ and a probability measure μ that consists in approximating μ by a sequence of measures $(\mu_N)_{N \geq 1}$ with finite support such that $\lim_{N \rightarrow +\infty} W_2(\mu, \mu_N) = 0$ (Oliker and Prussner, 1989; Cullen et al., 1991; Gangbo and McCann, 1996; Caffarelli et al., 1999; Mirebeau, 2015). Finally in image processing, semi-discrete transport has proven useful for texture synthesis and style transfer (Galerne et al., 2017, 2018; Leclaire and Rabin, 2020). We thus focus in this chapter on the semi-discrete setting, and show that we can improve the recent asymptotic bounds given in (Altschuler et al., 2022) under slightly stronger regularity assumptions on the source measure. In particular, we produce a non-asymptotic analysis of the dual solutions to problem (P_ε) in terms of ε , which may be important in itself for the resolution of semi-discrete optimal transport using ε -scaling techniques. We then deduce non-asymptotic bounds for $W_{2,\varepsilon}$ in terms of ε in this semi-discrete framework.

Outline. Section 7.2 succinctly recalls elements of semi-discrete (entropic) optimal transport from Chapters 2 and 4 and state our main results. Section 7.3 derives the ODE from which starts the proof of our main bound. This ODE presents two terms that both involve the entropic semi-discrete Kantorovich functional of Chapter 4. The strong-convexity estimate of Chapter 4 together with another estimate derived in Section 7.4 then allow to prove our main bound. This main bound admits two corollaries giving rates of convergences of entropic semi-discrete solutions to their non-regularized counterpart (proven in Section 7.5) and rates of convergence of the entropic cost to the classical cost (proven Section 7.6). Section 7.7 finally illustrates our theoretical results on simple one-dimensional numerical examples.

7.2 Convergence bounds for semi-discrete entropic optimal transport

7.2.1 Semi-discrete (Entropic) Optimal Transport

Let \mathcal{X} be a compact subset of \mathbb{R}^d and $\rho \in \mathcal{P}(\mathcal{X})$ be an absolutely continuous probability measure on \mathcal{X} . Let $\mathcal{Y} = \{y_1, \dots, y_N\} \subset \mathbb{R}^d$ be a set of N points in \mathbb{R}^d and let σ be the counting measure associated to this set, i.e. $\sigma = \sum_{i=1}^N \delta_{y_i}$. Let $\mu = \sum_{i=1}^N \mu_i \delta_{y_i} \in \mathcal{P}(\mathcal{Y})$ where for all $i, \mu_i \geq \underline{\mu} > 0$. Note that we will denote $R_{\mathcal{X}}, R_{\mathcal{Y}} > 0$ the smallest constants such that the $\mathcal{X} \subset B(0, R_{\mathcal{X}}), \mathcal{Y} \subset B(0, R_{\mathcal{Y}})$ respectively, as well as $\text{diam}(\mathcal{X}), \text{diam}(\mathcal{Y})$ the respective diameter of \mathcal{X}, \mathcal{Y} .

We noticed in Chapters 1 and 4 that developing the square $\|x - y\|^2$ in (P_ε) and using that γ belongs to $\Gamma(\rho, \mu)$, this problem is equivalent to the following regularized maximum correlation problem:

$$\max_{\gamma \in \Gamma(\rho, \mu)} \int_{\mathcal{X} \times \mathcal{Y}} \langle x|y \rangle d\gamma(x, y) - \varepsilon \text{KL}(\gamma | \rho \otimes \sigma), \quad (P'_\varepsilon)$$

with the relation $(P_{2\varepsilon}) = M_2(\rho) + M_2(\mu) - 2\varepsilon \mathcal{H}(\mu) - 2 \times (P'_\varepsilon)$, where $M_2(\cdot)$ denotes the second moment of a probability measure and $\mathcal{H}(\cdot)$ its Shannon entropy. By either ε -strong concavity (when $\varepsilon > 0$) or Brenier's theorem (Brenier, 1991) (when $\varepsilon = 0$, using that ρ is absolutely continuous), we can ensure that problem (P'_ε) admits a unique solution that we denote γ^ε . Moreover, we know from Chapters 1, 2 and 4 that (P'_ε) admits the following (semi-)dual formulation and strong duality holds (see also for instance Sections 2 of (Genevay et al., 2016; Bercu and Bigot, 2020)):

$$\min_{\psi \in \mathbb{R}^N} \int_{\mathcal{X}} \psi^{c, \varepsilon} d\rho + \langle \psi | \mu \rangle + \varepsilon,$$

where μ is conflated with the vector $(\mu_i)_{i=1, \dots, N} \in (\mathbb{R}_+^*)^N$ and where $\psi^{c, \varepsilon}$ corresponds to the (c, ε) -transform of ψ when $\varepsilon > 0$ and to its Legendre transform when $\varepsilon = 0$: $\forall x \in \mathcal{X}$,

$$\psi^{c, \varepsilon}(x) = \begin{cases} \varepsilon \log \left(\sum_{i=1}^N e^{\frac{\langle x|y_i \rangle - \psi_i}{\varepsilon}} \right) & \text{if } \varepsilon > 0, \\ \max_{i=1, \dots, N} \langle x|y_i \rangle - \psi_i = \psi^*(x) & \text{if } \varepsilon = 0. \end{cases}$$

This dual problem is invariant to addition of constant vectors to ψ . We fix this invariance by adding the constraint that $\langle \psi | \mathbb{1}_N \rangle = 0$ without any loss of generality, where $\mathbb{1}_N$ denotes the all-ones vector of \mathbb{R}^N . As in Section 2.2 of Chapter 2 and Section 4.3 of Chapter 4, introducing the semi-discrete (entropic) Kantorovich functional $\mathcal{K}_\rho^\varepsilon : \psi \mapsto \int_{\mathcal{X}} \psi^{c, \varepsilon} d\rho + \varepsilon$, one can then rewrite the dual formulation as

$$\min_{\psi \in \mathbb{R}^N, \langle \psi | \mathbb{1}_N \rangle = 0} \mathcal{K}_\rho^\varepsilon(\psi) + \langle \psi | \mu \rangle. \quad (D_\varepsilon)$$

From Sections 2.2 and 4.3 we know that the functional $\mathcal{K}_\rho^\varepsilon$ is convex on \mathbb{R}^N and strictly convex on $(\mathbb{1}_N)^\top$: problem (D_ε) admits a unique solution denoted ψ^ε (that we call later on a *potential*). We also know from Sections 2.2 and 4.3 that $\mathcal{K}_\rho^\varepsilon$ is \mathcal{C}^2 on \mathbb{R}^N , with first and second order derivatives available in (2.3), (2.5), (4.7) and (4.8). The unique solution ψ^ε of (D_ε) thus verifies the following first-order condition:

$$\nabla \mathcal{K}_\rho^\varepsilon(\psi^\varepsilon) = -\mu. \quad (7.1)$$

More precisely, this first-order condition means that for all $i \in \{1, \dots, N\}$, if $\varepsilon > 0$,

$$\mu(\{y_i\}) = \int_{x \in \mathcal{X}} e^{\frac{\langle x|y_i \rangle - \psi_i^\varepsilon - (\psi^\varepsilon)^{c, \varepsilon}(x)}{\varepsilon}} d\rho(x) \quad (7.2)$$

and if $\varepsilon = 0$, $\mu(\{y_i\}) = \int_{x \in \mathcal{X}} \mathbb{1}_{\text{Lag}_i(\psi^0)}(x) d\rho(x)$, where for any $\psi \in \mathbb{R}^N$, $\text{Lag}_i(\psi)$ denotes the i -th Laguerre cell w.r.t. ψ :

$$\text{Lag}_i(\psi) = \{x \in \mathcal{X} | \forall j, \langle x|y_i \rangle - \psi_i \geq \langle x|y_j \rangle - \psi_j\}.$$

Note that the Laguerre cells are convex polytopes intersected with \mathcal{X} and they define a tessellation of \mathcal{X} : $\bigcup_i \text{Lag}_i(\psi^\varepsilon) = \mathcal{X}$.

Finally, the primal-dual relationship that links the solution ψ^ε of problem (D_ε) to the solution γ^ε of problem (P'_ε) is the following: for all Borel set $A \subset \mathcal{X}$, all $i \in \{1, \dots, N\}$, if $\varepsilon > 0$,

$$\gamma^\varepsilon(A, \{y_i\}) = \int_{x \in A} e^{\frac{\langle x | y_i \rangle - \psi_i^\varepsilon - (\psi^\varepsilon)^{c, \varepsilon}(x)}{\varepsilon}} d\rho(x), \quad (7.3)$$

and if $\varepsilon = 0$, $\gamma^0(A, \{y_i\}) = \int_{x \in A} \mathbb{1}_{\text{Lag}_i(\psi^0)}(x) d\rho(x)$.

7.2.2 Non-asymptotic Behavior of Potentials

The authors of (Altschuler et al., 2022) recently tackled the question of the rate of convergence of $W_{2, \varepsilon}$ to W_2 in the specific semi-discrete setting. They showed that, mainly assuming that the absolutely continuous source ρ is such that the interior of its support is connected, the boundary of its support has zero Lebesgue measure, and ρ has positive density on the interior of its support, then the following asymptotic expansion holds (Theorem 1.1 in (Altschuler et al., 2022)):

$$W_{2, \varepsilon}^2(\rho, \mu) = W_2^2(\rho, \mu) + \varepsilon^2 \frac{\pi^2}{12} \sum_{i < j} \frac{w_{ij}}{\|y_i - y_j\|} + o(\varepsilon^2),$$

where $w_{ij} = \int_{\text{Lag}_i(\psi^0) \cap \text{Lag}_j(\psi^0)} \rho(x) d\mathcal{H}^{d-1}(x)$. In order to show this, they demonstrated that the convergence of ψ^ε to ψ^0 as ε goes to 0 happens at a rate faster than ε (Theorem 1.3 in (Altschuler et al., 2022)):

$$\lim_{\varepsilon \rightarrow 0} \frac{1}{\varepsilon} (\psi^\varepsilon - \psi^0) = \dot{\psi}^\varepsilon \Big|_{\varepsilon=0} = 0,$$

where $\dot{\psi}^\varepsilon = \frac{\partial}{\partial \varepsilon} \psi^\varepsilon$. As discussed in (Altschuler et al., 2022), this result is unexpected because false in general optimal transport and stems from the particular setting of semi-discrete optimal transport with a positive source. Quoting (Altschuler et al., 2022), we notice indeed that when ρ and μ are discrete, the quantity $\frac{1}{\varepsilon} (\psi^\varepsilon - \psi^0)$ may converge to a non-zero limit in general (see e.g. Proposition 5.5 of (Cominetti and Martín, 1994)). The authors of (Altschuler et al., 2022) also notice that the assumption that ρ is positive on the interior of its support is essential for this asymptotic result to hold, and that if ρ does not satisfy this assumption then it is possible to find examples where $\frac{1}{\varepsilon} (\psi^\varepsilon - \psi^0)$ may diverge (taking for instance ρ decaying to zero at different rates on opposite sides of one of the hyperplane boundaries $\text{Lag}_i(\psi^0) \cap \text{Lag}_j(\psi^0)$).

In this work, we show that the result of Theorem 1.3 in (Altschuler et al., 2022) can be extended and quantified to get a non-asymptotic control of $\dot{\psi}^\varepsilon$, i.e. not only when $\varepsilon \rightarrow 0$ but for $\varepsilon \in \mathbb{R}_+^*$. As in (Altschuler et al., 2022), we notice that regularity assumptions on the source measure ρ are necessary to proceed with such controls. In particular, we make the following assumption (stronger than the one in (Altschuler et al., 2022)):

Assumption 7.1. The compact set \mathcal{X} is convex. The source measure $\rho \in \mathcal{P}(\mathcal{X})$ is absolutely continuous and its density (also denoted ρ), is bounded away from zero and infinity, i.e. there exist m_ρ, M_ρ such that on \mathcal{X} ,

$$0 < m_\rho \leq \rho \leq M_\rho < +\infty.$$

Under this assumption and an Hölder continuity assumption on the density of ρ , we show the following behavior:

Theorem 7.2. *Let $\rho \in \mathcal{P}(\mathcal{X})$ satisfying Assumption 7.1 with an α -Hölder continuous density for some $\alpha \in (0, 1]$ and let $\mu \in \mathcal{P}(\mathcal{Y})$. Then for any $\varepsilon \leq 1$, $\alpha' \in (0, \alpha)$, the solutions ψ^ε to problem (D_ε) verify:*

$$\|\dot{\psi}^\varepsilon\|_2 \lesssim \varepsilon^{\alpha'},$$

where $\dot{\psi}^\varepsilon = \frac{\partial}{\partial \varepsilon} \psi^\varepsilon$ and \lesssim hides multiplicative constants that depend on $\mathcal{X}, \rho, \mathcal{Y}, \mu$. Besides, for any $\varepsilon \geq 1$,

$$\|\dot{\psi}^\varepsilon\|_2 \lesssim 1$$

Remark 7.3 (Constants). A (very) rough upper bound on the hidden constants is given by the quantity

$$\begin{aligned} \frac{N}{\underline{\mu}} \frac{M_\rho}{m_\rho} e^{R_{\mathcal{Y}} \text{diam}(\mathcal{X})} & \left(N R_{\mathcal{X}} \text{diam}(\mathcal{Y}) + \log \frac{1}{\underline{\mu}} \right. \\ & + N^2 M_\rho \text{diam}(\mathcal{X})^{d-1} \left(1 + \frac{C_\rho}{\delta^\alpha} + R_{\mathcal{X}} \text{diam}(\mathcal{Y}) + \log \frac{1}{\underline{\mu}} \right) \\ & \left. + N^3 M_\rho \frac{\text{diam}(\mathcal{X})^{d-2} \text{diam}(\mathcal{Y})^4}{\cos(\theta/2) \delta^4} \left(1 + R_{\mathcal{X}} \text{diam}(\mathcal{Y}) + \log \frac{1}{\underline{\mu}} \right) \right), \end{aligned}$$

up to a multiplicative constant that depends only on the dimension. In this formula, C_ρ is such that for any $x, x' \in \mathcal{X}$, $|\rho(x) - \rho(x')| \leq C_\rho \|x - x'\|^\alpha$, δ is the minimum distance between two points in \mathcal{Y} and θ is the maximum angle formed by three not-aligned points in \mathcal{Y} . The dependence on N is rather bad and it may be improved by replacing the N^2 term with N times the maximum number of $(d-1)$ -facets a Laguerre cell has in the tessellation $\bigcup_i \text{Lag}_i(\psi^0)$ and the N^3 term with N times the maximum number of $(d-2)$ -facets a Laguerre cell has in this tessellation.

Remark 7.4 (Assumptions on the source measure). The source measure ρ is assumed to satisfy some restrictive regularity assumptions, namely it should be at least absolutely continuous with bounded density supported on a compact convex set. These assumptions are required essentially to be able to apply at some point the strong convexity estimate of Theorem 4.4 from Chapter 4. We note that the convexity assumption made on the support of ρ may be relaxed to some extent as in Section 1.4: the bound of Theorem 7.2 holds more generally if ρ is absolutely continuous with bounded density supported on a compact set made of a finite union of convex sets and it satisfies a Poincaré-Wirtinger inequality.

The behavior of Theorem 7.2 is a consequence of the analysis of an ODE satisfied by the map $\varepsilon \mapsto \psi^\varepsilon$, and it is proven in Section 7.3. An immediate consequence of this result concerns the quantitative stability of the mapping $\varepsilon \mapsto \psi^\varepsilon$. It also gives quantitative convergence results for ψ^ε , $(\psi^\varepsilon)^{c,\varepsilon}$ and γ^ε toward their different limits as ε goes to zero – results that are reminiscent of the asymptotic ones of (Cominetti and Martín, 1994) in the study of solutions of exponentially penalized finite dimensional linear programs. The proof of the following corollary follows rather directly from Theorem 7.2 and it is deferred to Section 7.5.

Corollary 7.5. *Let $0 < \varepsilon' \leq \varepsilon \leq 1$. Under assumptions of Theorem 7.2, denote $\psi^{\varepsilon'}$, ψ^ε the solutions to problem (D_ε) with regularization $\varepsilon', \varepsilon$ respectively. Then for any $\alpha' \in (0, \alpha)$,*

$$\left\| \psi^\varepsilon - \psi^{\varepsilon'} \right\|_\infty \lesssim \varepsilon^{\alpha'} (\varepsilon - \varepsilon').$$

In particular, letting ε' go to 0 yields

$$\begin{aligned} \left\| \psi^\varepsilon - \psi^0 \right\|_\infty &\lesssim \varepsilon^{1+\alpha'}, \\ \left\| (\psi^\varepsilon)^{c, \varepsilon} - (\psi^0)^* \right\|_\infty &\lesssim \varepsilon. \end{aligned}$$

Additionally, for ρ -a.e. $x \in \mathcal{X}$,

$$\begin{aligned} \left| (\psi^\varepsilon)^{c, \varepsilon}(x) - (\psi^0)^*(x) \right| &\lesssim \varepsilon^{1+\alpha'}, \\ \left| \gamma^\varepsilon(x, \cdot) - \gamma^0(x, \cdot) \right| &\lesssim e^{-c_x/\varepsilon}, \end{aligned}$$

where $c_x = \min_{i \in \{1, \dots, N\}} \{(\psi^0)^*(x) - \langle x | y_i \rangle + \psi_i^0 \mid \langle x | y_i \rangle - \psi_i^0 \neq (\psi^0)^*(x)\} > 0$.

Remark 7.6 (ε -scaling). Corollary 7.5 may be a first step, in the semi-discrete context, towards a mathematical justification of ε -scaling or ε -scheduling techniques used in the numerical resolution of optimal transport. Such techniques, reported for instance in (Kosowsky and Yuille, 1994; Schmitzer, 2019; Feydy, 2020) in the context of Sinkhorn's algorithm for solving the assignment or discrete optimal transport problems using entropic regularization, are used to reduce the number of iterations necessary to compute a regularized solution. They consist in solving (P_ε) with a starting *large* regularization parameter ε^0 , and then gradually decrease the regularization parameter over the course of the optimization, with a geometric decrease – typically, $\varepsilon^{k+1} = \varepsilon^k/2$. The idea is that ψ^{ε^k} (or an approximation of it) is supposed to be a good starting point for an optimization algorithm that aims at estimating the solution $\psi^{\varepsilon^{k+1}}$. This technique was introduced for Bertsekas' auction algorithm (Bertsekas, 1981; Bertsekas and Eckstein, 1988) for the resolution of the assignment problem, and it proved to reduce the worst case complexity from $O\left(\frac{N^2}{\varepsilon}\right)$ to $O\left(N^3 \log\left(\frac{1}{\varepsilon}\right)\right)$ in order to get an ε -approximate solution, where N denotes the number of agents/tasks. Although successful in practice, similar reduction of the worst-case complexity of Sinkhorn's algorithm using the ε -scaling strategy could not be proved, see the discussions in (Schmitzer, 2019; Feydy, 2020) for more details.

Remark 7.7 (Exponential convergence of γ^ε). The convergence of $\gamma^\varepsilon(x, y_i)$ to $\gamma^0(x, y_i)$ at the rate $e^{-c_x/\varepsilon}$ for ρ -a.e. x and $i \in \{1, \dots, N\}$ matches in our semi-discrete setting the result of (Bernton et al., 2022) that showed this rate of convergence, only asymptotically and for (x, y_i) not in the support of γ^0 , but in a much more general setting.

7.2.3 Non-asymptotic Expansion of the Difference of Costs

Another consequence of the new bounds of Theorem 7.2 for $\varepsilon \leq 1$ is an improvement of the asymptotic result on the convergence of the difference of costs proven in (Altschuler et al., 2022), to the following tight non-asymptotic result:

Theorem 7.8. *Under assumptions of Theorem 7.2, for any $\alpha' \in (0, \alpha)$ and $\varepsilon \leq 1$,*

$$\left| W_{2, \varepsilon}^2(\rho, \mu) - W_2^2(\rho, \mu) - \varepsilon^2 \frac{\pi^2}{12} \sum_{i < j} \frac{w_{ij}}{\|y_i - y_j\|} \right| \lesssim \varepsilon^{2+\alpha'}.$$

This result and its tightness are proven in Section 7.6.

7.3 A governing ODE

Similarly to (Cominetti and Martín, 1994), we show Theorem 7.2 by leveraging the fact that $\varepsilon \mapsto \boldsymbol{\psi}^\varepsilon$ satisfies a specific ODE that is deduced from the stationary equation (7.1). Let's fix $\varepsilon > 0$ and recall in this case the expression of the entropic semi-discrete Kantorovich's functional $\mathcal{K}_\rho^\varepsilon: \mathbb{R}^N \rightarrow \mathbb{R}$: for all $\boldsymbol{\psi} \in \mathbb{R}^N$,

$$\mathcal{K}_\rho^\varepsilon(\boldsymbol{\psi}) = \int_{\mathcal{X}} \varepsilon \log \left(\sum_{i=1}^N \exp \left(\frac{\langle x | y_i \rangle - \boldsymbol{\psi}_i}{\varepsilon} \right) \right) d\rho(x) + \varepsilon. \quad (7.4)$$

From Lemma 4.2, we know that $\mathcal{K}_\rho^\varepsilon$ is a \mathcal{C}^2 function from \mathbb{R}^N to \mathbb{R} . Its derivatives when evaluated in $\boldsymbol{\psi} \in \mathbb{R}^N$ read from equations (4.7), (4.8):

$$\begin{aligned} \nabla \mathcal{K}_\rho^\varepsilon(\boldsymbol{\psi}) &= -\mathbb{E}_{x \sim \rho} \gamma_x^\varepsilon(\boldsymbol{\psi}), \\ \nabla^2 \mathcal{K}_\rho^\varepsilon(\boldsymbol{\psi}) &= \frac{1}{\varepsilon} \mathbb{E}_{x \sim \rho} \left(\text{diag}(\gamma_x^\varepsilon(\boldsymbol{\psi})) - \gamma_x^\varepsilon(\boldsymbol{\psi}) \gamma_x^\varepsilon(\boldsymbol{\psi})^\top \right), \end{aligned}$$

where for any $x \in \mathcal{X}$ and $\boldsymbol{\psi} \in \mathbb{R}^N$, $\gamma_x^\varepsilon(\boldsymbol{\psi})$ is a vector of \mathbb{R}^N whose components read for all $i \in \{1, \dots, N\}$

$$\gamma_x^\varepsilon(\boldsymbol{\psi})_i = \frac{\exp \left(\frac{\langle x | y_i \rangle - \boldsymbol{\psi}_i}{\varepsilon} \right)}{\sum_{j=1}^N \exp \left(\frac{\langle x | y_j \rangle - \boldsymbol{\psi}_j}{\varepsilon} \right)}. \quad (7.5)$$

As noticed in Section 4.3, one can interpret $x \mapsto \gamma_x^\varepsilon(\boldsymbol{\psi})_i$ as a smoothed version of the indicator function associated to the i -th Laguerre cell of $\boldsymbol{\psi}$: it represents the ratio of mass sent from x to y_i proposed by the candidate solution $\boldsymbol{\psi}$ to problem (D_ε) . One can also easily prove that for any $\boldsymbol{\psi} \in \mathbb{R}^N$, $\varepsilon \mapsto \nabla \mathcal{K}_\rho^\varepsilon(\boldsymbol{\psi})$ is a \mathcal{C}^1 function from \mathbb{R}_+^* to \mathbb{R}^N , with the formula

$$\frac{\partial}{\partial \varepsilon} (\nabla \mathcal{K}_\rho^\varepsilon)(\boldsymbol{\psi}) = \frac{1}{\varepsilon} \mathbb{E}_{x \sim \rho} (\text{diag}(\gamma_x^\varepsilon(\boldsymbol{\psi})) \log \gamma_x^\varepsilon(\boldsymbol{\psi}) - \gamma_x^\varepsilon(\boldsymbol{\psi}) \gamma_x^\varepsilon(\boldsymbol{\psi})^\top \log \gamma_x^\varepsilon(\boldsymbol{\psi})). \quad (7.6)$$

We can then show:

Proposition 7.9. *For any $\varepsilon \geq 0$, denote $\boldsymbol{\psi}^\varepsilon$ the solution to problem (D_ε) . The mapping $\varepsilon \mapsto \boldsymbol{\psi}^\varepsilon$ is a \mathcal{C}^1 function from \mathbb{R}_+^* to $(\mathbb{1}_N)^\perp$ that satisfies for any $\varepsilon > 0$ the ODE*

$$\nabla^2 \mathcal{K}_\rho^\varepsilon(\boldsymbol{\psi}^\varepsilon) \boldsymbol{\psi}^\varepsilon + \frac{\partial}{\partial \varepsilon} (\nabla \mathcal{K}_\rho^\varepsilon)(\boldsymbol{\psi}^\varepsilon) = 0. \quad (7.7)$$

Proof. Since for any $\varepsilon > 0$, $\boldsymbol{\psi} \mapsto \mathcal{K}_\rho^\varepsilon(\boldsymbol{\psi})$ is a \mathcal{C}^2 convex function from \mathbb{R}^N to \mathbb{R} , one can characterize $\boldsymbol{\psi}^\varepsilon$ with the first order condition (7.1):

$$\nabla \mathcal{K}_\rho^\varepsilon(\boldsymbol{\psi}^\varepsilon) = -\mu.$$

Using that $\boldsymbol{\psi} \mapsto \mathcal{K}_\rho^\varepsilon(\boldsymbol{\psi})$ is \mathcal{C}^2 on \mathbb{R}^N , $\varepsilon \mapsto \nabla \mathcal{K}_\rho^\varepsilon(\boldsymbol{\psi})$ is \mathcal{C}^1 on \mathbb{R}_+^* , and that $\mathcal{K}_\rho^\varepsilon$ is strictly convex on $(\mathbb{1}_N)^\perp$ (see Section 4.3), the implicit function theorem asserts that $\varepsilon \mapsto \boldsymbol{\psi}^\varepsilon$ is a \mathcal{C}^1 function from \mathbb{R}_+^* to $(\mathbb{1}_N)^\perp$. We can therefore differentiate the stationary equation (7.1) with respect to ε and obtain that $\boldsymbol{\psi}^\varepsilon$ satisfies ODE (7.7). \square

Using (7.7), one can see that controlling $\|\dot{\psi}^\varepsilon\|$ amounts to finding a lower bound on the p.s.d. matrix $\nabla^2 \mathcal{K}_\rho^\varepsilon(\psi^\varepsilon)$ and an upper bound on the second term $\frac{\partial}{\partial \varepsilon}(\nabla \mathcal{K}_\rho^\varepsilon)(\psi^\varepsilon)$. The strong-convexity estimate of Theorem 4.4 proven in Chapter 4 readily give the first lower bound. Let's recall the statement of this result here, specialized to the case $\psi = \psi^\varepsilon$:

Theorem (Theorem 4.4). *Let $\rho \in \mathcal{P}(\mathcal{X})$ satisfying Assumption 7.1 and let $\mu \in \mathcal{P}(\mathcal{Y})$. Then for any $\varepsilon > 0$, the solution ψ^ε to problem (D_ε) verifies for any $v \in \mathbb{R}^N$*

$$\mathbb{V}\text{ar}_\mu(v) \leq \left(e^{R_{\mathcal{Y}} \text{diam}(\mathcal{X})} \frac{M_\rho}{m_\rho} + \varepsilon \right) \langle v | \nabla^2 \mathcal{K}_\rho^\varepsilon(\psi^\varepsilon) v \rangle.$$

Remark 7.10 (Dependence on ε). We already know from Chapter 4 that this estimate allows to recover as $\varepsilon \rightarrow 0$ the similar strong-convexity estimate of Chapter 3 for the unregularized Kantorovich functional. This estimate may also be compared to two other similar strong-convexity estimates for entropic optimal transport, found in Theorem 4 of (Luise et al., 2019) (in the discrete context) and in Lemma A.1 of (Bercu and Bigot, 2020) (that is not explicit) that both diverge as ε goes to zero. Note also that as ε goes to ∞ , our estimate deteriorates: in this limit, $\mathcal{K}_\rho^\varepsilon$ gets *flat* around its minimum.

The second control we derive in this chapter gives a uniform bound on the second term of ODE (7.7). It is proven in Section 7.4.

Theorem 7.11. *Let $\rho \in \mathcal{P}(\mathcal{X})$ satisfying Assumption 7.1 with an α -Hölder continuous density for some $\alpha \in (0, 1)$ and let $\mu \in \mathcal{P}(\mathcal{Y})$. Then for any $\varepsilon \leq 1$, $\alpha' \in (0, \alpha)$, the solutions ψ^ε to problem (D_ε) verify:*

$$\left\| \frac{\partial}{\partial \varepsilon}(\nabla \mathcal{K}_\rho^\varepsilon)(\psi^\varepsilon) \right\|_\infty \lesssim \varepsilon^{\alpha'},$$

where \lesssim hides multiplicative constants that depend on $\mathcal{X}, \rho, \mathcal{Y}, \mu$. Besides, for any $\varepsilon \geq 1$,

$$\left\| \frac{\partial}{\partial \varepsilon}(\nabla \mathcal{K}_\rho^\varepsilon)(\psi^\varepsilon) \right\|_\infty \lesssim \frac{1}{\varepsilon}.$$

With these results, the proof of Theorem 7.2 falls directly.

Proof of Theorem 7.2. For any $\varepsilon > 0$, we can apply Proposition 7.9 to ψ^ε and observe the relation

$$\nabla^2 \mathcal{K}_\rho^\varepsilon(\psi^\varepsilon) \dot{\psi}^\varepsilon = -\frac{\partial}{\partial \varepsilon}(\nabla \mathcal{K}_\rho^\varepsilon)(\psi^\varepsilon).$$

Taking the scalar product of the last expression with $\dot{\psi}^\varepsilon$ this gives:

$$\langle \dot{\psi}^\varepsilon | \nabla^2 \mathcal{K}_\rho^\varepsilon(\psi^\varepsilon) \dot{\psi}^\varepsilon \rangle = -\langle \dot{\psi}^\varepsilon | \frac{\partial}{\partial \varepsilon}(\nabla \mathcal{K}_\rho^\varepsilon)(\psi^\varepsilon) \rangle.$$

Applying Theorem 4.4 with $v = \dot{\psi}^\varepsilon$ ensures that

$$\mathbb{V}\text{ar}_\mu(\dot{\psi}^\varepsilon) \leq -\left(e^{R_{\mathcal{Y}} \text{diam}(\mathcal{X})} \frac{M_\rho}{m_\rho} + \varepsilon \right) \langle \dot{\psi}^\varepsilon | \frac{\partial}{\partial \varepsilon}(\nabla \mathcal{K}_\rho^\varepsilon)(\psi^\varepsilon) \rangle. \quad (7.8)$$

Denote $\underline{\mu} > 0$ a positive real such that for all $i \in \{1, \dots, N\}$, $\mu(y_i) \geq \underline{\mu}$. Notice then that the facts that $\mu \geq \underline{\mu} \mathbf{1}_N$ and that $\langle \dot{\psi}^\varepsilon | \mathbf{1}_N \rangle = 0$ (because $\langle \psi^\varepsilon | \mathbf{1}_N \rangle = 0$) entail

$$\begin{aligned} \text{Var}_\mu(\dot{\psi}^\varepsilon) &= \min_{m \in \mathbb{R}} \left\| \dot{\psi}^\varepsilon - m \mathbf{1}_N \right\|_{L^2(\mu)}^2 \\ &\geq \underline{\mu} \min_{m \in \mathbb{R}} \left\| \dot{\psi}^\varepsilon - m \mathbf{1}_N \right\|_2^2 = \underline{\mu} \left\| \dot{\psi}^\varepsilon \right\|_2^2. \end{aligned}$$

Using this inequality together with the Cauchy-Schwartz inequality in equation (7.8) we thus have

$$\begin{aligned} \left\| \dot{\psi}^\varepsilon \right\|_2 \underline{\mu} &\leq \left(e^{R_{\mathcal{Y}} \text{diam}(\mathcal{X})} \frac{M_\rho}{m_\rho} + \varepsilon \right) \left\| \frac{\partial}{\partial \varepsilon} (\nabla \mathcal{K}_\rho^\varepsilon)(\psi^\varepsilon) \right\|_2 \\ &\leq N \left(e^{R_{\mathcal{Y}} \text{diam}(\mathcal{X})} \frac{M_\rho}{m_\rho} + \varepsilon \right) \left\| \frac{\partial}{\partial \varepsilon} (\nabla \mathcal{K}_\rho^\varepsilon)(\psi^\varepsilon) \right\|_\infty. \end{aligned}$$

Applying Theorem 7.11 to the last inequality yields the wanted result. \square

7.4 Proof of Theorem 7.11

In this section, we prove Theorem 7.11 that gives a uniform bound on the second term $\frac{\partial}{\partial \varepsilon} (\nabla \mathcal{K}_\rho^\varepsilon)(\psi^\varepsilon) \in \mathbb{R}^N$ in ODE (7.7). For simplicity, we will use γ_x^ε instead of $\gamma_x^\varepsilon(\psi^\varepsilon)$ since ψ^ε will be the only potential of interest. For any $j \in \{1, \dots, N\}$, we introduce the function

$$f_j^\varepsilon : x \in \mathcal{X} \mapsto \langle x | y_j \rangle - \psi_j^\varepsilon \in \mathbb{R}.$$

Notice that for any $j \in \{1, \dots, N\}$, $\gamma_{x,j}^\varepsilon$ satisfies the following equality and inequality:

$$\gamma_{x,j}^\varepsilon = \frac{\exp\left(\frac{f_j^\varepsilon(x)}{\varepsilon}\right)}{\sum_k \exp\left(\frac{f_k^\varepsilon(x)}{\varepsilon}\right)} \leq \exp\left(\frac{f_j^\varepsilon(x) - \max_\ell f_\ell^\varepsilon(x)}{\varepsilon}\right).$$

Finding a uniform bound on the second term of (7.7) then consists in finding a bound for any $i \in \{1, \dots, N\}$ on the quantity

$$\begin{aligned} \left[\frac{\partial}{\partial \varepsilon} (\nabla \mathcal{K}_\rho^\varepsilon)(\psi^\varepsilon) \right]_i &= \int_{\mathcal{X}} \frac{1}{\varepsilon} [(\text{diag}(\gamma_x^\varepsilon) - \gamma_x^\varepsilon (\gamma_x^\varepsilon)^\top) \log \gamma_x^\varepsilon]_i d\rho(x) \\ &= \int_{\mathcal{X}} \sum_{j \neq i} \left(\frac{f_i^\varepsilon(x) - f_j^\varepsilon(x)}{\varepsilon^2} \right) \gamma_{x,j}^\varepsilon \gamma_{x,i}^\varepsilon d\rho(x). \end{aligned} \quad (7.9)$$

Recall that the first order condition (7.2) entails that $e^{\frac{\psi_i^\varepsilon}{\varepsilon}} \mu_i = \int_{\mathcal{X}} e^{\frac{\langle x | y_i \rangle - (\psi^\varepsilon)^{c,\varepsilon}(x)}{\varepsilon}} d\rho(x)$, which ensures $\left| \psi_i^\varepsilon - \psi_j^\varepsilon \right| \leq R_{\mathcal{X}} |y_i - y_j| + \varepsilon \left| \log\left(\frac{\mu_i}{\mu_j}\right) \right|$ (see e.g. the proof of Proposition 7.13). Thus for any $x \in \mathcal{X}$ and $j \in \{1, \dots, N\}$,

$$\left| f_i^\varepsilon(x) - f_j^\varepsilon(x) \right| \leq 2R_{\mathcal{X}} \text{diam}(\mathcal{Y}) + \varepsilon \left| \log(\underline{\mu}) \right|.$$

Hence for $\varepsilon \geq 1$,

$$\left| \left[\frac{\partial}{\partial \varepsilon} (\nabla \mathcal{K}_\rho^\varepsilon)(\psi) \right]_i \right| \lesssim \frac{1}{\varepsilon}.$$

We now look for a more informative bound in the limit $\varepsilon \rightarrow 0$. The quantity (7.9) being an integral over \mathcal{X} , it will be in our interest to partition \mathcal{X} into different subdomains where we can control the integrand. To this end, the Laguerre tessellation $\bigcup_i \text{Lag}_i(\psi^\varepsilon)$ already provides a first interesting partition. Recall the definition of the Laguerre cells with our new notation:

$$\text{Lag}_i(\psi^\varepsilon) = \{x \in \mathcal{X} \mid \forall j, f_i^\varepsilon(x) \geq f_j^\varepsilon(x)\}.$$

Figure 7.1 gives an illustration of Laguerre cells, where the boundary of those cells are indicated by the plain black lines. In the control of (7.9), we will see that for x in the interior of $\text{Lag}_i(\psi^\varepsilon)$, $\gamma_{x,j}^\varepsilon$ is *very small* for any $j \neq i$, and conversely for x far from $\text{Lag}_i(\psi^\varepsilon)$, $\gamma_{x,i}^\varepsilon$ is *very small*. We will thus introduce two sets of points $\mathcal{X}_{i,\eta,+}^\varepsilon, \mathcal{X}_{i,\eta,-}^\varepsilon$ corresponding respectively to the points of $\text{Lag}_i(\psi^\varepsilon)$ and $\mathcal{X} \setminus \text{Lag}_i(\psi^\varepsilon)$ that are at a distance at least η from the boundary of $\text{Lag}_i(\psi^\varepsilon)$. These sets are illustrated in Figure 7.1 in green and blue respectively. Now for x close from the boundary of $\text{Lag}_i(\psi^\varepsilon)$ (i.e. $x \in \mathcal{X} \setminus (\mathcal{X}_{i,\eta,+}^\varepsilon \cup \mathcal{X}_{i,\eta,-}^\varepsilon)$), we will see that $\gamma_{x,i}^\varepsilon$ cannot be too small and there always exists a $j \neq i$ such that $\gamma_{x,j}^\varepsilon$ is not small neither. A finer treatment of those points close from the boundary of $\text{Lag}_i(\psi^\varepsilon)$ has to be carried out. For some $\zeta > 0$, we will first define a set $A_{i,\eta,\zeta}^\varepsilon$ of points that lie near the intersection between $\text{Lag}_i(\psi^\varepsilon)$ and another cell $\text{Lag}_j(\psi^\varepsilon)$, but that are at a distance at least $\sim \zeta$ from the other cells (i.e. $A_{i,\eta,\zeta}^\varepsilon$ excludes the *corners* of $\text{Lag}_i(\psi^\varepsilon)$). This set is represented in yellow in Figure 7.1. On this set, only $\gamma_{x,i}^\varepsilon$ and $\gamma_{x,j}^\varepsilon$ are *not small* and we show that their contributions to integral (7.9) get compensated by symmetry w.r.t. the interface $\text{Lag}_i(\psi^\varepsilon) \cap \text{Lag}_j(\psi^\varepsilon)$. Finally, we will denote the rest of \mathcal{X} by $B_{i,\eta,\zeta}^\varepsilon = \mathcal{X} \setminus (\mathcal{X}_{i,\eta,+}^\varepsilon \cup \mathcal{X}_{i,\eta,-}^\varepsilon \cup A_{i,\eta,\zeta}^\varepsilon)$: this set corresponds to the areas in red in Figure 7.1. We will control (7.9) on this set by leveraging two facts: its points are close from $\text{Lag}_i(\psi^\varepsilon)$ and its volume scales as $\sim \zeta^2$.




Figure 7.1: Partition of \mathcal{X} : $\mathcal{X}_{i,\eta,+}^\varepsilon$ is in green, $\mathcal{X}_{i,\eta,-}^\varepsilon$ is in blue, $A_{i,\eta,\zeta}^\varepsilon$ is in yellow and $B_{i,\eta,\zeta}^\varepsilon$ is in red.

We make precise the definitions of the above mentioned sets in the proof of the following proposition, that allows to get the following bound:

Proposition 7.12. *Under assumptions of Theorem 7.11, for $0 < \varepsilon \leq 1, i \in \{1, \dots, N\}$*

and for any $\eta, \zeta > 0$,

$$\begin{aligned} \left| \left[\frac{\partial}{\partial \varepsilon} (\nabla \mathcal{K}_\rho^\varepsilon)(\psi) \right]_i \right| &\lesssim \frac{1}{\varepsilon^2} e^{-\eta/\varepsilon} + \frac{\eta^{2+\alpha}}{\varepsilon^2} + \frac{\zeta^2}{\varepsilon^2} \left(\eta + e^{-\eta/\varepsilon} \right) \\ &\quad + \frac{1}{\varepsilon^2} e^{-\tilde{\zeta}/\varepsilon} \left(\eta + \varepsilon \eta e^{\eta/\varepsilon} - \varepsilon^2 (e^{\eta/\varepsilon} - 1) \right), \end{aligned}$$

where $\tilde{\zeta} = \zeta \delta - \frac{\text{diam}(\mathcal{Y})^2}{\delta} \eta$ and $\delta = \min_{i \neq j} \|y_i - y_j\| > 0$.

Before proving Proposition 7.12, let us mention that the proof of the " $\varepsilon \leq 1$ " side of Theorem 7.11 follows from this proposition and an arbitrage on the quantities ζ, η :

Proof of Theorem 7.11. Here we assume that $\varepsilon \leq 1$. In the inequality on $\left| \left[\frac{\partial}{\partial \varepsilon} (\nabla \mathcal{K}_\rho^\varepsilon)(\psi) \right]_i \right|$ provided by Proposition 7.12, we set $\zeta = \frac{\eta}{\delta} \left(\frac{\text{diam}(\mathcal{Y})^2}{\delta} + 2 \right)$ for some $\eta > 0$. This yields

$$\left| \left[\frac{\partial}{\partial \varepsilon} (\nabla \mathcal{K}_\rho^\varepsilon)(\psi) \right]_i \right| \lesssim \frac{\eta^{2+\alpha} + \eta^3}{\varepsilon^2} + \frac{e^{-\eta/\varepsilon}}{\varepsilon^2} \left(1 + \eta^2 + \varepsilon \eta - \varepsilon^2 + (\eta + \varepsilon^2) e^{-\eta/\varepsilon} \right).$$

Then, for any $\beta \in (\frac{2}{2+\alpha}, 1)$, choosing $\eta = \varepsilon^\beta$ yields

$$\left| \left[\frac{\partial}{\partial \varepsilon} (\nabla \mathcal{K}_\rho^\varepsilon)(\psi) \right]_i \right| \lesssim \varepsilon^{(2+\alpha)\beta-2} + \frac{e^{-1/\varepsilon^{1-\beta}}}{\varepsilon^2}.$$

With $\alpha' = (2 + \alpha)\beta - 2$, we get that for any $\alpha' \in (0, \alpha)$,

$$\left| \left[\frac{\partial}{\partial \varepsilon} (\nabla \mathcal{K}_\rho^\varepsilon)(\psi) \right]_i \right| \lesssim \varepsilon^{\alpha'} + \frac{e^{-1/\varepsilon^{\frac{\alpha-\alpha'}{2+\alpha}}}}{\varepsilon^2} \lesssim \varepsilon^{\alpha'}. \quad \square$$

Proof of Proposition 7.12. We introduce for any $i \in \{1, \dots, N\}$ and parameter $\eta > 0$ the sets:

$$\begin{aligned} \mathcal{X}_{i,\eta,+}^\varepsilon &= \{x \in \text{Lag}_i(\psi^\varepsilon) \mid \forall j \neq i, \frac{f_i^\varepsilon(x) - f_j^\varepsilon(x)}{\|y_i - y_j\|} \geq \eta\}, \\ \mathcal{X}_{i,\eta,-}^\varepsilon &= \{x \in \mathcal{X} \mid \forall j \in \arg \max_{\ell} f_\ell^\varepsilon(x), \frac{f_j^\varepsilon(x) - f_i^\varepsilon(x)}{\|y_j - y_i\|} \geq \eta\}, \end{aligned}$$

that correspond respectively to the points of $\text{Lag}_i(\psi^\varepsilon)$, $\mathcal{X} \setminus \text{Lag}_i(\psi^\varepsilon)$ that are at a distance at least η from the boundary of $\text{Lag}_i(\psi^\varepsilon)$ and that are illustrated in green and in blue in Figure 7.1. We then define for any $j \neq i$ the common boundary between $\text{Lag}_i(\psi^\varepsilon)$ and $\text{Lag}_j(\psi^\varepsilon)$:

$$H_{ij} = \text{Lag}_i(\psi^\varepsilon) \cap \text{Lag}_j(\psi^\varepsilon).$$

Next, for a parameter $\zeta > 0$, define the set of points of H_{ij} that are at a distance at least ζ from the other Laguerre cells:

$$H_{ij}^{-\zeta} = \{x^0 \in H_{ij} \mid \forall k \neq i, j, f_i^\varepsilon(x^0) = f_j^\varepsilon(x^0) \geq f_k^\varepsilon(x^0) + \zeta \max(\|y_i - y_k\|, \|y_j - y_k\|)\}.$$

Then define

$$A_{i,\eta,\zeta}^\varepsilon = \bigcup_{j \neq i} \{x^0 + t d_{ij}, x^0 \in H_{ij}^{-\zeta}, t \in [-\eta \|y_i - y_j\|, +\eta \|y_i - y_j\|]\}$$

where $d_{ij} = \frac{y_i - y_j}{\|y_i - y_j\|^2}$. This set corresponds to the areas in yellow in Figure 7.1. Define finally $B_{i,\eta,\zeta}^\varepsilon = \mathcal{X} \setminus (\mathcal{X}_{i,\eta,+}^\varepsilon \cup \mathcal{X}_{i,\eta,-}^\varepsilon \cup A_{i,\eta,\zeta}^\varepsilon)$: this set corresponds to the areas in red in Figure 7.1.

Control on $\mathcal{X}_{i,\eta,+}^\varepsilon$. For $x \in \mathcal{X}_{i,\eta,+}^\varepsilon$, we have for any $j \neq i$,

$$\gamma_{x,j}^\varepsilon \leq \exp\left(\frac{f_j^\varepsilon(x) - f_i^\varepsilon(x)}{\varepsilon}\right) \leq e^{-\eta\|y_i - y_j\|/\varepsilon} \leq e^{-\eta\delta/\varepsilon}.$$

This gives in equation (7.9) the control

$$\forall x \in \mathcal{X}_{i,\eta,+}^\varepsilon, \quad \sum_{j=1, j \neq i}^N \left(\frac{f_i^\varepsilon(x) - f_j^\varepsilon(x)}{\varepsilon^2}\right) \gamma_{x,j}^\varepsilon \gamma_{x,i}^\varepsilon \lesssim \frac{1}{\varepsilon^2} e^{-\eta\delta/\varepsilon} \lesssim \frac{1}{\varepsilon^2} e^{-\eta/\varepsilon}. \quad (7.10)$$

Control on $\mathcal{X}_{i,\eta,-}^\varepsilon$. For $x \in \mathcal{X}_{i,\eta,-}^\varepsilon$, we have

$$\gamma_{x,i}^\varepsilon \leq \exp\left(\frac{f_i^\varepsilon(x) - \max_j f_j^\varepsilon(x)}{\varepsilon}\right) \leq e^{-\eta\delta/\varepsilon}.$$

This gives in equation (7.9) the control

$$\forall x \in \mathcal{X}_{i,\eta,-}^\varepsilon, \quad \sum_{j=1, j \neq i}^N \left(\frac{f_i^\varepsilon(x) - f_j^\varepsilon(x)}{\varepsilon^2}\right) \gamma_{x,j}^\varepsilon \gamma_{x,i}^\varepsilon \lesssim \frac{1}{\varepsilon^2} e^{-\eta/\varepsilon}. \quad (7.11)$$

Control on $A_{i,\eta,\zeta}^\varepsilon$. For any $x \in A_{i,\eta,\zeta}^\varepsilon$, there exists $j \in \{1, \dots, N\}$, $x^0 \in H_{ij}^{-\zeta}$ and $t \in [-\eta\|y_i - y_j\|, +\eta\|y_i - y_j\|]$ such that

$$x = x^0 + td_{ij}.$$

For such a point, we have

$$\begin{aligned} f_i^\varepsilon(x) - f_j^\varepsilon(x) &= \langle x^0 + td_{ij} | y_i - y_j \rangle - \psi_i^\varepsilon + \psi_j^\varepsilon \\ &= f_i^\varepsilon(x^0) - f_j^\varepsilon(x^0) + t \langle d_{ij} | y_i - y_j \rangle \\ &= t. \end{aligned} \quad (7.12)$$

Moreover, for any $k \neq i, j$ we have by definition of $H_{ij}^{-\zeta}$

$$\begin{aligned} f_i^\varepsilon(x) - f_k^\varepsilon(x) &= f_i^\varepsilon(x^0) - f_k^\varepsilon(x^0) + t \langle d_{ij} | y_i - y_k \rangle \\ &\geq \zeta \|y_i - y_k\| - |t| \frac{\text{diam}(\mathcal{Y})}{\delta} \\ &\geq \zeta \delta - \frac{\text{diam}(\mathcal{Y})^2}{\delta} \eta := \tilde{\zeta}. \end{aligned} \quad (7.13)$$

In the same way,

$$f_j^\varepsilon(x) - f_k^\varepsilon(x) \geq \zeta \delta - \frac{\text{diam}(\mathcal{Y})^2}{\delta} \eta = \tilde{\zeta}. \quad (7.14)$$

The integral we want to control on $A_{i,\eta,\zeta}^\varepsilon$ reads:

$$\sum_j \int_{x^0 \in H_{ij}^{-\zeta}} \int_{t=0}^{\eta\|y_i - y_j\|} (g_i^\varepsilon(x^0 - td_{ij}) + g_i^\varepsilon(x^0 + td_{ij})) dt d\mathcal{H}^{d-1}(x^0) \quad (7.15)$$

where $g_i^\varepsilon(x) = \sum_{j=1, j \neq i}^N \left(\frac{f_i^\varepsilon(x) - f_j^\varepsilon(x)}{\varepsilon^2} \right) \gamma_{x,j}^\varepsilon \gamma_{x,i}^\varepsilon \rho(x)$.

Let's find an upper bound on $|g_i^\varepsilon(x^0 - td_{ij}) + g_i^\varepsilon(x^0 + td_{ij})|$. To simplify the notation, denote $x^- = x^0 - td_{ij}$ and $x^+ = x^0 + td_{ij}$. Using equation (7.12), we have the expression

$$g_i^\varepsilon(x^-) = \left(\frac{f_i^\varepsilon(x^-) - f_j^\varepsilon(x^-)}{\varepsilon^2} \right) \gamma_{x^-,j}^\varepsilon \gamma_{x^-,i}^\varepsilon \rho(x^-) \quad (7.16)$$

$$\begin{aligned} &+ \sum_{k \neq i,j} \left(\frac{f_i^\varepsilon(x^-) - f_k^\varepsilon(x^-)}{\varepsilon^2} \right) \gamma_{x^-,k}^\varepsilon \gamma_{x^-,i}^\varepsilon \rho(x^-) \\ &= \frac{-t}{\varepsilon^2} \gamma_{x^-,j}^\varepsilon \gamma_{x^-,i}^\varepsilon \rho(x^-) + S(x^-), \end{aligned} \quad (7.17)$$

where $S(x) = \sum_{k \neq i,j} \left(\frac{f_i^\varepsilon(x) - f_k^\varepsilon(x)}{\varepsilon^2} \right) \gamma_{x,k}^\varepsilon \gamma_{x,i}^\varepsilon \rho(x)$. Notice that from (7.13) and (7.14), for $k \neq i, j$, $\gamma_{x^-,k}^\varepsilon \leq e^{-\tilde{\zeta}/\varepsilon}$. This gives the bound

$$|S(x^-)| \lesssim \frac{1}{\varepsilon^2} e^{-\tilde{\zeta}/\varepsilon}. \quad (7.18)$$

Similarly, we have

$$g_i^\varepsilon(x^+) = \frac{t}{\varepsilon^2} \gamma_{x^+,j}^\varepsilon \gamma_{x^+,i}^\varepsilon \rho(x^+) + S(x^+), \quad (7.19)$$

where $S(x^+)$ verifies

$$|S(x^+)| \lesssim \frac{1}{\varepsilon^2} e^{-\tilde{\zeta}/\varepsilon}. \quad (7.20)$$

From equations (7.17) and (7.19), we thus have the control

$$\begin{aligned} |g_i^\varepsilon(x^-) + g_i^\varepsilon(x^+)| &\leq \frac{t}{\varepsilon^2} \left| \gamma_{x^+,j}^\varepsilon \gamma_{x^+,i}^\varepsilon \rho(x^+) - \gamma_{x^-,j}^\varepsilon \gamma_{x^-,i}^\varepsilon \rho(x^-) \right| + |S(x^-)| + |S(x^+)| \\ &\leq \frac{t}{\varepsilon^2} \left| \gamma_{x^+,j}^\varepsilon \gamma_{x^+,i}^\varepsilon - \gamma_{x^-,j}^\varepsilon \gamma_{x^-,i}^\varepsilon \right| \rho(x^+) + \frac{t}{\varepsilon^2} \gamma_{x^-,j}^\varepsilon \gamma_{x^-,i}^\varepsilon \left| \rho(x^+) - \rho(x^-) \right| \\ &\quad + |S(x^-)| + |S(x^+)| \end{aligned} \quad (7.21)$$

Now, notice that

$$\begin{aligned} \gamma_{x^+,i}^\varepsilon &= \frac{\exp\left(\frac{f_i^\varepsilon(x^+)}{\varepsilon}\right)}{\exp\left(\frac{f_i^\varepsilon(x^+)}{\varepsilon}\right) + \exp\left(\frac{f_j^\varepsilon(x^+)}{\varepsilon}\right) + \sum_{k \neq i,j} \exp\left(\frac{f_k^\varepsilon(x^+)}{\varepsilon}\right)} \\ &= \frac{1}{1 + e^{-t/\varepsilon} + S_i(x^+)}, \end{aligned}$$

where $S_\ell(x) = \sum_{k \neq i,j} \exp\left(\frac{f_k^\varepsilon(x) - f_\ell^\varepsilon(x)}{\varepsilon}\right)$. Similarly, we have

$$\begin{aligned} \gamma_{x^+,j}^\varepsilon &= \frac{1}{1 + e^{t/\varepsilon} + S_j(x^+)}, \\ \gamma_{x^-,i}^\varepsilon &= \frac{1}{1 + e^{t/\varepsilon} + S_i(x^-)}, \\ \gamma_{x^-,j}^\varepsilon &= \frac{1}{1 + e^{-t/\varepsilon} + S_j(x^-)}. \end{aligned}$$

Moreover, remark that for $\ell \in \{i, j\}$ and $x \in \{x^-, x^+\}$ we have the bound

$$S_\ell(x) \lesssim e^{-\tilde{\zeta}/\varepsilon}.$$

From these expressions we deduce the following bound:

$$\left| \gamma_{x^+,j}^\varepsilon \gamma_{x^+,i}^\varepsilon - \gamma_{x^-,j}^\varepsilon \gamma_{x^-,i}^\varepsilon \right| \lesssim e^{-\tilde{\zeta}/\varepsilon} e^{t/\varepsilon}. \quad (7.22)$$

Now using that ρ is α -Hölder, we know that there exists a constant $C_\rho > 0$ such that

$$|\rho(x^+) - \rho(x^-)| \leq C_\rho \|x^+ - x^-\|^\alpha = C_\rho \|2td_{ij}\|^\alpha \leq C_\rho \left(\frac{2}{\delta}\right)^\alpha t^\alpha. \quad (7.23)$$

Plugging the bounds (7.18), (7.20), (7.22) and (7.23) into (7.21) then yields:

$$|g_i^\varepsilon(x^+) + g_i^\varepsilon(x^-)| \lesssim \frac{1}{\varepsilon^2} \left(te^{-\tilde{\zeta}/\varepsilon} e^{t/\varepsilon} + t^{1+\alpha} + e^{-\tilde{\zeta}/\varepsilon} \right).$$

Injecting these bounds into integral (7.15) entails

$$\begin{aligned} \left| \int_{A_{i,\eta,\zeta}^\varepsilon} \frac{1}{\varepsilon} \sum_{j=1, j \neq i}^N \left(\frac{f_i^\varepsilon(x) - f_j^\varepsilon(x)}{\varepsilon^2} \right) \gamma_{x,j}^\varepsilon \gamma_{x,i}^\varepsilon \rho(x) dx \right| &\lesssim \int_0^{\eta \text{diam}(\mathcal{Y})} \frac{1}{\varepsilon^2} \left(te^{-\tilde{\zeta}/\varepsilon} e^{t/\varepsilon} + t^{1+\alpha} + e^{-\tilde{\zeta}/\varepsilon} \right) dt \\ &\lesssim \frac{\eta^{2+\alpha}}{\varepsilon^2} + \frac{1}{\varepsilon^2} e^{-\tilde{\zeta}/\varepsilon} \left(\eta + \varepsilon \eta e^{\eta/\varepsilon} - \varepsilon^2 (e^{\eta/\varepsilon} - 1) \right). \end{aligned} \quad (7.24)$$

Control on $B_{i,\eta,\zeta}^\varepsilon$. We first derive a uniform bound on the integrand

$$\sum_{j=1, j \neq i}^N \left(\frac{f_i^\varepsilon(x) - f_j^\varepsilon(x)}{\varepsilon^2} \right) \gamma_{x,j}^\varepsilon \gamma_{x,i}^\varepsilon$$

on the domain $B_{i,\eta,\zeta}^\varepsilon$, that is included in the η -neighborhood of $\text{Lag}_i(\psi^\varepsilon)$. For $x \in B_{i,\eta,\zeta}^\varepsilon$, for any $j \neq i$, either $|f_i^\varepsilon(x) - f_j^\varepsilon(x)| \leq \eta(\text{diam}(\mathcal{Y}) + 1)$, and in this case

$$\left| \frac{f_i^\varepsilon(x) - f_j^\varepsilon(x)}{\varepsilon^2} \right| \gamma_{x,j}^\varepsilon \gamma_{x,i}^\varepsilon \lesssim \frac{\eta}{\varepsilon^2},$$

or $|f_i^\varepsilon(x) - f_j^\varepsilon(x)| > \eta(\text{diam}(\mathcal{Y}) + 1)$, which entails $\gamma_{x,j}^\varepsilon \leq e^{-\eta/\varepsilon}$. To see this, denote $k \in \arg \max_\ell f_\ell^\varepsilon(x)$. Since x is in a η -neighborhood of $\text{Lag}_i(\psi^\varepsilon)$, we have

$$0 \leq f_k^\varepsilon(x) - f_i^\varepsilon(x) \leq \eta \|y_k - y_i\| \leq \eta \text{diam}(\mathcal{Y}).$$

Hence

$$|f_j^\varepsilon(x) - f_k^\varepsilon(x)| \geq ||f_j^\varepsilon(x) - f_i^\varepsilon(x)| - |f_k^\varepsilon(x) - f_i^\varepsilon(x)|| \geq \eta.$$

The inequality $\gamma_{x,j}^\varepsilon \leq e^{-\eta/\varepsilon}$ then comes from the fact that $\gamma_{x,j}^\varepsilon \leq \exp\left(\frac{f_j^\varepsilon(x) - f_k^\varepsilon(x)}{\varepsilon}\right)$. From these remarks we can therefore write for $x \in B_{i,\eta,\zeta}^\varepsilon$:

$$\sum_{j=1, j \neq i}^N \left(\frac{f_i^\varepsilon(x) - f_j^\varepsilon(x)}{\varepsilon^2} \right) \gamma_{x,j}^\varepsilon \gamma_{x,i}^\varepsilon \lesssim \frac{1}{\varepsilon^2} \left(\eta + e^{-\eta/\varepsilon} \right).$$

Finally, notice that $B_{i,\eta,\zeta}^\varepsilon$ is made of a union of *corners* of η -neighborhoods $\text{Lag}_i(\psi^\varepsilon)$, where *corner* is meant for intersection of 2 hyperplanes. There are at most $\binom{N-1}{2} \leq N^2$ such corners. Denote $\theta = \arg \max_{i,j,k} |\angle y_i y_j y_k| < \gamma$, i.e. the maximum angle that can be formed from a triplet of points in the support of the target that do not lie on a same line. Then the *corners* that constitute $B_{i,\eta,\zeta}^\varepsilon$ are actually included in *cylinders* of length at most $\text{diam}(\mathcal{X})$ and of radius at most $\frac{2\zeta}{\cos(\theta/2)}$, that is of volume at most $\frac{4\pi \text{diam}(\mathcal{X})^{d-2}}{\cos(\theta/2)^2} \zeta^2$. All these considerations allow us to write the following bound:

$$\left| \int_{B_{i,\eta,\zeta}^\varepsilon} \frac{1}{\varepsilon} \sum_{j=1, j \neq i}^N \left(\frac{f_i^\varepsilon(x) - f_j^\varepsilon(x)}{\varepsilon^2} \right) \gamma_{x,j}^\varepsilon \gamma_{x,i}^\varepsilon \rho(x) dx \right| \lesssim \frac{\zeta^2}{\varepsilon^2} \left(\eta + e^{-\eta/\varepsilon} \right). \quad (7.25)$$

Conclusion. Finally, using equations (7.10), (7.11), (7.24), (7.25) we get the wanted control:

$$\begin{aligned} \left| \left[\frac{\partial}{\partial \varepsilon} (\nabla \mathcal{K}_\rho^\varepsilon)(\psi) \right]_i \right| &= \left| \int_{\mathcal{X}} \frac{1}{\varepsilon} \sum_{j=1, j \neq i}^N \left(\frac{f_i^\varepsilon(x) - f_j^\varepsilon(x)}{\varepsilon^2} \right) \gamma_{x,j}^\varepsilon \gamma_{x,i}^\varepsilon \rho(x) dx \right| \\ &\lesssim \frac{1}{\varepsilon^2} e^{-\eta/\varepsilon} + \frac{\eta^{2+\alpha}}{\varepsilon^2} + \frac{\zeta^2}{\varepsilon^2} \left(\eta + e^{-\eta/\varepsilon} \right) + \frac{1}{\varepsilon^2} e^{-\tilde{\zeta}/\varepsilon} \left(\eta + \varepsilon \eta e^{\eta/\varepsilon} - \varepsilon^2 (e^{\eta/\varepsilon} - 1) \right). \quad \square \end{aligned}$$

7.5 Proof of Corollary 7.5

Before proving Corollary 7.5, let us prove the following proposition that ensures the convergence of ψ^ε to ψ^0 as ε goes to zero. This fact holds in much more general settings (Nutz and Wiesel, 2021) and we include its proof in our context only for completeness.

Proposition 7.13. *The solutions ψ^ε to problem (D_ε) verify:*

$$\lim_{\varepsilon \rightarrow 0} \psi^\varepsilon = \psi^0.$$

Proof. Let's first prove that for any $\varepsilon \geq 0$, the solution ψ^ε to problem (D_ε) verifies:

$$\|\psi^\varepsilon\|_\infty \leq R_{\mathcal{X}} \text{diam}(\mathcal{Y}) + \varepsilon \log(1/\underline{\mu}).$$

For $\varepsilon > 0$, recall that the first order condition (7.2) entails that

$$e^{\frac{\psi_i^\varepsilon}{\varepsilon}} \mu_i = \int_{\mathcal{X}} e^{\frac{\langle x | y_i \rangle - (\psi^\varepsilon)^{c,\varepsilon}(x)}{\varepsilon}} d\rho(x).$$

The right-hand side of this equality corresponds to a $R_{\mathcal{X}}$ -Lipschitz function of y_i . Therefore we have the bound:

$$|\psi_i^\varepsilon - \psi_j^\varepsilon| \leq R_{\mathcal{X}} |y_i - y_j| + \varepsilon \left| \log\left(\frac{\mu_i}{\mu_j}\right) \right| \leq R_{\mathcal{X}} \text{diam}(\mathcal{Y}) + \varepsilon \log(1/\underline{\mu}). \quad (7.26)$$

Now recall that $\langle \psi^\varepsilon | \mathbb{1}_N \rangle = 0$, which means that the components of $\psi^\varepsilon \in \mathbb{R}^N$ take both positive and negative values. This entails for any $i \in \{1, \dots, N\}$:

$$|\psi_i^\varepsilon| = |\psi_i^\varepsilon - 0| \leq \max_j |\psi_i^\varepsilon - \psi_j^\varepsilon| \leq R_{\mathcal{X}} \text{diam}(\mathcal{Y}) + \varepsilon \log(1/\underline{\mu}).$$

When $\varepsilon = 0$, the bound (7.26) comes from the fact that ψ^0 is a $R_{\mathcal{X}}$ -Lipschitz function from \mathcal{Y} to \mathbb{R} because the dual solution is equal to its biconjugate: for any $i \in \{1, \dots, N\}$, $\psi_i^0 = ((\psi^0)^*)^*(y_i)$. We conclude similarly to the case $\varepsilon > 0$ to ensure $\|\psi^0\|_{\infty} \leq R_{\mathcal{X}} \text{diam}(\mathcal{Y})$.

Now consider a sequence $(\varepsilon_k)_k > 0$ such that $\lim_{k \rightarrow \infty} \varepsilon_k = 0$. By what precedes, the sequence $(\psi^{\varepsilon_k})_k$ is bounded and one can extract a converging subsequence (that we do not relabel). Notice now that for any $x \in \mathcal{X}$ and $\psi \in \mathbb{R}^N$, the (c, ε) -transform $\psi^{c, \varepsilon}(x)$ corresponds to a rescaled LogSumExp (or smooth maximum) of the vector $(\langle x|y_i \rangle - \psi_i)_{i=1, \dots, N}$:

$$\psi^{c, \varepsilon}(x) = \varepsilon LSE \left(\frac{(\langle x|y_i \rangle - \psi_i)_{i=1, \dots, N}}{\varepsilon} \right),$$

where $LSE(z_1, \dots, z_N) = \log(\exp(z_1) + \dots + \exp(z_N))$. Bounds on LSE allow us to write that for any $\psi \in \mathbb{R}^N$, $x \in \mathcal{X}$ and $\varepsilon > 0$ we have

$$\psi^*(x) \leq \psi^{c, \varepsilon}(x) \leq \psi^*(x) + \varepsilon \log N,$$

where we recall that $\psi^*(x) = \max_{i=1, \dots, N} \langle x|y_i \rangle - \psi_i$ denotes the Legendre transform of ψ evaluated in x . Thus if we consider $k \in \mathbb{N}$, we have by optimality of ψ^0 , ψ^{ε_k} for their respective problems the inequalities:

$$\begin{aligned} \langle (\psi^0)^* | \rho \rangle + \langle \psi^0 | \mu \rangle &\leq \langle (\psi^{\varepsilon_k})^* | \rho \rangle + \langle \psi^{\varepsilon_k} | \mu \rangle \\ &\leq \langle (\psi^{\varepsilon_k})^{c, \varepsilon_k} | \rho \rangle + \langle \psi^{\varepsilon_k} | \mu \rangle + \varepsilon_k \\ &\leq \langle (\psi^0)^{c, \varepsilon_k} | \rho \rangle + \langle \psi^0 | \mu \rangle + \varepsilon_k \\ &\leq \langle (\psi^0)^* | \rho \rangle + \langle \psi^0 | \mu \rangle + \varepsilon_k(1 + \log N). \end{aligned}$$

Hence we have the limit:

$$\lim_{k \rightarrow \infty} \langle (\psi^{\varepsilon_k})^* | \rho \rangle + \langle \psi^{\varepsilon_k} | \mu \rangle = \langle (\psi^0)^* | \rho \rangle + \langle \psi^0 | \mu \rangle.$$

By uniqueness of the solution of the unregularized problem on $(\mathbf{1}_N)^\perp$, this proves that

$$\lim_{k \rightarrow \infty} \psi^{\varepsilon_k} = \psi^0.$$

There is thus only one accumulation point for the bounded sequence $(\psi^{\varepsilon_k})_k$, which shows that the whole sequence converges to this point. \square

Proof of Corollary 7.5. Following Theorem 7.2, let $C > 0$ (depending on $\mathcal{X}, \rho, \mathcal{Y}, \mu$) be such that $\|\dot{\psi}^\tau\|_2 \leq C\tau^{\alpha'}$ for $\tau \in [\varepsilon', \varepsilon]$. Then notice

$$\|\psi^\varepsilon - \psi^{\varepsilon'}\|_{\infty} \leq \|\psi^\varepsilon - \psi^{\varepsilon'}\|_2 = \left\| \int_{\varepsilon'}^{\varepsilon} \dot{\psi}^\tau d\tau \right\|_2 \leq \int_{\varepsilon'}^{\varepsilon} \|\dot{\psi}^\tau\|_2 d\tau \leq C\varepsilon^{\alpha'}(\varepsilon - \varepsilon').$$

Letting ε' go to 0 and using Proposition 7.13 yields

$$\|\psi^\varepsilon - \psi^0\|_{\infty} \leq C\varepsilon^{1+\alpha'}.$$

For the second result, we use

$$\|(\psi^\varepsilon)^{c, \varepsilon} - (\psi^0)^*\|_{\infty} \leq \|(\psi^\varepsilon)^{c, \varepsilon} - (\psi^0)^{c, \varepsilon}\|_{\infty} + \|(\psi^0)^{c, \varepsilon} - (\psi^0)^*\|_{\infty}$$

One can easily show with the definition of the (c, ε) -transform that $\|\psi^\varepsilon - \psi^0\|_\infty \leq C\varepsilon^{1+\alpha'}$ entails

$$\|(\psi^\varepsilon)^{c,\varepsilon} - (\psi^0)^{c,\varepsilon}\|_\infty \leq C\varepsilon^{1+\alpha'}.$$

On the other hand, $\|(\psi^0)^{c,\varepsilon} - (\psi^0)^*(x)\|_\infty \leq \varepsilon \log N$ is a LogSumExp property (see the above proof of Proposition 7.13). This property can be refined to get to the third result: we have for all $x \in \mathcal{X}$

$$(\psi^0)^*(x) \leq (\psi^0)^{c,\varepsilon}(x) = (\psi^0)^*(x) + \varepsilon \log \left(\sum_{j=1}^N e^{\frac{\langle x|y_j \rangle - \psi_j^0 - (\psi^0)^*(x)}{\varepsilon}} \right).$$

But for ρ -almost every $x \in \mathcal{X}$, there is only one $i \in \{1, \dots, N\}$ that satisfies $(\psi^0)^*(x) = \langle x|y_i \rangle - \psi_i^0$. Thus for such x , denoting $c_x = \min_{j \neq i} (\langle x|y_i \rangle - \psi_i^0) - (\langle x|y_j \rangle - \psi_j^0) > 0$, we have

$$\sum_{j=1}^N e^{\frac{\langle x|y_j \rangle - \psi_j^0 - (\psi^0)^*(x)}{\varepsilon}} \leq 1 + (N-1)e^{-c_x/\varepsilon}.$$

We thus get:

$$(\psi^0)^*(x) \leq (\psi^0)^{c,\varepsilon}(x) \leq (\psi^0)^*(x) + (N-1)\varepsilon e^{-c_x/\varepsilon}.$$

From this we deduce that for ρ -a.e. $x \in \mathcal{X}$,

$$|(\psi^0)^{c,\varepsilon}(x) - (\psi^0)^*(x)| \lesssim \varepsilon e^{-c_x/\varepsilon} \lesssim \varepsilon^{1+\alpha'}.$$

Finally, we use the notation of Section 7.4 that denotes

$$\frac{d\gamma^\varepsilon}{d\rho \otimes \sigma}(x, y_i) = \gamma_{x,i}^\varepsilon.$$

Let $x \in \mathcal{X}$ be such that $c_x > 0$. Notice that for any $i \in \{1, \dots, N\}$,

$$\gamma_{x,i}^0 = \begin{cases} 1 & \text{if } \langle x|y_i \rangle - \psi_i^0 \geq \langle x|y_j \rangle - \psi_j^0 \quad \forall j, \\ 0 & \text{else.} \end{cases}$$

If $\gamma_{x,i}^0 = 0$, then using $\|\psi^\varepsilon - \psi^0\|_\infty \leq C\varepsilon^{1+\alpha'}$ and $\sum_j e^{\frac{\langle x|y_j \rangle - \psi_j^0}{\varepsilon}} \geq e^{\frac{(\psi^0)^*(x)}{\varepsilon}}$,

$$|\gamma_{x,i}^\varepsilon - \gamma_{x,i}^0| = \gamma_{x,i}^\varepsilon = \frac{e^{\frac{\langle x|y_i \rangle - \psi_i^\varepsilon}{\varepsilon}}}{\sum_j e^{\frac{\langle x|y_j \rangle - \psi_j^\varepsilon}{\varepsilon}}} \leq e^{2C\varepsilon^{\alpha'}} \frac{e^{\frac{\langle x|y_i \rangle - \psi_i^0}{\varepsilon}}}{\sum_j e^{\frac{\langle x|y_j \rangle - \psi_j^0}{\varepsilon}}} \leq e^{2C\varepsilon^{\alpha'}} e^{-c_x/\varepsilon} \lesssim e^{-c_x/\varepsilon}.$$

If $\gamma_{x,i}^0 = 1$, then

$$|\gamma_{x,i}^\varepsilon - 1| = \frac{\sum_{j \neq i} e^{\frac{\langle x|y_j \rangle - \psi_j^\varepsilon}{\varepsilon}}}{\sum_j e^{\frac{\langle x|y_j \rangle - \psi_j^\varepsilon}{\varepsilon}}} \leq e^{2C\varepsilon^{\alpha'}} \frac{\sum_{j \neq i} e^{\frac{\langle x|y_j \rangle - \psi_j^0}{\varepsilon}}}{\sum_j e^{\frac{\langle x|y_j \rangle - \psi_j^0}{\varepsilon}}} \leq e^{2C\varepsilon^{\alpha'}} \sum_{j \neq i} e^{\frac{\langle x|y_j \rangle - \psi_j^0 - (\psi^0)^*(x)}{\varepsilon}}.$$

We thus get, using that $(\psi^0)^*(x) = \langle x|y_i \rangle - \psi_i^0$ and that $c_x = \min_{j \neq i} (\langle x|y_i \rangle - \psi_i^0) - (\langle x|y_j \rangle - \psi_j^0) > 0$,

$$|\gamma_{x,i}^\varepsilon - \gamma_{x,i}^0| = |\gamma_{x,i}^\varepsilon - 1| \leq e^{2C\varepsilon^{\alpha'}} (N-1)e^{-c_x/\varepsilon} \lesssim e^{-c_x/\varepsilon}.$$

This proves that for ρ -a.e. $x \in \mathcal{X}$ and all $i \in \{1, \dots, N\}$, $|\gamma_{x,i}^\varepsilon - \gamma_{x,i}^0| \lesssim e^{-c_x/\varepsilon}$. \square

7.6 Proof of Theorem 7.8 and its tightness

7.6.1 Expansion of the difference of costs (Theorem 7.8)

The proof of Theorem 7.8 follows very closely the proof of Theorem 1.1 in (Altschuler et al., 2022) and we thus make numerous mentions of results from this paper.

Proof of Theorem 7.8. Let $\varepsilon \in (0, 1]$ and recall that $W_{2,\varepsilon}^2(\rho, \mu)$ is computed using the solution of the regularized maximum correlation problem (P'_ε) with regularization parameter $\frac{\varepsilon}{2}$:

$$W_{2,\varepsilon}^2(\rho, \mu) = \mathbb{E}_{(x,y) \sim \gamma^{\varepsilon/2}} \|x - y\|^2.$$

Now notice that as in Lemma 5.2 in (Altschuler et al., 2022) we can write by strong duality

$$\begin{aligned} \mathbb{E}_{\gamma^0} \langle x|y \rangle &= \mathbb{E}_{x \sim \rho} (\boldsymbol{\psi}^0)^*(x) + \mathbb{E}_{y \sim \mu} \boldsymbol{\psi}^0(y) \\ &= \mathbb{E}_{(x,y) \sim \gamma^{\varepsilon/2}} ((\boldsymbol{\psi}^0)^*(x) + \boldsymbol{\psi}^0(y)). \end{aligned}$$

Therefore the difference of costs reads

$$\begin{aligned} W_{2,\varepsilon}^2(\rho, \mu) - W_2^2(\rho, \mu) &= \mathbb{E}_{\gamma^{\varepsilon/2}} \|x - y\|^2 - \mathbb{E}_{\gamma^0} \|x - y\|^2 \\ &= 2 \left(\mathbb{E}_{\gamma^0} \langle x|y \rangle - \mathbb{E}_{\gamma^{\varepsilon/2}} \langle x|y \rangle \right) \\ &= 2 \mathbb{E}_{(x,y) \sim \gamma^{\varepsilon/2}} ((\boldsymbol{\psi}^0)^*(x) + \boldsymbol{\psi}^0(y) - \langle x|y \rangle) \\ &= 2 \sum_{i,j} \int_{\text{Lag}_i(\boldsymbol{\psi}^0)} ((\boldsymbol{\psi}^0)^*(x) + \boldsymbol{\psi}_j^0 - \langle x|y_j \rangle) \gamma_{x,j}^{\varepsilon/2} d\rho(x) \\ &= \sum_{i,j} \int_{\text{Lag}_i(\boldsymbol{\psi}^0)} \Delta_{ij}(x) \gamma_{x,j}^{\varepsilon/2} d\rho(x), \end{aligned}$$

where we denoted for $x \in \text{Lag}_i(\boldsymbol{\psi}^0)$,

$$\Delta_{ij}(x) = 2 \left((\boldsymbol{\psi}^0)^*(x) + \boldsymbol{\psi}_j^0 - \langle x|y_j \rangle \right) = 2 \left(\langle x|y_i - y_j \rangle - \boldsymbol{\psi}_i^0 + \boldsymbol{\psi}_j^0 \right) \geq 0.$$

Using Theorem 7.2 and its Corollary 7.5 we know that there exists $C > 0$ depending on $\mathcal{X}, \rho, \mathcal{Y}, \mu$ such that for $\alpha' \in (0, \alpha)$,

$$\left\| \boldsymbol{\psi}^{\varepsilon/2} - \boldsymbol{\psi}^0 \right\|_\infty \leq C \left(\frac{\varepsilon}{2} \right)^{1+\alpha'}.$$

From this bound, using that $\gamma_{x,j}^{\varepsilon/2} = \frac{\exp\left(\frac{\langle x|y_j \rangle - \boldsymbol{\psi}_j^{\varepsilon/2}}{\varepsilon/2}\right)}{\sum_\ell \exp\left(\frac{\langle x|y_\ell \rangle - \boldsymbol{\psi}_\ell^{\varepsilon/2}}{\varepsilon/2}\right)}$ we deduce the bounds

$$e^{-2C(\varepsilon/2)^{\alpha'}} \frac{e^{-\Delta_{ij}(x)/\varepsilon}}{\sum_\ell e^{-\Delta_{i\ell}(x)/\varepsilon}} \leq \gamma_{x,j}^{\varepsilon/2} \leq e^{2C(\varepsilon/2)^{\alpha'}} \frac{e^{-\Delta_{ij}(x)/\varepsilon}}{\sum_\ell e^{-\Delta_{i\ell}(x)/\varepsilon}}.$$

Hence we deduce the following control:

$$\begin{aligned} &\left| W_{2,\varepsilon}^2(\rho, \mu) - W_2^2(\rho, \mu) - \sum_{i,j} \int_{\text{Lag}_i(\boldsymbol{\psi}^0)} \Delta_{ij}(x) \frac{e^{-\Delta_{ij}(x)/\varepsilon}}{\sum_\ell e^{-\Delta_{i\ell}(x)/\varepsilon}} d\rho(x) \right| \\ &\lesssim \varepsilon^{\alpha'} \sum_{i,j} \int_{\text{Lag}_i(\boldsymbol{\psi}^0)} \Delta_{ij}(x) \frac{e^{-\Delta_{ij}(x)/\varepsilon}}{\sum_\ell e^{-\Delta_{i\ell}(x)/\varepsilon}} d\rho(x). \end{aligned}$$

Lemma 6.4 found in (Altschuler et al., 2022) then asserts that

$$\sum_{i,j} \int_{\text{Lag}_i(\psi^0)} \Delta_{ij}(x) \frac{e^{-\Delta_{ij}(x)/\varepsilon}}{\sum_{\ell} e^{-\Delta_{i\ell}(x)/\varepsilon}} d\rho(x) = \varepsilon^2 \frac{\pi^2}{12} \sum_{i < j} \frac{w_{ij}}{\|y_i - y_j\|} + o(\varepsilon^2).$$

Directly injecting this into the last control would lead to the exact same asymptotic result as the one given in Theorem 1.1 in (Altschuler et al., 2022). We refine this last asymptotic development to a non-asymptotic one by leveraging the fact that we assumed the source ρ to be α -Hölder continuous.

For any $i, j \in \{1, \dots, N\}$, denote $I_{ij} = \int_{\text{Lag}_i(\psi^0)} \Delta_{ij}(x) \frac{e^{-\Delta_{ij}(x)/\varepsilon}}{\sum_{\ell} e^{-\Delta_{i\ell}(x)/\varepsilon}} d\rho(x)$. Then notice that for $i \neq j$,

$$I_{ij} \leq \int_{\text{Lag}_i(\psi^0)} \Delta_{ij}(x) \frac{e^{-\Delta_{ij}(x)/\varepsilon}}{1 + e^{-\Delta_{ij}(x)/\varepsilon}} d\rho(x) \quad (7.27)$$

Similarly to (Altschuler et al., 2022), introduce for $a > 0$

$$S_{ij}(a) = \{x \in \text{Lag}_i(\psi^0) \mid \forall k \neq i, j, \Delta_{ik}(x) \geq a\}.$$

Notice that $S_{ij}(0) = \text{Lag}_i(\psi^0)$, and that for some $a > 0$:

$$\begin{aligned} I_{ij} &\geq \int_{S_{ij}(a)} \Delta_{ij}(x) \frac{e^{-\Delta_{ij}(x)/\varepsilon}}{1 + (N-2)e^{-a/\varepsilon} + e^{-\Delta_{ij}(x)/\varepsilon}} d\rho(x) \\ &\geq \int_{S_{ij}(a)} \Delta_{ij}(x) \frac{e^{-\Delta_{ij}(x)/\varepsilon}}{c(a) + e^{-\Delta_{ij}(x)/\varepsilon}} d\rho(x), \end{aligned} \quad (7.28)$$

for $c(a) = 1 + (N-2)e^{-a/\varepsilon} > 1$. The quantity I_{ij} is thus bounded by integrals of the form

$$\int_{S_{ij}(a)} \Delta_{ij}(x) \frac{e^{-\Delta_{ij}(x)/\varepsilon}}{c + e^{-\Delta_{ij}(x)/\varepsilon}} d\rho(x)$$

for some $a \geq 0$ and $c \geq 1$. Let's find a non-asymptotic control of such integrals in terms of ε .

Recall that for $x \in \text{Lag}_i(\psi^0)$, $\Delta_{ij}(x) = 2 \left(\langle x | y_i - y_j \rangle - \psi_i^0 + \psi_j^0 \right)$. The coarea formula then ensures:

$$\int_{S_{ij}(a)} \Delta_{ij}(x) \frac{e^{-\Delta_{ij}(x)/\varepsilon}}{c + e^{-\Delta_{ij}(x)/\varepsilon}} d\rho(x) = \frac{1}{2 \|y_i - y_j\|} \int_0^\infty t \frac{e^{-t/\varepsilon}}{c + e^{-t/\varepsilon}} h_{ij}(t; a) dt \quad (7.29)$$

where we denoted $h_{ij}(t; a) = \int_{S_{ij}(a) \cap (\Delta_{ij})^{-1}(t)} \rho(x) d\mathcal{H}^{d-1}(x)$ similarly to (Altschuler et al., 2022) (one can already notice that $h_{ij}(0; 0) = w_{ij}$). Notice then from Lemma 6.2 in (Altschuler et al., 2022) that

$$\begin{aligned} &\left| \int_0^\infty t \frac{e^{-t/\varepsilon}}{c + e^{-t/\varepsilon}} h_{ij}(t; a) dt - \varepsilon^2 h_{ij}(0; a) (-\text{Li}_2(-1/c)) \right| \\ &= \varepsilon^2 \left| \int_0^\infty u \frac{e^{-u}}{c + e^{-u}} (h_{ij}(\varepsilon u; a) - h_{ij}(0; a)) du \right|, \end{aligned}$$

where Li_2 denotes the dilogarithm function. We now focus on the difference $h_{ij}(\varepsilon u; a) - h_{ij}(0; a)$:

$$h_{ij}(\varepsilon u; a) - h_{ij}(0; a) = \int_{S_{ij}(a) \cap (\Delta_{ij})^{-1}(\varepsilon u)} \rho(x) d\mathcal{H}^{d-1}(x) - \int_{S_{ij}(a) \cap (\Delta_{ij})^{-1}(0)} \rho(x) d\mathcal{H}^{d-1}(x)$$

One can notice that there exists a set $R_{ij}(\varepsilon u)$, that is a subset of an hyperplane, and whose $(d-1)$ -area is (at most) linear in εu , such that either

$$S_{ij}(a) \cap (\Delta_{ij})^{-1}(\varepsilon u) = (S_{ij}(a) \cap (\Delta_{ij})^{-1}(0) + (\varepsilon u)n_{ij}) \cup R_{ij}(u),$$

or

$$S_{ij}(a) \cap (\Delta_{ij})^{-1}(\varepsilon u) = (S_{ij}(a) \cap (\Delta_{ij})^{-1}(0) + (\varepsilon u)n_{ij}) \setminus R_{ij}(u),$$

where $n_{ij} = \frac{y_i - y_j}{\|y_i - y_j\|}$. Hence we have:

$$\begin{aligned} |h_{ij}(\varepsilon u; a) - h_{ij}(0; a)| &\leq \int_{S_{ij}(a) \cap (\Delta_{ij})^{-1}(0)} |\rho(x + (\varepsilon u)n_{ij}) - \rho(x)| d\mathcal{H}^{d-1}(x) \\ &\quad + \int_{R_{ij}(\varepsilon u)} \rho(x) d\mathcal{H}^{d-1}(x) \end{aligned}$$

Recalling that ρ is α -Hölder continuous, we have $|\rho(x + (\varepsilon u)n_{ij}) - \rho(x)| \lesssim \varepsilon^\alpha u^\alpha$. Hence we deduce

$$|h_{ij}(\varepsilon u; a) - h_{ij}(0; a)| \lesssim \varepsilon^\alpha u^\alpha + \varepsilon u.$$

We thus have shown that

$$\left| \int_0^\infty t \frac{e^{-t/\varepsilon}}{c + e^{-t/\varepsilon}} h_{ij}(t; a) dt - \varepsilon^2 h_{ij}(0; a) (-\text{Li}_2(-1/c)) \right| \lesssim \varepsilon^2 \int_0^\infty u \frac{e^{-u}}{c + e^{-u}} (\varepsilon^\alpha u^\alpha + \varepsilon u) du \lesssim \varepsilon^{2+\alpha},$$

where we used that $c \geq 1$ and $\varepsilon \leq 1$.

We finally bound the distance between $h_{ij}(0; a) (-\text{Li}_2(-1/c))$ and $h_{ij}(0; 0) (-\text{Li}_2(-1))$. We have the following inequality:

$$\begin{aligned} |h_{ij}(0; a) (-\text{Li}_2(-1/c)) - h_{ij}(0; 0) (-\text{Li}_2(-1))| &\leq |-\text{Li}_2(-1)| |h_{ij}(0; a) - h_{ij}(0; 0)| \\ &\quad + |h_{ij}(0; a)| |-\text{Li}_2(-1/c) - (-\text{Li}_2(-1))| \end{aligned}$$

The quantity $|h_{ij}(0; a) - h_{ij}(0; 0)|$ obviously scales linearly with a . Then one can notice that on $[1, c]$, the function $t \mapsto -\text{Li}_2(-1/t)$ is $(-\text{Li}_2(-1))$ -Lipschitz. These facts ensure the following control:

$$|h_{ij}(0; a) (-\text{Li}_2(-1/c)) - h_{ij}(0; 0) (-\text{Li}_2(-1))| \lesssim a + |c - 1|.$$

This allows to write

$$\left| \int_0^\infty t \frac{e^{-t/\varepsilon}}{c + e^{-t/\varepsilon}} h_{ij}(t; a) dt - \varepsilon^2 h_{ij}(0; 0) (-\text{Li}_2(-1)) \right| \lesssim \varepsilon^{2+\alpha} + \varepsilon^2 (a + |c - 1|).$$

Thus, setting $a = \varepsilon^\alpha$, we get $c(a) = 1 + (N-2)e^{-1/\varepsilon^{1-\alpha}}$ and

$$\left| \int_0^\infty t \frac{e^{-t/\varepsilon}}{c(a) + e^{-t/\varepsilon}} h_{ij}(t; a) dt - \varepsilon^2 h_{ij}(0; 0) (-\text{Li}_2(-1)) \right| \lesssim \varepsilon^{2+\alpha}.$$

Thus, using the controls (7.27), (7.28) on I_{ij} together with the formula of (7.29) and the last control, we have:

$$\left| I_{ij} - \varepsilon^2 \frac{h_{ij}(0; 0)}{2 \|y_i - y_j\|} (-\text{Li}_2(-1)) \right| = \left| I_{ij} - \varepsilon^2 \frac{w_{ij}}{\|y_i - y_j\|} \frac{\pi^2}{24} \right| \lesssim \varepsilon^{2+\alpha}.$$

Eventually, recalling the bound

$$\left| W_{2,\varepsilon}^2(\rho, \mu) - W_2^2(\rho, \mu) - \sum_{i,j} I_{ij} \right| \lesssim \varepsilon^{\alpha'} \sum_{i,j} I_{ij},$$

we obtain the wanted control for $\varepsilon \leq 1$:

$$\left| W_{2,\varepsilon}^2(\rho, \mu) - W_2^2(\rho, \mu) - \varepsilon^2 \frac{\pi^2}{12} \sum_{i < j} \frac{w_{ij}}{\|y_i - y_j\|} \right| \lesssim \varepsilon^{2+\alpha}. \quad \square$$

7.6.2 Tightness of Theorem 7.8

We now show that Theorem 7.8 is tight on a simple one-dimensional example.

Theorem 7.14. *In Theorem 7.8, there exists ρ, μ such that for $\varepsilon \leq 1$,*

$$\left| W_{2,\varepsilon}^2(\rho, \mu) - W_2^2(\rho, \mu) - \varepsilon^2 \frac{\pi^2}{12} \sum_{i < j} \frac{w_{ij}}{\|y_i - y_j\|} \right| \gtrsim \varepsilon^{2+\alpha}.$$

Proof. Once again, we rely on results from (Altschuler et al., 2022) (Section 3), where the following formula for the difference of costs for the transport between a continuous symmetric density ρ supported on $[-1, 1]$ and the target $\mu = \frac{1}{2}(\delta_{\{-1\}} + \delta_{\{+1\}})$ is given:

$$W_{2,\varepsilon}^2(\rho, \mu) - W_2^2(\rho, \mu) = 8 \int_0^1 \frac{x}{1 + e^{4x/\varepsilon}} \rho(x) dx.$$

We consider, for $\alpha \in (0, 1]$ the following α -Hölder density for the source:

$$\rho(x) = \frac{1 + \alpha}{2\alpha} (1 - |x|^\alpha) \chi_{[-1,1]}.$$

We can then derive the difference between the suboptimality and its asymptote for $\varepsilon \leq 1$:

$$\begin{aligned}
\left| W_{2,\varepsilon}^2(\rho, \mu) - W_2^2(\rho, \mu) - \frac{\pi^2 \rho(0)}{24} \varepsilon^2 \right| &= \left| 8 \int_0^1 \frac{x}{1 + e^{4x/\varepsilon}} \rho(x) dx - \rho(0) \frac{\varepsilon^2}{2} (-\text{Li}_2(-1)) \right| \\
&\geq \left| 8 \int_0^1 \frac{x}{1 + e^{4x/\varepsilon}} (\rho(x) - \rho(0)) dx \right| \\
&\quad - \left| 8 \rho(0) \int_0^1 \frac{x}{1 + e^{4x/\varepsilon}} dx - \rho(0) \frac{\varepsilon^2}{2} (-\text{Li}_2(-1)) \right| \\
&= \left| 8 \frac{1 + \alpha}{2\alpha} \int_0^1 \frac{x^{1+\alpha}}{1 + e^{4x/\varepsilon}} dx \right| - \left| \rho(0) \frac{\varepsilon^2}{2} \int_{4/\varepsilon}^\infty \frac{te^{-t}}{1 + e^{-t}} dt \right| \\
&\geq \frac{4(1 + \alpha)}{\alpha} \left(\frac{\varepsilon}{4} \right)^{2+\alpha} \int_0^{4/\varepsilon} \frac{t^{1+\alpha}}{1 + e^t} dt - 4\rho(0)\varepsilon e^{-4/\varepsilon} \\
&\geq \frac{4(1 + \alpha)}{\alpha} \left(\frac{\varepsilon}{4} \right)^{2+\alpha} \int_0^4 \frac{t^{1+\alpha}}{1 + e^t} dt - 4\rho(0)\varepsilon e^{-4/\varepsilon} \\
&\geq \frac{2(1 + \alpha)}{\alpha} \left(\frac{2}{(1 + e^4)(2 + \alpha)} \varepsilon^{2+\alpha} - \varepsilon e^{-4/\varepsilon} \right).
\end{aligned}$$

Thus for ε small enough we get:

$$\left| W_{2,\varepsilon}^2(\rho, \mu) - W_2^2(\rho, \mu) - \frac{\pi^2 \rho(0)}{24} \varepsilon^2 \right| \gtrsim \varepsilon^{2+\alpha}. \quad \square$$

7.7 Numerical illustrations

The code that generated the illustrations of this section is available at <https://github.com/alex-delalande/potentials-entropic-sd-ot>.

7.7.1 Difference of Costs

Figure 7.2 gives an illustration of Theorem 7.8 for a target $\mu = \frac{1}{2}(\delta_{-1} + \delta_1)$ and for four different source measures: 1. Lebesgue: $\rho(x) \propto \chi_{[-1,1]}(x)$; 2. Rescaled Gaussian: $\rho(x) \propto e^{-x^2/2\sigma^2} \chi_{[-1,1]}(x)$; 3. Rescaled Laplace: $\rho(x) \propto e^{-|x|} \chi_{[-1,1]}(x)$; 4. $\frac{1}{2}$ -Hölder density: $\rho(x) \propto (1 - |x|^{1/2}) \chi_{[-1,1]}(x)$. For all these sources, we plot the absolute value of the difference of costs minus its asymptote as functions of ε and compare their rate of convergence to zero as ε goes to zero to the rate predicted in Theorem 7.8. The difference of costs is computed using the following formula given in Section 3 of (Altschuler et al., 2022):

$$W_{2,\varepsilon}^2(\rho, \mu) - W_2^2(\rho, \mu) = 8 \int_0^1 \frac{x}{1 + e^{4x/\varepsilon}} \rho(x) dx.$$

Note that in these examples, $\varepsilon \mapsto \psi^\varepsilon$ is constant because of the symmetry of the problems. One can notice that for the cases of a Lebesgue or rescaled Gaussian source, the convergence of the difference of costs to its asymptote seems faster than the guaranteed ε^3 of Theorem 7.8. However, one can observe that Theorem 7.8 seems to give tight rates of convergence in the cases of a rescaled Laplace source or a $\frac{1}{2}$ -Hölder source.

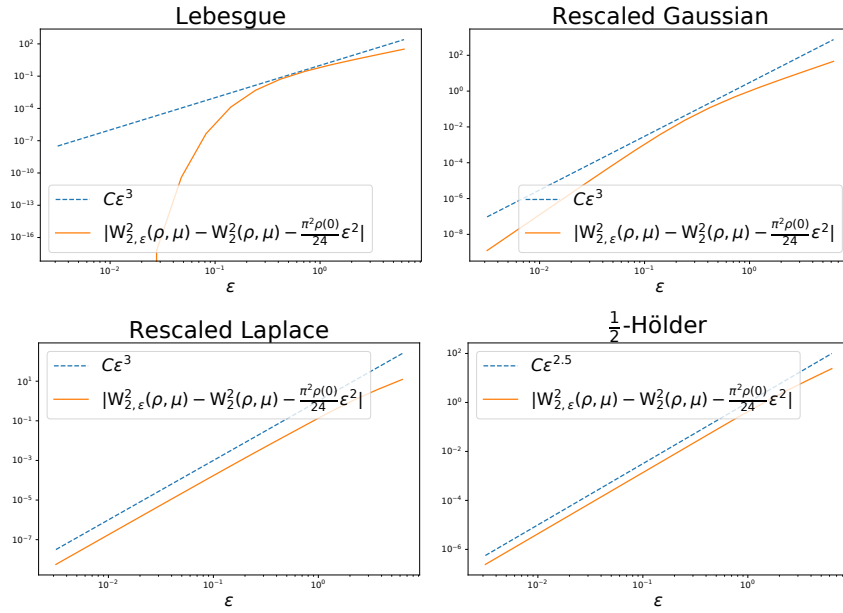


Figure 7.2: Convergence of the difference of costs to its asymptote for the four different sources. The target is $\mu = \frac{1}{2}(\delta_{-1} + \delta_1)$.

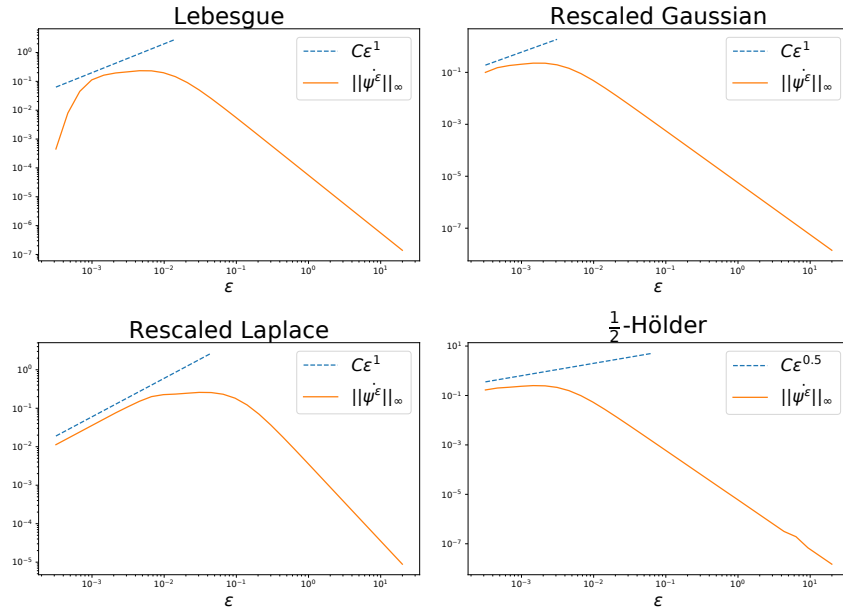


Figure 7.3: Behavior of $\epsilon \mapsto \|\psi^\epsilon\|_\infty$ for the 4 different sources of Section 7.7.1 and $\mu = \frac{1}{5} \sum_{i=1}^5 \delta_{y_i}$ with $(y_i)_{i=1,\dots,5}$ randomly chosen.




Figure 7.4: Convergence of ψ^ε to ψ^0 for the four same examples as in Figure 7.3.

7.7.2 Behavior of $\varepsilon \mapsto \psi^\varepsilon$

Figures 7.3 and 7.4 give an illustration of Theorem 7.2 and its Corollary 7.5 respectively. We consider the same one dimensional sources as in the preceding section (up to restriction of their support to limit numerical errors). We consider a target μ with 5 support points randomly chosen in the support of the source. We compute ψ^ε using the L-BFGS-B quasi-Newton method from SciPy (Virtanen et al., 2020), where all integrals (appearing for instance in gradient computations) are also approximated using this package. Figure 7.3 represents the behavior of $\|\dot{\psi}^\varepsilon\|_\infty$ with respect to ε and compares the empirical results to the theoretical rates of Theorem 7.2. The derivative $\dot{\psi}^\varepsilon$ is computed as the only solution in $(\mathbf{1}_N)^\perp$ to the linear system induced by the ODE (7.7). Note that in Figure 7.3, the *long-time* bound $\|\dot{\psi}^\varepsilon\|_\infty \lesssim 1$ for $\varepsilon \geq 1$ seems to be loose, but this is specific to the setting where the target is uniform and this bound seems tight in general¹. The *short-time* case $\varepsilon < 1$ however yields in the case of the Laplace and $\frac{1}{2}$ -Hölder sources practical rates that seem to match the theoretical rate $\|\dot{\psi}^\varepsilon\|_\infty \lesssim \varepsilon^\alpha$. Figure 7.4 gives an illustration of the convergence of ψ^ε to ψ^0 . One can observe that the Lebesgue and Gaussian sources seem to enjoy faster rates of convergence than our theoretical rates. However, the Laplace and $\frac{1}{2}$ -Hölder sources seem to yield potentials ψ^ε that converge to ψ^0 as fast as predicted in Corollary 7.5.

¹Further experiments with a non-uniform target led to match empirically the long-time bounds of Theorem 7.2, see the GitHub repository.

Part III

Numerical applications: the Linearized Optimal Transport framework

Linearized optimal transport and applications

Abstract

In the Linearized Optimal Transport framework of (Wang et al., 2013), probability measures from $(\mathcal{P}_2(\mathbb{R}^d), W_2)$ are embedded into a Hilbert space $L^2(\rho, \mathbb{R}^d)$ where ρ is a given reference probability measure. In Chapter 5, we studied theoretically this embedding and showed under regularity assumption on ρ that it does not distort too much the 2-Wasserstein distance in general. In this chapter, we gather numerical illustrations showcasing how the LOT embedding performs in practice. We first compare its distance approximation performances to what we predicted in Chapter 5. Then, we give examples of how this embedding may be used to extend some classical Euclidean data analysis methods to measure data, consistently with the Wasserstein geometry. Namely, we tackle K -means and dictionary learning problems in the Wasserstein space.

8.1 Introduction

Numerous data analysis problems involve the comparison of point clouds, i.e. sets of points that lie in a metric space and for which the spatial distribution is of interest. Seeing the point clouds as discrete probability measures in a metric space, it is natural to compare them using Wasserstein distances. These distances have indeed been successfully used in a variety of applications in machine learning (Canas and Rosasco, 2012; Arjovsky et al., 2017; Gordaliza et al., 2019; Genevay et al., 2018; Flamary et al., 2018; Alaux et al., 2019) and in statistics (Weed and Berthet, 2019; Cazelles et al., 2018; Bigot et al., 2019a; Ramdas et al., 2015). In the discrete setting, many efficient algorithms have been proposed to compute or approximate the Wasserstein distances, such as Sinkhorn-Knopp and auction algorithms – see (Peyré and Cuturi, 2019) and references therein. However efficient these algorithms are, they still represent high computational costs when dealing with large databases of point clouds. For instance, when there are k point clouds, $\frac{1}{2}k(k-1)$ optimal transport problems must be solved to get the distance matrix of the database. In addition, such methods provide good approximations of Wasserstein distances but they do not allow for the direct use of machine learning algorithms based on the Wasserstein geometry. In (Wang et al., 2013), the Linearized Optimal Transport (LOT) framework

was introduced in order to circumvent both of these problems. In this framework, for a fixed absolutely continuous reference probability measure $\rho \in \mathcal{P}_{2,a.c.}(\mathbb{R}^d)$, an explicit embedding of the metric space $(\mathcal{P}_2(\mathbb{R}^d), W_2)$ into the Hilbert space $L^2(\rho, \mathbb{R}^d)$ is computed. This embedding may be defined as follows:

$$\begin{cases} (\mathcal{P}_2(\mathbb{R}^d), W_2) & \rightarrow L^2(\rho, \mathbb{R}^d), \\ \mu & \mapsto T_\mu, \end{cases}$$

where for any $\mu \in \mathcal{P}_2(\mathbb{R}^d)$, T_μ denotes the optimal transport map from ρ to μ . By Brenier's theorem ((Brenier, 1991), see Theorem 1.12), this embedding is well defined. Moreover, as already mentioned in Section 5.1 of Chapter 5 (see in particular §5.1.1), this embedding is reverse-Lipschitz and continuous. In Chapter 5, we have quantified the continuity of this embedding under restrictions on its domain of definition and under some regularity assumptions on the reference ρ . For instance, Theorem 5.12 asserts that assuming that the reference ρ is supported on a compact convex set \mathcal{X} and that its density is bounded away from zero and infinity, for any compact set $\mathcal{Y} \subset \mathbb{R}^d$, the LOT embedding is bi-Hölder continuous on $(\mathcal{P}_2(\mathcal{Y}), W_2)$:

$$\forall \mu, \nu \in \mathcal{P}_2(\mathcal{Y}), \quad W_2(\mu, \nu) \leq \|T_\mu - T_\nu\|_{L^2(\rho, \mathbb{R}^d)} \lesssim W_2(\mu, \nu)^{1/6},$$

where \lesssim hides a multiplicative constant that depends on d, ρ, \mathcal{X} and \mathcal{Y} . A LOT approach may have a cheaper computational cost than usual optimal transport based approaches used for data analysis on measure datasets. For instance, the computation of the distance matrix of a database of k point clouds in this framework only requires the resolution of k optimal transport problems in order to compute the embeddings, and a number of $\frac{1}{2}k(k-1)$ Hilbertian (Euclidean in practice) distance computations, which are generally cheap to compute. Second, the LOT framework presents the advantage of enabling the use of the classical Hilbertian statistical toolbox on families of probability measures while keeping some features of the Wasserstein geometry.

Outline. As mentioned in Section 5.1, the LOT framework has already found many successful applications (see this section for references). In this chapter, we give complementary illustrations of the LOT embedding. In Section 8.2, we illustrate the theoretical results of Chapter 5 and observe the metric distortion induced by the LOT embedding on some two-dimensional examples. We also mention how the LOT embedding may be used to perform barycenter approximation in the Wasserstein space. Then in Section 8.3, we give two example extensions of classical Hilbertian data analysis methods to probability measures within the LOT framework. These extensions concern K -means and dictionary learning problems in the Wasserstein space.

Setting. Most of the examples in this chapter are two-dimensional. In this case, the source measure ρ is fixed and corresponds to the Lebesgue measure on the the unit square, i.e. $\rho \equiv 1$ on $\mathcal{X} = [0, 1] \times [0, 1]$. The databases of probability measures we consider are only made of discrete probability measures that are taken uniform unless stated otherwise. In this context, the images of the above defined LOT embedding correspond to semi-discrete optimal transport maps that we estimate with a damped Newton's algorithm (Kitagawa et al., 2019). The embeddings T_μ are infinite dimensional objects that are approximated by their block approximation over a uniform block partition of \mathcal{X} : for m a positive integer defining the blocks side $\frac{1}{m}$, the blocks are defined




Figure 8.1: $W_{2,\rho}$ vs. W_2 between point clouds sampled from Gaussian, Mixture of 4 Gaussian and Uniform distributions.

by $\mathcal{X}_{s,t} = \left[\frac{s-1}{m}, \frac{s}{m}\right] \times \left[\frac{t-1}{m}, \frac{t}{m}\right]$ for $s, t \in \{1, \dots, m\}$ and T_μ is approximated by the vector $\mathbf{T}_\mu := (\rho(\mathcal{X}_{s,t})^{-1/2} \int_{\mathcal{X}_{s,t}} T_\mu d\rho)_{s,t \in \{1, \dots, m\}}$ of size $2m^2$. We can notice that the stability results of Chapter 5 on the maps $T_\mu \in L^2(\rho, \mathbb{R}^d)$ can be directly applied to the vectors \mathbf{T}_μ . Indeed these vectors correspond to the projections of T_μ on a subspace of $L^2(\rho, \mathbb{R}^d)$ of piece-wise constant functions on \mathcal{X} : as a projection this mapping is 1-Lipschitz, which allows to write $\|\mathbf{T}_\mu - \mathbf{T}_\nu\|_2 \leq \|T_\mu - T_\nu\|_{L^2(\rho, \mathbb{R}^d)}$.

Remark 8.1. We can note that in dimension d , the approximation \mathbf{T}_μ is of size dm^d : this limits the use of this approximation to small values of d . Lighter representations of the map T_μ could however be considered. In particular one could leverage the fact that, in this semi-discrete context, T_μ is piece-wise constant. Moreover, it is defined properly by a dual potential that can be seen as a vector of size equal to the number of points in the support of μ (see Chapter 2).

8.2 Behavior of the LOT embedding

The code that generated the experiments of this section can be found at https://github.com/alex-delalande/stability_ot_maps_and_linearization_wasserstein_space/.

8.2.1 Distance approximation

We first compare the LOT distance $W_{2,\rho}(\mu, \nu) = \|T_\mu - T_\nu\|_{L^2(\rho, \mathbb{R}^d)}$ against $W_2(\mu, \nu)$ in specific settings to illustrate the Hölder stability results of Chapter 5. We consider three different settings corresponding to three different families of distributions. In each setting, 50 point clouds of 300 points are sampled, each from a random distribution that belongs to the given family, and pairwise W_2 and $W_{2,\rho}$ distances on the 50 point clouds are computed (W_2 distances are computed exactly with the network simplex algorithm implemented in the Python Optimal Transport library (Flamary et al., 2021)). The distances $\|T_\mu - T_\nu\|_{L^2(\rho, \mathbb{R}^d)}$ are approximated with $\|\mathbf{T}_\mu - \mathbf{T}_\nu\|_2$ with $m = 500$. The three families of distributions we consider are all truncated to belong to $\mathcal{P}([0, 1]^2)$ and are built from Gaussian distributions, mixtures of four Gaussian distributions and the uniform distribution. Note that for each point cloud sampling in the two first settings the parameters of the sampled distribution are selected randomly. We report in Figure 8.1 the comparisons between $W_{2,\rho}$ and W_2 . We observe that, as expected, $W_{2,\rho}$ is always greater than W_2 but does not distort too much the 2-Wasserstein distance. These figures do not allow however to investigate whether or not the Hölder exponent of $\frac{1}{6}$ of Theorem 5.12 is optimal.




Figure 8.2: (Left) Target measures, push-forwards of the maps $T_k = \nabla\phi_k$ where $\phi_{\text{Disk}}(x, y) := 0.25(x + y) + 0.07(|x + y|^{3/2} + |x - y|^{3/2})$, $\phi_{\text{Cross}}(x, y) := 0.5(x + y) + 0.04 \max(4(x + y - 1)^2 + 0.5(2x - 1)^2 + 0.5(2y - 1)^2, 4(x - y)^2 + 0.5(2x - 1)^2 + 0.5(2y - 1)^2)$ and $\phi_{\text{Square}}(x, y) := 0.5(x^2 + y^2)$ (Right) Sampling distance $\|T_\mu - T_{\mu_N}\|_{L^2(\rho, \mathbb{R}^d)}$.




Figure 8.3: From left to right: densities of the sampled Gaussian, Mixture of 4 Gaussians and Uniform distributions and sampling distance $\|T_\mu - T_{\mu_N}\|_{L^2(\rho, \mathbb{R}^d)}$ as a function of N

8.2.2 Statistical behavior

In practice, a measure of interest $\mu \in \mathcal{P}(\mathcal{X})$ (where $\mathcal{X} = [0, 1]^2$) may be unknown and one may only have access to samples $(x_i)_{i=1, \dots, N}$ from this distribution. Denoting $\mu_N = \frac{1}{N} \sum_{i=1}^N \delta_{x_i}$ the corresponding empirical distribution, it is well-known that μ_N converges weakly almost surely to μ as N goes to infinity (Varadarajan, 1958). In dimension $d = 2$, Theorem 1 of (Fournier and Guillin, 2015) indicates that the rate of this convergence in expected Wasserstein distance is at least in $N^{-1/2}$: there exists a constant $C > 0$ such that

$$\mathbb{E}W_2^2(\mu, \mu_N) \leq CN^{-1/2}.$$

(Note that this is a worst case bound: for instance, when $\mu = \rho = \lambda_{[0, 1]^2}$, this bound is known to improve asymptotically to $\frac{\log N}{N}$ (Ambrosio et al., 2019)). From Theorem 5.12, the empirical transport map T_{μ_N} should thus approach T_μ in $L^2(\rho, \mathbb{R}^d)$ distance at a rate at least $(N^{-1/4})^{1/6}$. We observe this by plotting the quantity $\|T_\mu - T_{\mu_N}\|_{L^2(\rho, \mathbb{R}^d)}$ as a function of N in again three different settings, where now the ground truth maps T_μ are prescribed. The three maps are chosen as gradients of convex functions and transport the unit square to measures resembling a disk, a cross and a square (Figure 8.2). For the different values of N the experiments are repeated 25 times and the standard deviations define the shaded areas surrounding the curves. In the right hand side log-plot of Figure 8.2, we observe slopes of about -0.4 , which corresponds to faster rates than the expected (and worst case) rate $-\frac{1}{24}$. In a more statistical context, we observe in Figure 8.3 the same quantities when the target measures are a Gaussian, a mixture of 4 Gaussians and the uniform distribution, all truncated to \mathcal{X} . Since the ground truth maps T_μ are unknown in these cases, we approximate them with a map T_{μ_M} for $M = 10000$. Again, we observe slopes of about -0.4 .




Figure 8.4: LOT barycenters of four point clouds. Weights $(\lambda_s)_s$ are bilinear with respect to the corners of the square.

8.2.3 Barycenter approximation

Computing means and barycenters is often necessary in many data analysis tasks. For point cloud or more generally measure data, we have seen in Chapter 6 that the Wasserstein distance is a natural choice to define such barycenters. We recall that for $(\mu_s)_{s=1,\dots,S}$ a set of S discrete probability measures (corresponding to S point clouds), the barycenter of $(\mu_s)_{s=1,\dots,S}$ with non-negative weights $(\lambda_s)_{s=1,\dots,S}$ summing to one is the solution of the following minimization problem:

$$\min_{\mu} \sum_{s=1}^S \lambda_s W_2^2(\mu, \mu_s).$$

This problem does not have an explicit solution and its solution needs to be computed every time the weights are changed. Using transport maps from a reference measure ρ , it is natural to consider

$$\mu = \left(\sum_{s=1}^S \lambda_s T_{\mu_s} \right)_{\#} \rho$$

as the barycenter of the $(\mu_s)_{s=1,\dots,S}$, and one can indeed check that this μ minimizes

$$\sum_s \lambda_s \|T_{\mu} - T_{\mu_s}\|_{L^2(\rho, \mathbb{R}^d)}^2.$$

We illustrate this idea with the computation of barycenters of four point clouds in Figure 8.4. Again, operations are performed on the vectorized transport maps \mathbf{T}_{μ} with $m = 200$. These barycenters are in general not equal to their Wasserstein counterparts but they seem to retain the geometric information contained in the point clouds.

8.3 Applications to clustering, interpolation and classification in the Wasserstein space

In this final section, we give two examples of Euclidean data analysis methods that can easily be extended to probability measures within the LOT framework.

8.3.1 Wasserstein K -means

We give here an illustrative application of the proven bi-Hölder stability results of Chapter 5 to a classical data analysis task of clustering in the space of probability measures endowed with the Wasserstein distance.

Several applications involve decomposing a finite family of probability measures into clusters, using the p -Wasserstein distances as a guiding geometry. These applications include classification of images (Fukunaga and Kasai, 2021; Ye et al., 2017; Ho et al., 2017), of tomographic projections (Rao et al., 2020), of time-series (Yang et al., 2018; Ho et al., 2017), of textual documents (Ye et al., 2017) or of cloud regime histograms (Staib and Jegelka, 2017). In order to keep the computational cost reasonable, many of these works involve approximations that are not analyzed theoretically in terms of clustering quality. The objective of this section is to establish a (simple) guarantee on the clustering quality that one can obtain in the Linearized Optimal Transport framework, relying on Theorem 5.14.

K -means clustering in metric spaces. In a general metric space (M, d_M) , the K -Means clustering of N data points $x_1, \dots, x_N \in M$ consists in finding a set of K centroids (also called *means*) $\mathbf{c} = (c_1, \dots, c_K) \in M^K$ minimizing the *distortion*

$$R_M(\mathbf{c}) = \frac{1}{N} \sum_{i=1}^N \min_{j=1, \dots, K} d_M(x_i, c_j)^2.$$

A set of centroids $\mathbf{c} \in M^K$ induces a clustering of the set $\{x_i\}_{1 \leq i \leq N}$ using Voronoi cells:

$$\forall j \in \{1, \dots, K\}, \quad V_j(\mathbf{c}) = \{x_i \mid \forall \ell \in \{1, \dots, K\}, d_M(x_i, c_j) \leq d_M(x_i, c_\ell)\}.$$

The K -Means clustering problem thus consists in finding a clustering of the data at hand that minimizes the within-cluster variance, obtained by minimizing R_M . A popular algorithm used to find a local minimum of R_M is Lloyd's algorithm (Lloyd, 1982). However, because of its non-convexity, minimizing the objective R_M is in general NP-hard (Mahajan et al., 2012). Performance guarantees in terms of clustering error for K -means algorithms can still be obtained, but only under assumptions of well-clusterability of the data and good initialization of the centroids. A popular initialization method is offered by the K -Means++ (Arthur and Vassilvitskii, 2007) randomized seeding technique, and is described in Algorithm 1.

The (random) output $\hat{\mathbf{c}} = (\hat{c}_1, \dots, \hat{c}_K)$ of Algorithm 1 comes with the following guarantee in expectation in term of the K -Means objective (Nielsen and Sun, 2019, Theorem 4):

$$\mathbb{E}_{\hat{\mathbf{c}}} R_M(\hat{\mathbf{c}}) \leq 16(\log K + 2) \min_{\mathbf{c} \in M^K} R_M(\mathbf{c}). \quad (8.1)$$

Algorithm 1: K -Means++ randomized seeding technique (Arthur and Vassilvitskii, 2007)

Data: Dataset $\{x_1, \dots, x_N\}$

Result: Set of initialized centroids $\hat{\mathbf{c}}$

Select an initial centroid \hat{c}_1 uniformly at random from $\{x_1, \dots, x_N\}$ and set

$\mathcal{C}_1 = \{\hat{c}_1\}$;

for $j \leftarrow 2$ **to** K **do**

Select $\hat{c}_j = x_i \in \{x_1, \dots, x_N\}$ with probability $\frac{d_M(x_i, \mathcal{C}_{j-1})^2}{\sum_{i'=1}^N d_M(x_{i'}, \mathcal{C}_{j-1})^2}$;

Set $\mathcal{C}_j = \mathcal{C}_{j-1} \cup \{\hat{c}_j\}$;

end

where $\mathbb{E}_{\hat{\mathbf{c}}}$ denotes the expectation with respect to the random variable $\hat{\mathbf{c}}$ following the law described in Algorithm 1. Combined with clusterability assumptions, this estimate yields performance guarantees in terms of clustering quality when using $\hat{\mathbf{c}}$ as a set of centroids.

Theorem 8.2. *Let $x_1, \dots, x_N \in M$ partitioned into K' disjoint non-empty clusters $(C_j)_{j=1, \dots, K'}$, and assume that there exists $\epsilon > 0$ and $f \geq 1$ such that for all $\ell, \ell' \in \{1, \dots, K'\}$ with $\ell \neq \ell'$,*

$$\begin{aligned} \forall (x_i, x_j) \in C_\ell \times C_\ell, \quad d_M(x_i, x_j) &\leq \epsilon, \\ \forall (x_i, x_j) \in C_\ell \times C_{\ell'}, \quad d_M(x_i, x_j) &\geq f\epsilon. \end{aligned}$$

Denote $\hat{\mathbf{c}}$ the set of centroids given by the K -Means++ seeding and $\tau_{err}(\hat{\mathbf{c}})$ the proportion of misclassified points using the centroids $\hat{\mathbf{c}}$, i.e.

$$\tau_{err}(\hat{\mathbf{c}}) = \frac{1}{N} \sum_{j=1}^K \min_{j'=1, \dots, K'} |V_j(\hat{\mathbf{c}}) \cap (C_{j'})^c|.$$

Then for $g \geq 1$, we have

$$\mathbb{P}_{\hat{\mathbf{c}}} \left(\tau_{err}(\hat{\mathbf{c}}) \leq \frac{1}{g} \right) \geq 1 - \frac{16(\log K + 2)g}{f^2}$$

Proof. First notice that Markov's inequality and the K -Means++ guarantee (8.1) ensure

$$\mathbb{P}_{\hat{\mathbf{c}}} \left(R_M(\hat{\mathbf{c}}) \geq \frac{f^2}{g} \min_{\mathbf{c} \in M} R_M(\mathbf{c}) \right) \leq \frac{g \mathbb{E}_{\hat{\mathbf{c}}} R_M(\hat{\mathbf{c}})}{f^2 \min_{\mathbf{c} \in M} R_M(\mathbf{c})} \leq \frac{16(\log K + 2)g}{f^2}.$$

Thus with probability greater than $1 - \frac{16(\log K + 2)g}{f^2}$, we have $R_M(\hat{\mathbf{c}}) \leq \frac{f^2}{g} \min_{\mathbf{c} \in M} R_M(\mathbf{c})$. However, the clusterability assumptions imply that

$$\min_{\mathbf{c} \in M} R_M(\mathbf{c}) \leq \epsilon^2 \quad \text{and} \quad R_M(\hat{\mathbf{c}}) \geq \tau_{err}(\hat{\mathbf{c}}) f^2 \epsilon^2.$$

Hence with probability greater than $1 - \frac{16(\log K + 2)g}{f^2}$, $\tau_{err}(\hat{\mathbf{c}}) f^2 \epsilon^2 \leq \frac{f^2}{g} \epsilon^2$. \square




Figure 8.5: (Top) Example images of uppercase A, B, and C from the EMNIST dataset (Cohen et al., 2017). (Bottom) Push-forwards of the 6 centroids.

K -Means with Linearized OT. The performance guarantee given in Theorem 8.2 holds for a general metric space (M, d_M) and it is thus possible to choose $(M, d_M) = (\mathcal{P}_2(\mathbb{R}^d), W_2)$. A dataset $\mu_1, \dots, \mu_N \in (\mathcal{P}_2(\mathbb{R}^d), W_2)$ satisfying the clusterability assumptions of Theorem 8.2 can therefore be clustered with the K -Means++ seeding at the cost of solving $O(KN)$ optimal transport problems in order to compute the seeding probabilities. For settings with relatively large values of N or K , these computations might be prohibitive and it might be interesting to choose to embed $(\mathcal{P}_2(\mathbb{R}^d), W_2)$ into $L^2(\rho; \mathbb{R}^d)$ using the mapping $\mu \mapsto T_\mu$. Indeed, doing this requires solving only N optimal transport problems to compute the embeddings and then computing $O(KN)$ Hilbertian distances. Performance guarantees for K -Means++ seeding on the dataset $T_{\mu_1}, \dots, T_{\mu_N} \in L^2(\rho; \mathbb{R}^d)$ can still be obtained from Theorem 8.2 by assuming stronger clusterability assumptions on the dataset $\mu_1, \dots, \mu_N \in (\mathcal{P}_2(\mathbb{R}^d), W_2)$. Indeed, assume first that μ_1, \dots, μ_N belong to a subset $S \subset \mathcal{P}_2(\mathbb{R}^d)$ that is such that there exists $(C_S, \alpha_S) \in \mathbb{R}_+^* \times (0, 1)$ such that

$$\forall \mu, \nu \in S, \quad W_2(\mu, \nu) \leq \|T_\mu - T_\nu\|_{L^2(\rho; \mathbb{R}^d)} \leq C_S W_2(\mu, \nu)^{\alpha_S}.$$

Theorem 5.14 and Corollary 5.16 give examples of such sets S . Then assume that μ_1, \dots, μ_N is partitioned into K disjoint non-empty clusters $(C_j)_{j=1, \dots, K}$, and that there exists $\epsilon > 0$ and $f \geq 1$ such that for all $\ell, \ell' \in \{1, \dots, K\}$ with $\ell \neq \ell'$,

$$\begin{aligned} \forall (\mu_i, \mu_j) \in C_\ell \times C_\ell, \quad W_2(\mu_i, \mu_j) &\leq \left(\frac{\epsilon}{C_S}\right)^{1/\alpha_S}, \\ \forall (\mu_i, \mu_j) \in C_\ell \times C_{\ell'}, \quad W_2(\mu_i, \mu_j) &\geq f\epsilon. \end{aligned}$$

Then one can apply Theorem 8.2 to the dataset $T_{\mu_1}, \dots, T_{\mu_N} \in L^2(\rho; \mathbb{R}^d)$.

We give an illustration of the described techniques by performing a clustering of images of handwritten letters. We extract from the EMNIST dataset (Cohen et al., 2017) 2,000 images of each of the uppercase letters A, B and C (Figure 8.5) and convert each image into a discrete probability measure of $\mathcal{X} = [0, 1]^2$ by treating the pixels as Dirac masses. Taking ρ to be the Lebesgue measure on \mathcal{X} , we then compute for each measure its optimal transport map from ρ to itself. We eventually perform a clustering with the K -means++ algorithm on the vectorized maps, looking for $K = 6$ clusters. The push-forwards of the 6 centroids are displayed in Figure 8.5 where each *class* seems to be recovered and where *intra-class* variation might be explained from the use or not of italic writing.

8.3.2 Wasserstein dictionary learning

In this final subsection, we propose to approximately solve Wasserstein dictionary learning problems, introduced in (Schmitz et al., 2018), using the LOT embedding and classical Euclidean linear dictionary learning tools. The presentation of dictionary learning methods we make in this section is heavily inspired from (Schmitz et al., 2018). Note that the code for the different experiments of this subsection is available at https://github.com/alex-delalande/linearized_wasserstein_dictionary_learning.

Consider a classical Euclidean data analysis problem on a dataset of N elements living in \mathbb{R}^d . It is well-known that when the dimension d is *large* with respect to the number of samples N , the inference of statistical information from the dataset tends to be difficult. This is often referred to as the *curse of dimensionality* phenomenon, which essentially corresponds to the fact that volumes in \mathbb{R}^d increase exponentially with d , so that any (statistical) reconstruction from samples need in general a number of samples that is exponential in the dimension to achieve a given precision. Many modern days datasets take the form of high dimensional data (think for instance of images with millions of pixels or genomics data with millions of gene features). However in many cases, a dataset can be embedded almost isometrically into a low-dimensional Euclidean space (which corresponds to Johnson–Lindenstrauss lemma) or more generally into a low-dimensional manifold. Consider for instance a three seconds high-resolution video of a swinging pendulum. This video is naturally encoded with a finite set of images (say $N = 3 \times 30$ for a 30 FPS camera) containing each millions of pixels ($d \sim 10^6$). However, because of the periodic nature of the shot, this set of images can easily be thought of as a sequence of points that describe a one dimensional circle in a proper embedding space. Dimensionality reduction techniques aim at finding such low-dimensional embeddings or representations of high-dimensional datasets. Such representations have the advantage of needing less memory storage in a computer than higher-dimensional representations, and they may help alleviate the above mentioned curse of dimensionality for subsequent data analysis. The most famous of these dimensionality reduction techniques comprise principal component analysis (PCA), independent component analysis, manifold learning techniques, autoencoders, or dictionary learning methods. In this section, we focus on the last of these methods and describe how the LOT framework may help extend it straightforwardly to measure data.

Dictionary learning. Many dimensionality reduction methods propose to project a given dataset into a fixed predefined orthogonal basis. This is for instance the case of Fourier or wavelet transforms (Mallat, 2009). The projection basis can be thought of as a *dictionary* comprising *words* in terms of which the dataset may be expressed. Dictionary learning methods go beyond this approach and propose to *learn* the dictionary into which the data is projected, so that its words are well-chosen to represent the dataset at hand. A bit less formally, for a dataset gathered in a matrix $X \in \mathbb{R}^{d \times N}$, one tries to find a dictionary $D \in \mathbb{R}^{d \times k}$ of k *atoms* in \mathbb{R}^d (or words) and a list of *codes* $\Lambda \in \mathbb{R}^{k \times N}$ (or projection coordinates) such that $X \approx D\Lambda$. The similarity between X and $D\Lambda$ may be measured in various ways, but the Frobenius norm constitutes a usual choice of measure of similarity. Making this choice, the dictionary learning problem corresponds to the optimization problem

$$\min_{D \in \mathbb{R}^{d \times k}, \Lambda \in \mathbb{R}^{k \times N}} \|X - D\Lambda\|_F^2. \quad (8.2)$$

When no additional constraint is imposed on D or Λ , solving this problem corresponds to computing a low-rank approximation of X . By the Eckart–Young–Mirsky theorem, this problem admits a closed form solution in terms of the singular value decomposition of X . In practice, one may want to solve the dictionary learning problem with D or Λ satisfying additional constraints. For instance, one may impose that the codes Λ satisfy some sparsity properties in order to select only a few atoms in each representation of the data, in which case one recovers the sparse PCA problem (d’Aspremont et al., 2007). One may otherwise impose that both D and Λ have positive coefficients, in which case one recovers the Nonnegative Matrix Factorization (NMF) problem (Lee and Seung, 1999). From a computational perspective, problem (8.2) is convex separably in D and Λ , so that it is amenable in most cases to (projected) coordinate descent.

Other variants of the dictionary learning problem arise when considering alternative measures of similarity, such as for instance the Kullback-Leibler divergence (Lee and Seung, 1999) or Wasserstein distances (Rolet et al., 2016). Finally, more variants of the dictionary learning problem can be obtained by considering different ways of performing the *reconstruction* from the dictionary and the codes. In the classical dictionary learning problem, the reconstruction $D\Lambda$ is linear in the atoms gathered in D since linear combinations of columns of D are expected to give good reconstruction of columns of X . This reconstruction term $D\Lambda$ may however be replaced by a non-linear reconstruction term $\mathcal{R}(D, \Lambda)$, where \mathcal{R} is a non-linear function of its inputs. This can be particularly relevant for datasets that lie on Riemannian manifolds, where linear interpolations are expected to be replaced by geodesic interpolations (see (Schmitz et al., 2018) for more references on this topic). In the following, we consider the Wasserstein dictionary learning variant introduced in (Schmitz et al., 2018). In this problem, the dataset is made of probability measures equipped with a Wasserstein distance. As such, the atoms of the dictionary are expected to also be probability measures, and reconstruction of input measures from these atoms is not expected to be linear. In the same way, the similarity between a measure and its reconstruction may not be measured in Frobenius norm.

Wasserstein dictionary learning. For a given family of $k \geq 1$ probability measures $D = (d_j)_{1 \leq j \leq k} \in \mathcal{P}_2(\mathbb{R}^d)^k$ and any vector $\lambda \in \Delta_k$, where Δ_k denotes the k -simplex

$$\Delta_k = \{p \in (\mathbb{R}_+)^k \mid \sum_{j=1}^k p_j = 1\},$$

introduce the 2-Wasserstein reconstruction operator

$$\mathcal{R}_{W_2}(D, \lambda) \in \arg \min_{\mu} \sum_{j=1}^k \lambda_j W_2^2(\mu, d_j).$$

This operator associates to the family of measures D a Wasserstein barycenter with weights λ (Agueh and Carlier, 2011). Such barycenter always exists so that \mathcal{R}_{W_2} is well-defined. The Wasserstein dictionary learning problem, introduced in (Schmitz et al., 2018), may be stated as follows. Let $(\mu_i)_{1 \leq i \leq N}$ be a dataset of $N \geq 1$ probability measures supported over \mathbb{R}^d with finite second-order moments. Then solve

$$\min_{D \in \mathcal{P}_2(\mathbb{R}^d)^k, \Lambda \in (\Delta_k)^N} \sum_{i=1}^N W_2^2(\mu_i, \mathcal{R}_{W_2}(D, \Lambda_i)). \quad (8.3)$$

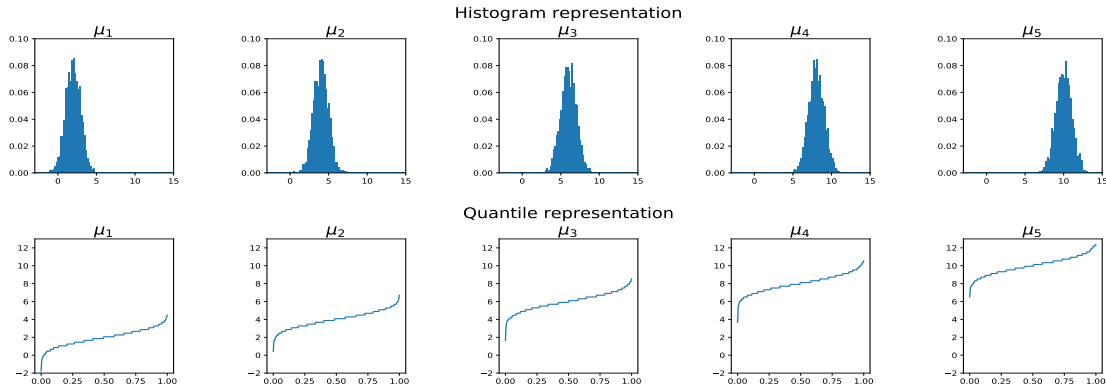


Figure 8.6: (Top) Histograms of each μ_i , computed with 100 bins over the interval $[-4, 16]$. (Bottom) 200-Quantiles of each μ_i .




Figure 8.7: Results for dictionary learning on $(\mu_i)_{1 \leq i \leq 5}$ using their histogram representation. (Left) Atoms. (Right) Original (in green) and reconstructed (in red) histograms.

As with classical dictionary learning, the resolution of this problem can be seen as way to reduce the *dimensionality* of the input dataset $(\mu_i)_{1 \leq i \leq N}$, by representing each μ_i by a vector of codes $\Lambda_i \in \Delta_k$.

A one dimensional example. The relevance of considering the Wasserstein geometry for the dictionary learning problem may easily be illustrated with the following one dimensional example. Let $(\mu_i)_{1 \leq i \leq 5}$ be a dataset of $N = 5$ probability measures. Assume that for each $i \in \{1, \dots, 5\}$, $\mu_i = \frac{1}{m} \sum_{j=1}^m \delta_{x_{i,j}}$ is a discrete probability measure, built from $m = 2000$ points $(x_{i,j})_{1 \leq j \leq m}$ sampled from a Gaussian $\mathcal{N}(2i, 1)$ of mean $2i$ and variance 1. The normalized histograms of each μ_i are gathered in the top row of Figure 8.6.

One may identify each μ_i with its histogram and notice that these five probability measures are approximately translated versions of one another. We wonder if we can recover computationally this structure and look for a dictionary of $k = 2$ atoms such that each μ_i corresponds to a convex combination of these atoms. We first do this in the *histogram domain*, i.e. we represent each μ_i with a vector of dimension $d = 100$ comprising the value of its normalized histogram computed with 100 bins over the interval $[-4, 16]$. We then solve (8.2) with the constraint that each Λ_i belongs to Δ_2 using a projected coordinate descent (see below for more details on the used algorithm). The learned atoms (represented as histograms) as well as the proposed reconstruction of each μ_i are displayed in Figure 8.7.

From Figure 8.7, we observe a general poor reconstruction quality, in particular for the *intermediate* measures μ_2, μ_3, μ_4 . This could be expected: convex combinations of unimodal Gaussian densities do not yield unimodal Gaussian densities but rather




Figure 8.8: Results (in the histogram domain) for dictionary learning on $(\mu_i)_{1 \leq i \leq 5}$ using their quantile representation. (Left) Atoms. (Right) Original (in green) and reconstructed (in red) histograms.

densities of mixtures of Gaussians. The L^2 comparison of the densities thus fails to recover the translations that explain the dataset variations. On the opposite, the 2-Wasserstein distance between a measure and its translation simply corresponds to the norm of this translation. We thus consider the Wasserstein dictionary learning problem of (8.3) and look for $k = 2$ measures that solve this problem. In our one-dimensional setting, the 2-Wasserstein distance between measures μ, ν is known to be equal to the $L^2([0, 1])$ distance between their quantile functions F_μ^{-1}, F_ν^{-1} (i.e. their inverse cumulative distribution function). As such, solving (8.3) for the dataset $(\mu_i)_{1 \leq i \leq 5}$ is approximately equivalent to solving (8.2) for the dataset $(\mathbf{F}_{\mu_i}^{-1})_{1 \leq i \leq 5}$, where each $\mathbf{F}_{\mu_i}^{-1}$ is a vectorized representation of $F_{\mu_i}^{-1}$, computed from a uniform evaluation of $F_{\mu_i}^{-1}$ on 200 points on $[0, 1]$. These uniformly sampled versions of the quantile functions are represented in the bottom row of Figure 8.6. We thus solve (8.3) on the dataset $(\mu_i)_{1 \leq i \leq 5}$ with $k = 2$ and report in Figure 8.8 the learned atoms as well as the reconstructions (all represented in the histogram domain). We observe in this figure that the two learned atoms seem to correspond to the leftmost and rightmost measures of the dataset. The convex combinations of their quantiles, in the quantile domain, naturally yield good reconstruction results. This could be expected from the quantile representations given in the bottom row of Figure 8.6, where quantiles of each μ_i for $i \in \{2, 3, 4\}$ can easily be seen to be recovered as a convex combination of quantiles of μ_1 and μ_5 .

Linearized Wasserstein dictionary learning. The above example of translated unidimensional Gaussians illustrates how the Wasserstein dictionary learning formulation (8.3) of (Schmitz et al., 2018) might be relevant to retrieve the geometric variations that underlie a dataset of probability measures. In the one-dimensional setting, we were able to solve easily (8.3) by leveraging the fact that W_2 is Hilbertian, so that we could reformulate (8.3) as a classical dictionary learning problem of form (8.2). However, for probability measures supported over \mathbb{R}^d with $d \geq 2$, the 2-Wasserstein distance is not Hilbertian and we cannot find equivalence between (8.3) and (8.2) in general. In (Schmitz et al., 2018), the authors propose to solve (8.3) relying on entropic regularization, automatic differentiation and a quasi-Newton method to update simultaneously the atoms and codes. More precisely, Wasserstein distances are replaced by their entropic approximation for a certain fixed entropic regularization parameter (see the introduction of Chapter 7), so that the estimation of Wasserstein distances and barycenters can be done using a finite number of iterations of Sinkhorn-like updates (Peyré and Cuturi, 2019). Being given a set of codes and atoms, these considerations allow to compute the reconstructions and the cost in (8.3). This cost is then minimized using a quasi-Newton descent (L-BFGS), where gradients are estimated using automatic differentiation through the Sinkhorn loops. Although showing successful results in (Schmitz et al., 2018), this approach may not always

be satisfactory in practice. First, it may be quite involved to implement and it does not allow to use existing efficient algorithms tailored for usual Euclidean linear dictionary learning. Second, an entropic regularization parameter must be chosen and constitutes a trade-off between sharpness of the outputs and duration of execution (with additional implementation difficulties when the value of this parameter is small). Third, once a dictionary is learned, the computation of the codes associated to a new probability measure (not present in the training dataset) necessitates the resolution of a new Wasserstein barycenter problem, which may not ensure rapidness of use of the method at test time. Here, we propose instead to leverage again (8.2) in order to approximately solve (8.3). In the one-dimensional case, we observed above that solving (8.3) for a dataset $(\mu_i)_{1 \leq i \leq N}$ is equivalent to solving (8.2) for the dataset $(F_{\mu_i}^{-1})_{1 \leq i \leq N}$. Denoting $\rho = \lambda_{[0,1]}$ the Lebesgue measure on $[0, 1]$, it is well-known that the quantile function $F_{\mu_i}^{-1}$ is a non-decreasing mapping that satisfies $(F_{\mu_i}^{-1})_{\#}\rho = \mu_i$, that is by Brenier's theorem $F_{\mu_i}^{-1} = T_{\mu_i}$ is the optimal transport map from ρ to μ_i . We propose here to extend this idea to dimension greater than one and to represent each measure from the dataset by its LOT embedding: for $d \geq 1$, denote ρ the Lebesgue measure on the unit cube $[0, 1]^d$ and consider a dataset $(\mu_i)_{1 \leq i \leq N} \in \mathcal{P}_2(\mathbb{R}^d)$. For all $i \in \{1, \dots, N\}$, denote T_{μ_i} the optimal transport map from ρ to μ_i and \mathbf{T}_{μ_i} its vectorized version with a certain grid size parameter $m \geq 1$ (see the introduction of this section for the definition of this vectorization). We then propose to approximately solve the Wasserstein dictionary learning problem (8.3) on the dataset $(\mu_i)_{1 \leq i \leq N}$ by solving the following *Linearized Wasserstein dictionary learning*:

$$\min_{D \in \mathbb{R}^{dm^d \times k}, \Lambda \in (\Delta_k)^N} \sum_{i=1}^N \|\mathbf{T}_{\mu_i} - D\Lambda_i\|_2^2 = \|\mathbf{T} - D\Lambda\|_F^2, \quad (8.4)$$

where $\mathbf{T} \in \mathbb{R}^{dm^d \times N}$ is a matrix whose i -th column contains the vectorized transport map \mathbf{T}_{μ_i} . Intuitively, the atoms in D are expected to represent the most salient transport maps of the dataset. However, there is no reason in general that, for a given set of codes $\Lambda \in (\Delta_k)^N$, a minimizer D of problem (8.4) yields atoms that are themselves transport maps. In the following, we will enforce this by imposing that each column of D reads as a convex combination of the elements of the dataset $(\mathbf{T}_{\mu_i})_{1 \leq i \leq N}$. Indeed, convex combinations of gradients of convex functions are gradients of convex functions, so that atoms built from convex combinations of optimal transport maps are optimal transport maps. We thus introduce a weight matrix $W \in (\Delta_n)^k$ such that

$$D = D(W) = \mathbf{T}W \in \mathbb{R}^{dm^d \times k}.$$

Hence our *Linearized Wasserstein dictionary learning* problem reads

$$\min_{W \in (\Delta_n)^k, \Lambda \in (\Delta_k)^N} \|\mathbf{T} - D(W)\Lambda\|_F^2, \quad (8.5)$$

Another advantage of choosing atoms of D as convex combinations of elements of the dataset is that the total size of the unknowns in (8.5) is $2Nk$, which is less in general than the size $(dm^d + N)k$ of the unknowns in (8.4). The objective $\mathcal{L}(W, \Lambda) = \|\mathbf{T} - D(W)\Lambda\|_F^2$ appearing in (8.5) is convex separately in W and Λ , but not jointly in (W, D) . As such, we propose to use a projected coordinate descent for its minimization, taking a gradient step with respect to W and Λ alternatively and projecting their value to their respective simplex $(\Delta_n)^k$ and $(\Delta_k)^N$ using the projection algorithm of (Duchi et al., 2008). The

following gradients can easily be computed:

$$\begin{aligned}\partial_W \mathcal{L}(W, \Lambda) &= -2\mathbf{T}^\top (\mathbf{T} - \mathbf{T}W\Lambda)\Lambda^\top, \\ \partial_\Lambda \mathcal{L}(W, \Lambda) &= -2W^\top \mathbf{T}^\top (\mathbf{T} - \mathbf{T}W\Lambda).\end{aligned}$$

Both of these gradients can easily be set to zero in closed form when the other variable is fixed. We will leverage this only for the update of Λ , and we will use a simple gradient step for the update of W . This proved to yield better results in practice, and it may be explained intuitively by the fact that setting gradients to zero rapidly lead to a local minima of \mathcal{L} , i.e. it enforces a rapid *definitive* choice of both Λ and W , which might not be optimal when starting from random values for both these unknowns. We report our overall projected coordinate descent algorithm for the resolution of (8.5) in Algorithm 2. Note that in this algorithm, the projection onto simplex is done using the algorithm of (Duchi et al., 2008) and for a matrix M , the matrix M^\dagger denotes its pseudo-inverse. We now illustrate the use of the linearized Wasserstein dictionary learning method on two concrete examples.

Algorithm 2: Projected coordinate descent for linearized Wasserstein dictionary learning.

Data: Dataset of vectorized Monge maps $\mathbf{T} \in \mathbb{R}^{dm^d \times N}$, number of atoms k , step-size α , number of iterations n_{iter} .

Result: Dictionary $D \in \mathbb{R}^{dm^d \times k}$, codes $\Lambda \in \mathbb{R}^{k \times N}$

Select W uniformly at random in $[0, 1]^{N \times k}$;

Project each column of W onto Δ_N ;

Select Λ uniformly at random in $[0, 1]^{k \times N}$;

Project each column of Λ onto Δ_k ;

for $t \leftarrow 1$ **to** n_{iter} **do**

$W \leftarrow W + 2\alpha \mathbf{T}^\top (\mathbf{T} - \mathbf{T}W\Lambda)\Lambda^\top$;

Project each column of W onto Δ_N ;

$D \leftarrow \mathbf{T}W$;

$\Lambda \leftarrow (D^\top D)^\dagger D^\top \mathbf{T}$;

Project each column of Λ onto Δ_k ;

end

Cardiac MRI frames. This example is extracted from (Schmitz et al., 2018). In this work, the authors consider a sequence of MRI frames that summarize a beating heart cycle. Seeing each MRI frame (which is a black and white image) as a two-dimensional probability measure over a pixel grid, the authors of (Schmitz et al., 2018) use their Wasserstein dictionary learning algorithm to proceed with the search of a dictionary of four atoms, each atom being expected to represent a key frame of the sequence. Because the 4-dimensional learned codes Λ_i lie in the simplex Δ_4 , the authors of (Schmitz et al., 2018) propose to visualize each code Λ_i in a three-dimensional space (ignoring for instance the last coordinate) without any loss of information. Doing so, they notice that the frames describe a cycle when represented with these barycentric coordinates. As such, the Wasserstein dictionary learning method is able to recover the periodic nature of the sequence of frames. We proceed here with the same task and consider $N = 30$ MRI




Figure 8.9: (Top) Five example cardiac MRI frames from the sequence. (Bottom) Reconstructions of the first row frames from a learned linearized Wasserstein dictionary with $k = 4$ atoms.

frames of a beating heart cycle¹ (note that we could not find the same sequence of images as the one considered in (Schmitz et al., 2018)). We report in the first row of Figure 8.9 some examples of these frames. For $i \in \{1, \dots, 30\}$, we then consider the i -th image as a discrete probability measure μ_i supported over the pixel grid and where each pixel intensity indicates the weight of the corresponding Dirac mass. The LOT embeddings and corresponding vectorized transport maps $(\mathbf{T}_{\mu_i})_{1 \leq i \leq N}$ are then computed as described in the beginning of this chapter, with a grid size parameter $m = 1500$. We finally solve the linearized Wasserstein dictionary learning problem defined above on the dataset $(\mathbf{T}_{\mu_i})_{1 \leq i \leq N}$ using Algorithm 2, with $k = 4$, $\alpha = 7 \times 10^{-6}$ and $n_{iter} = 50$. The bottom row of Figure 8.9 displays example reconstructions. Comparing these reconstructions to the original images in the first row indicates that the method seems to yield satisfactory results. We report in Figure 8.10 the learned atoms (represented as probability measures, i.e. as black and white images, obtained from the push-forwards of ρ by the learned atoms) and the path followed by the sequence in the barycentric coordinates obtained from the codes. This path approximately follows edges of the 3-simplex: this could be expected from the fact that in formulation (8.5), the atoms of the dictionary are by design *convex combinations* of elements of the dataset. Figure 8.10 shows that the learned atoms actually correspond to elements of the dataset since there are frames at each vertex of the 3-simplex. This figure also shows that the whole sequence may be roughly recovered from Wasserstein geodesic interpolations between the four atoms (linear interpolations in the right representation of Figure 8.10 actually correspond to linear interpolations between transport maps, i.e. they correspond to *generalized* Wasserstein geodesic interpolations).

Face recognition. This last example reproduces the face recognition experiment of (Sandler and Lindenbaum, 2011) and (Rolet et al., 2016) on the ORL dataset (Samaria and Harter, 1994). This dataset contains 400 face pictures of 40 different people (with 10 pictures per person). We give in Figure 8.11 example pictures from this dataset.

In (Sandler and Lindenbaum, 2011) and (Rolet et al., 2016), the authors propose to use a dictionary learning method to perform face recognition on the ORL dataset. To do this, they evenly split the whole dataset set into a training set and a test set of 200 pictures each and each containing 5 pictures of the same person. This splitting is performed randomly and affects the subsequent classification performances. As such, only the best performances are reported in (Sandler and Lindenbaum, 2011) and (Rolet

¹found here: <https://www.mhvi.com/wp-content/uploads/2018/09/MRI-anim.gif>.




Figure 8.10: (Left) Learned atoms represented as images. (Right) Path followed by the frames in barycentric coordinates obtained from the codes (each point corresponds to a frame i and its coordinates $((\Lambda_i)_1, (\Lambda_i)_2, (\Lambda_i)_3)$ are the three first coordinates of $\Lambda_i \in \Delta_4$).




Figure 8.11: Example pictures from the ORL dataset (after downsampling), with people's ID shown above the pictures.

et al., 2016) and we will proceed as well in what follows. On the training set, a dictionary of a given number of atoms is learned together with the 200 *training codes* representing the pictures of the training set in the learned dictionary basis. Then, face recognition on the test set is performed as follows: each test picture is projected in the learned dictionary basis, yielding a *test code*. An ID is then attributed to each test picture using the ID of the training code that is the closest from its test code.

Seeing the pictures as vectors in a Euclidean space already allows to use classical dictionary learning methods (formulated in (8.2)) to perform face recognition as we just described. However, the L^2 comparison of images needs a perfect alignment of the faces as well as similar poses to be effective. This is because the L^2 distance between an image and its translation can be very large even for small translations. The consideration of a Wasserstein distance to compare the pictures is thus preferred in (Sandler and Lindenbaum, 2011) and (Rolet et al., 2016), where dictionary learning is performed using such distance to measure similarity (but where reconstruction of histograms is performed in a linear way). We proceed with the same face recognition task using our linearized Wasserstein dictionary learning formulation (8.5) solved with Algorithm 2. As in (Sandler and Lindenbaum, 2011) and (Rolet et al., 2016), the pictures from the original ORL dataset are downsampled (by a factor of 3) and seen as discrete probability measures supported over the pixel grid where each pixel intensity indicates the weight attributed to the corresponding Dirac mass. Again, we represent these probability measure using their LOT embedding (with reference $\rho = \lambda_{[0,1] \times [0,1]}$) and use the vectorization of the transport maps described in the beginning of this chapter, with a discretization parameter $m = 50$. We use the same procedure for splitting the dataset into training and testing sets as well as to perform face recognition. Dictionary learning with $k \geq 1$ atoms is performed on the training set $\mathbf{T}_{train} \in \mathbb{R}^{2m^2 \times 200}$ using algorithm 2, which yields a dictionary $D \in \mathbb{R}^{2m^2 \times k}$ and a set of codes $\Lambda_{train} \in \mathbb{R}^{k \times 200}$. Projection of the test set $\mathbf{T}_{test} \in \mathbb{R}^{2m^2 \times 200}$ on the

Value of k (number of atoms)	2	5	8	10	20	30	40	50	
(Sandler and Lindenbaum, 2011)	8.5	70.5	87.5	94.5	90.5	95	96.5	97	$\leq 20\text{min}$
(Rolet et al. (2016), $\gamma = 1/30$)				93	95.5	97	96.5	96	$\sim 20\text{s}$
(Rolet et al. (2016), $\gamma = 1/50$)				91	95	95	97	94.5	$\sim 90\text{s}$
Linearized Wasserstein dictionary learning	14.5	78.5	95.5	97	97.5	98	98	97.5	$< 7\text{s}$
Natural images	18	76.5	94	94	94.5	96	96	95.5	$< 0.3\text{s}$
Method	Classification accuracy (%)								Run time

Table 8.1: Face recognition performances on the ORL dataset.

dictionary atoms $D \in \mathbb{R}^{2m^2 \times k}$ is done through

$$\Lambda_{test} := (D^\top D)^\dagger D^\top \mathbf{T}_{test} \in \mathbb{R}^{k \times 200},$$

and each column of Λ_{test} is then projected onto Δ_k . Finally, as in (Sandler and Lindenbaum, 2011), each column of Λ_{train} and Λ_{test} are normalized so as to have a unit Euclidean norm before performing nearest neighbor search to assign IDs to test images. We report in Table 8.1 the obtained classification results and compare them to the results of (Sandler and Lindenbaum, 2011) and (Rolet et al., 2016), as well as to the results obtained from the same dictionary learning approach applied to natural images with a L^2 measure of similarity and linear reconstruction. We also report in the last column the execution time of each method (the language and CPU used for these implementations are respectively: Matlab on a 2.5GHz Intel Core 2 Quad for (Sandler and Lindenbaum, 2011), Matlab on a 2.4 GHz Intel Quad core i7 for (Rolet et al., 2016) and Python on a 2.6 GHz Intel Core i7 6-Core for our methods). We can first note that the straightforward linear dictionary learning method with a L^2 measure of similarity, which can be taken as a baseline, already yields very good performance results. This might be due to the fact that in the ORL dataset, faces are already well centered and there are few pose variations. Moreover, the execution time of this baseline method is very short, since there is no optimal transport computation. However, the performances of this L^2 method should be affected from small geometric variations applied to the dataset, while such variations should not affect much the methods based on optimal transport distances. We note that our linearized Wasserstein dictionary learning approach consistently yields good classification results and short run times with respect to the other optimal transport based approaches. This may motivate the use of the LOT framework for this kind of data analysis tasks.

Conclusion

In the optimal transport community, the issue of the quantitative stability of optimal transport quantities could be thought of as *the elephant in the room*: in most applications, it is assumed that such quantities can be approximated from approximations of the problem data, but the blatant lack of quantitative guarantees to back this assumption is hardly ever mentioned. In this thesis, we have given the first piece of answer to this quantitative stability question. Nonetheless, this issue remains far from being completely understood.

In Part I, we have derived explicit strong convexity estimates for the dual of the optimal transport problem, exclusively in the quadratic Euclidean setting and relying mainly on the Brunn-Minkowski and cousin inequalities. There remains many open questions surrounding this sole issue. A first question concerns the necessary and sufficient conditions on a source measure to ensure strong convexity estimates for its associated Kantorovich functional. We conjectured that being absolutely continuous (with bounds on the density) and satisfying a Poincaré-Wirtinger inequality should represent such necessary and sufficient conditions, but this remains to be proven. Similar conditions can also be sought after for the class of potentials on which a Kantorovich functional satisfies strong convexity estimates. Another natural inquiry concerns the extension of the strong convexity estimates of Part I to other domains and ground costs. The *semi-discrete* approach of Chapter 2 has a strong geometrical flavor and relies on the convexity of the Kantorovich potentials that is peculiar to quadratic optimal transport. Similarly, the *continuous* approach of Chapter 3 essentially builds on the Brascamp-Lieb inequality, giving a crucial role to the convexity of the Kantorovich potentials. The adaptation of these two first proofs to more general contexts of costs and domains may thus not easily be carried out. The *entropic* proof of Chapter 4 based on the Prékopa-Leindler inequality is probably the most prone to adaptations to other costs and domains: it relies essentially on algebraic manipulations and there are neither geometric estimates nor use of the convexity of potentials. The simplicity gained in these derivations from the entropic regularization of the original transport problem may motivate the investigation of other types of regularizations as an approach for proving strong convexity estimates in optimal transport. It could also be relevant to examine if similar estimates can be found in unbalanced optimal transport. Finally, there are natural algorithmic consequences of the strong convexity estimates of Part I that have not been investigated in this thesis, concerning in particular the analysis of Newton methods used in the resolution of semi-discrete optimal transport problems.

In Part II, we have built on the estimates from Part I to derive quantitative stability bounds for certain optimal transport quantities. In Chapter 5, we have given bi-Hölder stability estimates for optimal transport maps with respect to their target measure. In the derivation of these estimates, a series of inequalities has been used where each of the inequality was shown to be tight in terms of exponents. However, we do not know whether the final Hölder exponents for the stability of optimal transport maps w.r.t. their target are tight: this constitutes an open question. We also saw that we could interpret the results of Chapter 5 as embeddability guarantees for parts of the 2-Wasserstein space

$(\mathcal{P}_2(\mathbb{R}^d), W_2)$ in a L^2 space with a controlled bi-Hölder distortion. A previous strong negative result showed that the whole 2-Wasserstein space cannot admit such an embedding and, consistently with this result, we have found an embedding with controlled distortion only on (large) subsets of the 2-Wasserstein space. A remaining open question is that of the largest possible subset of $(\mathcal{P}_2(\mathbb{R}^d), W_2)$ that admits a bi-Hölder embedding in a L^2 space. In particular, we found distortion guarantees for a L^2 embedding of p -Wasserstein balls $B_{W_p}(\delta_0, M_p)$ for some $p > d$ and $M_p > 0$, but it would seem more natural to find similar guarantees for the same embedding of 2-Wasserstein balls $B_{W_2}(\delta_0, M_2)$ with $M_2 > 0$. In Chapter 6, we have derived quantitative stability estimates for barycenters in the 2-Wasserstein space with respect to their marginals. These estimates have been derived under compactness assumptions and it would be natural to wonder whether these results can be extended to non-compact settings. However, the current proof strongly relies on the compactness assumption. One could also wonder if such estimates could be found for Wasserstein barycenters defined using other Wasserstein distances than the quadratic one. This constitutes a difficult question. First, the strong convexity estimates of Part I should also be extended to these settings, and we saw in the preceding paragraph that such extensions may not be straightforward. Second, we relied in Chapter 6 on a quantitative stability estimate for the push-forward operation under an optimal transport map: there, the fact that such map reads as the gradient of a convex function appeared to be important in the proof. The extension of this result could thus also stir up some trouble. It is finally worth noticing that the results of chapters 5 and 6 both suggest natural multi-scale strategies for the numerical resolution of optimal transport problems, where the probability measures of interest are discretized and the quality of the discretization is improved over the course of the iterations of an optimal transport solver. Eventually, we have exposed in Chapter 7 convergence bounds for semi-discrete entropic optimal transport solutions with respect to the regularization parameter. As mentioned in the chapter, these bounds could help justify ε -scaling strategies where the regularization parameter is gradually decreased during the numerical resolution of the regularized optimal transport problem. Another interesting numerical approach for the resolution of optimal transport, not limited to the semi-discrete setting, is suggested in Chapter 7. Indeed, by differentiating the optimality condition of the dual entropic problem with respect to the regularization parameter, we have exposed an ODE. This ODE suggests a central path following approach reminiscent of interior point methods: the regularized optimal transport problem could be solved for an initial large value of the regularization parameter and the solution would then be updated using the ODE in order to move along a path of regularized solutions toward a solution of the unregularized problem. It is however not clear how to carry out the analysis of such an approach and whether it would yield significant computational advantages, the system represented by the ODE being possibly of large size depending on the context.

We have finally dedicated Part III to the exposition of illustrations and applications of the *Linearized Optimal Transport* (LOT) framework. We have seen in this final part that the bounds of Chapter 5 on the metric distortion induced by the LOT embedding could be used to prove performance guarantees on data analysis algorithms within this framework. There remains however important computational challenges associated to this approach. Indeed, existing algorithms allow an efficient computation of semi-discrete optimal transport maps in small dimensions ($d = 2$ or 3), but there are no such efficient algorithms in higher dimensions. Several stochastic optimization methods have been proposed, but their convergence can be quite slow in practice. Another difficulty posed by

large values of the dimension lies in the computational representation of optimal transport maps. In Chapter 8, we opted for a uniform sampling over a grid. This approach does not scale well to large dimensions, and other representations should be considered, leveraging in particular the dual Kantorovich potentials whose representations do not depend on the ambient dimension.

Part IV

Appendix

Optimal transport facts

This chapter gathers some (well-known) optimal transport facts that are useful (but tangential) to this thesis. The presentation done here is more general than the one done in Chapter 1 because not limited to the quadratic cost. No new result is presented in this chapter and we refer to the following monographs and chapters (from which this chapter is inspired) for more general presentations of theoretical and computational aspects related to the field: (Villani, 2003, 2008), (Santambrogio, 2015), (Peyré and Cuturi, 2019) and (Mérigot and Thibert, 2021).

A.1 Monge and Kantorovich formulations

A.1.1 Monge formulation.

Let $\rho, \mu \in \mathcal{P}(\mathbb{R}^d)$ that represent two different distributions of mass and a cost function $c : \mathbb{R}^d \times \mathbb{R}^d \rightarrow \mathbb{R}_+$ that encodes with $c(x, y)$ the *cost* of transporting a unit of mass from a location x to a location y . Monge's formulation of the optimal transport problem with ground cost c is the following: look for a transport map $T : \mathbb{R}^d \rightarrow \mathbb{R}^d$ that solves the following non-convex optimization problem:

$$\inf_{T_{\#}\rho=\mu} \int_{\mathbb{R}^d} c(x, T(x)) d\rho(x), \quad (\text{A.1})$$

where $T_{\#}\rho$ corresponds to the push-forward measure of ρ by T , which satisfies $T_{\#}\rho(A) = \rho(T^{-1}(A))$ for every ρ -measurable set A .

A.1.2 Kantorovich formulation.

Kantorovich's formulation of the optimal transport problem with ground cost c is the following convex optimization problem:

$$\inf_{\gamma \in \Gamma(\rho, \mu)} \int_{\mathbb{R}^d \times \mathbb{R}^d} c(x, y) d\gamma(x, y). \quad (\text{A.2})$$

Problem (A.2) is a relaxation of (A.1) in the sense that if there exists a transport map T between ρ and μ , then one can build an admissible transport plan γ_T from it: $\gamma_T := (\text{id}, T)_{\#}\rho \in \Gamma(\rho, \mu)$. Such transport plan verifies in particular

$$\int_{\mathbb{R}^d \times \mathbb{R}^d} c(x, y) d\gamma_T(x, y) = \int_{\mathbb{R}^d} c(x, T(x)) d\rho(x),$$

so that one always has (A.2) \leq (A.1).

Remark A.1. It is natural to wonder under what conditions on ρ, μ and c do the Monge and Kantorovich formulations coincide, i.e. when do we have (A.2) = (A.1)? A well-known case corresponds to the case where the source ρ is absolutely continuous and c is of the form $c(x, y) = h(x - y)$ for some strictly convex function h , see Theorem 1.17 of (Santambrogio, 2015) for a precise statement.

Under very mild conditions on the cost function c , an optimal transport plan always exists:

Theorem A.2 (Villani (2008), Theorem 4.1). *Let $c : \mathbb{R}^d \times \mathbb{R}^d \rightarrow \mathbb{R} \cup \{+\infty\}$ be lower semi-continuous and assume that there exists upper semi-continuous functions $a \in L^1(\rho), b \in L^1(\mu)$ such that*

$$\forall (x, y) \in \mathbb{R}^d, \quad c(x, y) \geq a(x) + b(y).$$

Then problem (A.2) admits a solution.

A.2 Dual formulation

The following dual formulation of (A.2) can be recovered using the method of Lagrange multipliers for the constraints:

$$\sup_{(\phi, \psi) \in L_c^1(\rho, \mu)} \int_{\mathbb{R}^d} \phi d\rho + \int_{\mathbb{R}^d} \psi d\mu, \quad (\text{A.3})$$

where $L_c^1(\rho, \mu) = \{(\phi, \psi) \in L^1(\rho) \times L^1(\mu) \mid \forall (x, y) \in \mathbb{R}^d, \phi(x) + \psi(y) \leq c(x, y)\}$. From the definition of $L_c^1(\rho, \mu)$, it is immediate to check that weak duality holds, i.e. (A.3) \leq (A.2). Before we mention cases of strong-duality (i.e. when (A.3) = (A.2)) and existence of solutions to (A.3), we first notice that problem (A.3) may be turned easily into an unconstrained maximization problem. Indeed, for a given $\phi \in L^1(\rho)$ and $y \in \mathbb{R}^d$, it is interesting when maximizing (A.3) to choose $\psi(y)$ as large as possible and satisfying the constraint in $L_c^1(\rho, \mu)$, i.e. to choose

$$\psi(y) = \inf_{x \in \mathbb{R}^d} c(x, y) - \phi(x).$$

One can think of a similar choice for ϕ when ψ is fixed and such choices correspond to taking ψ and ϕ as c/\bar{c} -transforms of each other:

Definition A.3. Let $c : \mathbb{R}^d \times \mathbb{R}^d \rightarrow (-\infty, +\infty]$. The c -transform of $\phi : \mathbb{R}^d \rightarrow \mathbb{R} \cup \{+\infty\}$ is the function ϕ^c defined by

$$\forall y \in \mathbb{R}^d, \quad \phi^c(y) = \inf_{x \in \mathbb{R}^d} c(x, y) - \phi(x).$$

The \bar{c} -transform of $\psi : \mathbb{R}^d \rightarrow \mathbb{R} \cup \{+\infty\}$ is the function $\psi^{\bar{c}}$ defined by

$$\forall x \in \mathbb{R}^d, \quad \psi^{\bar{c}}(x) = \inf_{y \in \mathbb{R}^d} c(x, y) - \psi(y).$$

A function $\psi : \mathbb{R}^d \rightarrow \mathbb{R} \cup \{+\infty\}$ is then said \bar{c} -concave if there exists ϕ such that $\psi = \phi^c$, and similarly, $\phi : \mathbb{R}^d \rightarrow \mathbb{R} \cup \{+\infty\}$ is said c -concave if there exists ψ such that $\phi = \psi^{\bar{c}}$.

Remark A.4. If the family of functions $y \mapsto c(x, y)$ all share the same modulus of continuity for any $x \in \mathbb{R}^d$, then the c -transform ϕ^c also shares this modulus of continuity as an infimum over this (translated) family. The same holds for the \bar{c} -transform $\psi^{\bar{c}}$ and the modulus of continuity of $x \mapsto c(x, y)$ for any $y \in \mathbb{R}^d$.

In very general cases, such choices can always be made. It allows then to reformulate (A.3) as an unconstrained maximization problem and to show that strong-duality holds:

Theorem A.5 (Villani (2008), Theorem 5.9). *Let $c : \mathbb{R}^d \times \mathbb{R}^d \rightarrow \mathbb{R} \cup \{+\infty\}$ be lower semi-continuous and assume that there exists upper semi-continuous functions $a \in L^1(\rho), b \in L^1(\mu)$ such that*

$$\forall (x, y) \in \mathbb{R}^d, \quad c(x, y) \geq a(x) + b(y).$$

Then strong-duality holds and (A.3) may be replaced by an unconstrained problem:

$$(A.2) = (A.3) = \sup_{\phi \in L^1(\rho)} \int_{\mathbb{R}^d} \phi d\rho + \int_{\mathbb{R}^d} \phi^c d\mu = \sup_{\psi \in L^1(\mu)} \int_{\mathbb{R}^d} \psi^{\bar{c}} d\rho + \int_{\mathbb{R}^d} \psi.$$

Under slightly more restrictive assumptions on the cost function, one can then show existence of solution to (A.3). Further, it is possible to characterize these solutions and their links with solutions to (A.2) using again the notion of c/\bar{c} -transform of Definition A.3 and the related notion of c -subdifferential:

Definition A.6. Let $\phi : \mathbb{R}^d \rightarrow \mathbb{R} \cup \{+\infty\}$. Its c -subdifferential is the set defined by

$$\partial_c \phi = \{(x, y) \in \mathbb{R}^d \times \mathbb{R}^d \mid \phi(x) + \phi^c(y) = c(x, y)\}.$$

The c -subdifferential of ϕ at point $x \in \mathbb{R}^d$ is then

$$\partial_c \phi(x) = \{y \in \mathbb{R}^d \mid (x, y) \in \partial_c \phi\}.$$

Theorem A.7 (Villani (2008), Theorem 5.9). *Let $c : \mathbb{R}^d \times \mathbb{R}^d \rightarrow \mathbb{R}$ be lower semi-continuous and assume that there exists $c_1 \in L^1(\rho), c_2 \in L^1(\mu)$ and upper semi-continuous functions $a \in L^1(\rho), b \in L^1(\mu)$ such that*

$$\forall (x, y) \in \mathbb{R}^d, \quad a(x) + b(y) \leq c(x, y) \leq c_1(x) + c_2(y).$$

Then both (A.2) and (A.3) admit solutions, so that

$$(A.2) = \min_{\gamma \in \Gamma(\rho, \mu)} \int_{\mathbb{R}^d \times \mathbb{R}^d} c(x, y) d\gamma(x, y) = \max_{\phi \in L^1(\rho)} \int_{\mathbb{R}^d} \phi d\rho + \int_{\mathbb{R}^d} \phi^c d\mu = (A.3),$$

where one might impose that ϕ is c -concave. If in addition, a, b and c are continuous, then any $\gamma \in \Gamma(\rho, \mu)$ and c -concave $\phi \in L^1(\rho)$ are optimal for (A.2), (A.3) respectively if and only if γ is concentrated on $\partial_c \phi$.

The specialization of Theorem A.7 to specific choices of cost functions allows us to get the following well-known results.

Kantorovich-Rubinstein duality formula. When $c(x, y) = \|x - y\|$, it is clear that the sets of c - and \bar{c} -concave functions coincide and correspond to the set of 1-Lipschitz functions. A consequence of this fact and of Theorem A.7 is the following formula, known as *Kantorovich–Rubinstein duality result* (Kantorovich and Rubinstein, 1958):

Proposition A.8 (Villani (2008), Particular Case 5.15). *Let $\rho, \mu \in \mathcal{P}_1(\mathbb{R}^d)$. Then*

$$\min_{\gamma \in \Gamma(\rho, \mu)} \int_{\mathbb{R}^d \times \mathbb{R}^d} \|x - y\| d\gamma(x, y) = \max_{f \in 1\text{-Lip}(\mathbb{R}^d)} \int_{\mathbb{R}^d} f d\rho - \int_{\mathbb{R}^d} f d\mu.$$

This formula shows in particular that the optimal transport cost for the ground cost $c(x, y) = \|x - y\|$ corresponds to a dual norm on the space of signed measures with vanishing total mass.

Quadratic cost and Legendre transforms. For the *quadratic cost* $c(x, y) = \|x - y\|^2$, the notions of c/\bar{c} -transforms may be related to the usual notion of Legendre transform or convex conjugate of convex analysis (see Remark 1.6). Indeed, it was already noticed in Section 1.1 that working with the quadratic cost $c(x, y) = \|x - y\|^2$ is equivalent to working with the (negative) bilinear cost $c(x, y) = -\langle x | y \rangle$. With such cost, it is easy to check that the c -transform (which is equal to the \bar{c} -transform) verifies for any $\phi : \mathbb{R}^d \rightarrow \mathbb{R} \cup \{+\infty\}$,

$$\phi^c = \phi^{\bar{c}} = -(-\phi)^*.$$

These considerations, together with Theorem A.7 and a manipulation of the signs, allow to recover the semi-dual formulation of the quadratic optimal transport described in Section 1.1.

A.3 Stability of solutions

In order to study the stability properties of solutions to problems (A.1), (A.2) and (A.3), we need a notion of convergence for probability measures supported over \mathbb{R}^d and we will consider the one in duality with $\mathcal{C}_b(\mathbb{R}^d)$:

Definition A.9 (Weak convergence of probability measures). A sequence $(\rho_n)_{n \geq 0} \in \mathcal{P}(\mathbb{R}^d)$ is said to converge weakly to $\rho \in \mathcal{P}(\mathbb{R}^d)$ (denoted $\rho_n \rightharpoonup \rho$) if and only if for any $\phi \in \mathcal{C}_b(\mathbb{R}^d)$, $\int_{\mathbb{R}^d} \phi d\rho_n \rightarrow \int_{\mathbb{R}^d} \phi d\rho$.

Thanks to Theorem A.7 and the optimality conditions it yields for solutions to (A.2) and (A.3), it is possible to show the following stability result in the compact setting:

Theorem A.10 (Villani (2008), Theorem 5.19, and Santambrogio (2015), Theorems 1.51 and 1.52). *Let Ω be a compact subset of \mathbb{R}^d and $c : \Omega \times \Omega \rightarrow \mathbb{R}$ continuous. Let $(\rho_n)_{n \geq 0}, (\mu_n)_{n \geq 0}$ be sequences in $\mathcal{P}(\Omega)$ and assume that $\rho_n \rightharpoonup \rho \in \mathcal{P}(\Omega)$ and $\mu_n \rightharpoonup \mu \in \mathcal{P}(\Omega)$. Then for each $n \geq 0$, denoting $\gamma_n \in \Gamma(\rho_n, \mu_n)$ an optimal transport plan and $\phi_n \in L^1(\rho_n)$ a Kantorovich potential for the transport problem between ρ_n and μ_n w.r.t. c , one has, up to a subsequence:*

(i) $\gamma_n \rightharpoonup \gamma \in \Gamma(\rho, \mu)$, where γ is an optimal transport plan between ρ and μ w.r.t. c .

(ii) $\min_{\gamma \in \Gamma(\rho_n, \mu_n)} \int c d\gamma \rightarrow \min_{\gamma \in \Gamma(\rho, \mu)} \int c d\gamma$.

(iii) $\|\phi_n - \phi\|_{L^\infty(\Omega)} \rightarrow 0$ and $\|\phi_n^c - \phi^c\|_{L^\infty(\Omega)} \rightarrow 0$, where $\phi \in L^1(\rho)$ is a Kantorovich potential between ρ and μ w.r.t. c .

In Theorem A.10, (i) is proven using Prokhorov's theorem and the optimality conditions of Theorem A.7 to show existence and optimality of a limit γ . Point (ii) is then a direct consequence of point (i). Point (iii) is finally obtained by exhibiting a limit ϕ using that c -concave functions have the same modulus of continuity as c (see Remark A.4) and the Arzelà-Ascoli theorem, the optimality of the limit being shown eventually using point (ii). Theorem A.10 thus ensures the stability of problems (A.2), (A.3); and it can be used to show the stability of Monge's original problem (A.1) in the case (A.1) = (A.2):

Corollary A.11 (Villani (2008), Corollary 5.21). *With the same assumptions and notation of Theorem A.10, further assume that there exist measurable maps $T_n, T : \Omega \rightarrow \Omega$ such that*

$$\gamma_n = (\text{id}, T_n)_\# \rho_n, \quad \gamma = (\text{id}, T)_\# \rho,$$

and γ is the unique optimal transport plan between ρ and μ . Then

$$\forall \varepsilon > 0, \quad \rho_n(\{x \in \Omega, \|T_n(x) - T(x)\| > \varepsilon\}) \xrightarrow{n \rightarrow +\infty} 0.$$

In particular, if $\rho_n = \rho$ for all n , then T_n converges to T in ρ -probability.

A.4 Wasserstein distances and spaces

The value of (A.2) represents the *cost* of transporting ρ to μ when the ground cost is c . Thus, intuitively, this value gives a quantitative idea of how similar ρ and μ are, in the sense that the *cheaper* it is to transport ρ to μ , the more *similar* they are likely to be. When the cost is the p -th power of the Euclidean distance, this intuition is made rigorous since the value of (A.2) defines an actual distance between probability measures:

Definition A.12 (Wasserstein distances). Let $p \geq 1$. For any two probability measures $\rho, \mu \in \mathcal{P}_p(\mathbb{R}^d)$, the Wasserstein distance of order p between ρ and μ is defined by

$$W_p(\rho, \mu) = \left(\min_{\gamma \in \Gamma(\rho, \mu)} \int_{\mathbb{R}^d \times \mathbb{R}^d} \|x - y\|^p d\gamma(x, y) \right)^{1/p}.$$

One can quite easily check that W_p satisfies the axioms of a distance over $\mathcal{P}_p(\mathbb{R}^d)$, see e.g. Proposition 5.1 of (Santambrogio, 2015).

Remark A.13. For any $1 \leq p \leq q$, one has $W_p \leq W_q$ using Jensen inequality.

Having defined a distance on $\mathcal{P}_p(\mathbb{R}^d)$, we can then naturally define the following metric space:

Definition A.14 (Wasserstein spaces). The Wasserstein space of $\Omega \subseteq \mathbb{R}^d$ of order $p \geq 1$ is the metric space $(\mathcal{P}_p(\Omega), W_p)$ of probability measures supported over Ω with finite p -th moment, endowed with the Wasserstein distance of order p .

For $\Omega \subseteq \mathbb{R}^d$, the metric space $(\mathcal{P}_p(\Omega), W_p)$ is naturally endowed with the topology induced by the metric W_p . An interesting consequence of the stability result of Theorem A.10 is that whenever Ω is compact, this topology coincides with the weak topology on $\mathcal{P}_p(\Omega)$:

Theorem A.15 (Santambrogio (2015), Theorem 5.10). *If $\Omega \subset \mathbb{R}^d$ is compact and $p \geq 1$, in the space $(\mathcal{P}_p(\Omega), W_p)$ we have $\rho_n \rightharpoonup \rho$ if and only if $W_p(\rho_n, \rho) \rightarrow 0$.*

Wasserstein distances are thus said to metrize the notion of weak-convergence of compactly supported probability measures. In the case of non-compactly supported probability measures, they metricize a stronger notion of convergence:

Theorem A.16 (Santambrogio (2015), Theorem 5.11). *Let $p \geq 1$. In the space $(\mathcal{P}_p(\mathbb{R}^d), W_p)$ we have $W_p(\rho_n, \rho) \rightarrow 0$ if and only if $\rho_n \rightharpoonup \rho$ and $M_p(\rho_n) \rightarrow M_p(\rho)$.*

A final attractive feature of Wasserstein distances of order $p \geq 1$ is that they define a geometry on $\mathcal{P}_p(\mathbb{R}^d)$. In fact, one can show that for $p \geq 1$, the metric space $(\mathcal{P}_p(\mathbb{R}^d), W_p)$ is a geodesic space:

Theorem A.17 (Santambrogio (2015), Theorem 5.27). *Let $p \geq 1$ and $\rho, \mu \in \mathcal{P}_p(\mathbb{R}^d)$. Let $\gamma \in \Gamma(\rho, \mu)$ be an optimal transport plan w.r.t. the cost $c(x, y) = \|x - y\|^p$ and define the maps $\pi_t : (x, y) \mapsto (1 - t)x + ty$ for $t \in [0, 1]$. A constant speed geodesic connecting ρ and μ in $(\mathcal{P}_p(\mathbb{R}^d), W_p)$ is the curve $\rho_t := (\pi_t)_\# \gamma$. In particular if there exists an optimal transport map T such that $\gamma = (id, T)_\# \rho$, the curve can be obtained as $\rho_t = ((1 - t) + tT)_\# \rho$.*

Bibliography

- Adams, S., Dirr, N., Peletier, M. A., and Zimmer, J. (2011). From a Large-Deviations Principle to the Wasserstein Gradient Flow: A New Micro-Macro Passage. *Communications in Mathematical Physics*, 307(3):791.
- Agueh, M. and Carlier, G. (2011). Barycenters in the Wasserstein space. *SIAM Journal on Mathematical Analysis*, 43(2):904–924.
- Ahidar-Coutrix, A., Le Gouic, T., and Paris, Q. (2020). Convergence rates for empirical barycenters in metric spaces: curvature, convexity and extendable geodesics. *Probability Theory and Related Fields*, 177(1):323–368.
- Alaux, J., Grave, E., Cuturi, M., and Joulin, A. (2019). Unsupervised Hyper-alignment for Multilingual Word Embeddings. In *International Conference on Learning Representations*.
- Altschuler, J., Weed, J., and Rigollet, P. (2017). Near-linear time approximation algorithms for optimal transport via sinkhorn iteration. In Guyon, I., Luxburg, U. V., Bengio, S., Wallach, H., Fergus, R., Vishwanathan, S., and Garnett, R., editors, *Advances in Neural Information Processing Systems 30*, pages 1964–1974. Curran Associates, Inc.
- Altschuler, J. M. and Boix-Adsera, E. (2021). Wasserstein barycenters can be computed in polynomial time in fixed dimension. *Journal of Machine Learning Research*, 22(44):1–19.
- Altschuler, J. M., Niles-Weed, J., and Stromme, A. J. (2022). Asymptotics for semidiscrete entropic optimal transport. *SIAM Journal on Mathematical Analysis*, 54(2):1718–1741.
- Ambrosio, L., Gigli, N., and Savaré, G. (2008). *Gradient flows: in metric spaces and in the space of probability measures*. Springer Science & Business Media.
- Ambrosio, L., Stra, F., and Trevisan, D. (2019). A PDE approach to a 2-dimensional matching problem. *Probability Theory and Related Fields*, 173(1):433–477.
- Andoni, A., Naor, A., and Neiman, O. (2018). Snowflake universality of Wasserstein spaces. *Annales Scientifiques de l’Ecole Normale Supérieure*, 51(3):657–700.
- Arjovsky, M., Chintala, S., and Bottou, L. (2017). Wasserstein generative adversarial networks. In Precup, D. and Teh, Y. W., editors, *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, pages 214–223, International Convention Centre, Sydney, Australia. PMLR.
- Arthur, D. and Vassilvitskii, S. (2007). K-means++: The advantages of careful seeding. In *Proceedings of the Eighteenth Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA ’07, page 1027–1035, USA. Society for Industrial and Applied Mathematics.

- Attouch, H. (1984). *Variational convergence for functions and operators*. Applicable Mathematics Series. Pitman (Advanced Publishing Program), Boston, MA.
- Attouch, H., Buttazzo, G., and Michaille, G. (2014). *Variational analysis in Sobolev and BV spaces: applications to PDEs and optimization*, volume 17 of *MOS-SIAM Series on Optimization*. Society for Industrial and Applied Mathematics.
- Aurenhammer, F., Hoffmann, F., and Aronov, B. (1998). Minkowski-type theorems and least-squares clustering. *Algorithmica*, 20(1):61–76.
- Bakry, D., Gentil, I., and Ledoux, M. (2014). *Analysis and Geometry of Markov Diffusion operators*. Grundlehren der mathematischen Wissenschaften, Vol. 348. Springer.
- Ball, K. (1997). An elementary introduction to modern convex geometry. *Flavors of geometry*, 31:1–58.
- Barthe, F. (1997). *Inégalités fonctionnelles et géométriques obtenues par transport de mesures*. PhD thesis, Université de Marne-la-Vallée.
- Basu, S., Kolouri, S., and Rohde, G. K. (2014). Detecting and visualizing cell phenotype differences from microscopy images using transport-based morphometry. *Proceedings of the National Academy of Sciences*, 111(9):3448–3453.
- Benamou, J.-D., Carlier, G., Cuturi, M., Nenna, L., and Peyré, G. (2015). Iterative bregman projections for regularized transportation problems. *SIAM Journal on Scientific Computing*, 37(2):A1111–A1138.
- Benamou, J.-D., Carlier, G., Mérigot, Q., and Oudet, E. (2016). Discretization of functionals involving the Monge–Ampère operator. *Numerische mathematik*, 134(3):611–636.
- Bercu, B. and Bigot, J. (2020). Asymptotic distribution and convergence rates of stochastic algorithms for entropic optimal transportation between probability measures. *Annals of Statistics*.
- Berman, R. J. (2020). Convergence rates for discretized monge–ampère equations and quantitative stability of optimal transport. *Foundations of Computational Mathematics*.
- Berman, R. J. and Berndtsson, B. (2013). Real Monge–Ampère equations and Kähler–Ricci solitons on toric log-Fano varieties. *Annales de la Faculté des sciences de Toulouse: Mathématiques*, 22(4):649–711.
- Bernard, P. and Buffoni, B. (2007). Optimal mass transportation and Mather theory. *Journal of the European Mathematical Society*, 9:85–121.
- Bernton, E. (2018). Langevin Monte Carlo and JKO splitting. In Bubeck, S., Perchet, V., and Rigollet, P., editors, *Proceedings of the 31st Conference On Learning Theory*, volume 75 of *Proceedings of Machine Learning Research*, pages 1777–1798. PMLR.
- Bernton, E., Ghosal, P., and Nutz, M. (2022). Entropic Optimal Transport: Geometry and Large Deviations. *Duke Mathematical Journal*.
- Bertsekas, D. P. (1981). A new algorithm for the assignment problem. *Mathematical Programming*, 21(1):152–171.

- Bertsekas, D. P. and Eckstein, J. (1988). Dual coordinate step methods for linear network flow problems. *Mathematical Programming*, 42(1):203–243.
- Bigot, J., Cazelles, E., and Papadakis, N. (2019a). Central limit theorems for entropy-regularized optimal transport on finite spaces and statistical applications. *Electronic Journal of Statistics*, 13(2):5120 – 5150.
- Bigot, J., Cazelles, E., and Papadakis, N. (2019b). Penalization of Barycenters in the Wasserstein Space. *SIAM Journal on Mathematical Analysis*, 51(3):2261–2285.
- Bigot, J., Gouet, R., Klein, T., and López, A. (2018). Upper and lower risk bounds for estimating the Wasserstein barycenter of random measures on the real line. *Electronic Journal of Statistics*, 12(2):2253 – 2289.
- Bigot, J. and Klein, T. (2018). Characterization of barycenters in the Wasserstein space by averaging optimal transport maps. *ESAIM: PS*, 22:35–57.
- Bobkov, S. G. and Ledoux, M. (2000). From Brunn-Minkowski to Brascamp-Lieb and to logarithmic Sobolev inequalities. *Geometric and Functional Analysis*, 10(5):1028–1052.
- Boissard, E., Le Gouic, T., and Loubes, J.-M. (2015). Distribution’s template estimate with Wasserstein metrics. *Bernoulli*, 21(2):740 – 759.
- Bonneel, N., Peyré, G., and Cuturi, M. (2016). Wasserstein Barycentric Coordinates: Histogram Regression Using Optimal Transport. *ACM Transactions on Graphics*, 35(4):71:1–71:10.
- Bouchitté, G. and Buttazzo, G. (2001). Characterization of optimal shapes and masses through Monge-Kantorovich equation. *Journal of the European Mathematical Society*, 3(2):139–168.
- Boyd, S. and Vandenberghe, L. (2004). *Convex Optimization*. Cambridge University Press.
- Brascamp, H. and Lieb, E. (1976). On extensions of the Brunn-Minkowski and Prékopa-Leindler theorems, including inequalities for log concave functions, and with an application to the diffusion equation. *Journal of Functional Analysis*, 22(4):366–389.
- Brenier, Y. (1987). Décomposition polaire et réarrangement monotone des champs de vecteurs. (Polar decomposition and increasing rearrangement of vector fields). *C. R. Acad. Sci., Paris, Sér. I*, 305:805–808.
- Brenier, Y. (1989). The least action principle and the related concept of generalized flows for incompressible perfect fluids. *Journal of the American Mathematical Society*, 2(2):225–255.
- Brenier, Y. (1991). Polar factorization and monotone rearrangement of vector-valued functions. *Communications on Pure and Applied Mathematics*, 44(4):375–417.
- Brenier, Y. (1999). Minimal geodesics on groups of volume-preserving maps and generalized solutions of the euler equations. *Communications on Pure and Applied Mathematics*, 52(4):411–452.
- Brézis, H. (1973). *Opérateurs maximaux monotones et semi-groupes de contractions dans les espaces de Hilbert*. American Elsevier Publishing Co., Inc., New York.

- Brunn, H. (1887). *Über Ovale und Eiflächen*. Inaugural Dissertation, München.
- Buttazzo, G., De Pascale, L., and Gori-Giorgi, P. (2012). Optimal-transport formulation of electronic density-functional theory. *Phys. Rev. A*, 85:062502.
- Caffarelli, L. A., Feldman, M., and McCann, R. J. (2002). Constructing optimal maps for Monge’s transport problem as a limit of strictly convex costs. *Journal of the American Mathematical Society*, 15(1):1–26.
- Caffarelli, L. A., Kochengin, S. A., and Oliker, V. I. (1999). Problem of reflector design with given far-field scattering data. *Monge Ampère equation: applications to geometry and optimization*, 226:13.
- Caffarelli, L. A. and Oliker, V. I. (2008). Weak solutions of one inverse problem in geometric optics. *Journal of Mathematical Sciences*, 154(1):39–49.
- Cai, T., Cheng, J., Craig, N., and Craig, K. (2020). Linearized optimal transport for collider events. *Phys. Rev. D*, 102:116019.
- Canas, G. and Rosasco, L. (2012). Learning probability measures with respect to optimal transport metrics. *Advances in Neural Information Processing Systems*, 4.
- Carlen, E. A., Cordero-Erausquin, D., and Lieb, E. H. (2013). Asymmetric covariance estimates of Brascamp–Lieb type and related inequalities for log-concave measures. *Annales de l’Institut Henri Poincaré, Probabilités et Statistiques*, 49(1):1 – 12.
- Carlier, G., Chernozhukov, V., and Galichon, A. (2016). Vector quantile regression: An optimal transport approach. *The Annals of Statistics*, 44(3):1165 – 1192.
- Carlier, G., Delalande, A., and Merigot, Q. (2022). Quantitative Stability of Barycenters in the Wasserstein Space. *arXiv:2209.10217*.
- Carlier, G., Eichinger, K., and Kroshnin, A. (2021). Entropic-Wasserstein Barycenters: PDE Characterization, Regularity, and CLT. *SIAM Journal on Mathematical Analysis*, 53(5):5880–5914.
- Carlier, Guillaume, Oberman, Adam, and Oudet, Edouard (2015). Numerical methods for matching for teams and Wasserstein barycenters. *ESAIM: M2AN*, 49(6):1621–1642.
- Cazelles, E., Seguy, V., Bigot, J., Cuturi, M., and Papadakis, N. (2018). Geodesic PCA versus log-PCA of histograms in the Wasserstein space. *SIAM J. Sci. Comput.*, 40(2):429–456.
- Chazal, F., Cohen-Steiner, D., Lieutier, A., Mérigot, Q., and Thibert, B. (2017). Inference of curvature using tubular neighborhoods. In *Modern Approaches to Discrete Curvature*, pages 133–158. Springer.
- Chazal, F., Cohen-Steiner, D., and Mérigot, Q. (2010). Boundary measures for geometric inference. *Foundations of Computational Mathematics*, 10(2):221–240.
- Chernozhukov, V., Galichon, A., Hallin, M., and Henry, M. (2017). Monge–Kantorovich depth, quantiles, ranks and signs. *Annals of Statistics*, 45(1):223–256.

- Chewi, S., Maunu, T., Rigollet, P., and Stromme, A. (2020). Gradient descent algorithms for Bures-Wasserstein barycenters. In Abernethy, J. and Agarwal, S., editors, *Proceedings of Thirty Third Conference on Learning Theory*, volume 125 of *Proceedings of Machine Learning Research*, pages 1276–1304. PMLR.
- Cohen, G., Afshar, S., Tapson, J., and van Schaik, A. (2017). EMNIST: an extension of MNIST to handwritten letters.
- Colombo, P., Staerman, G., Piantanida, P., and Clavel, C. (2021). Automatic Text Evaluation through the Lens of Wasserstein Barycenters. In *EMNLP 2021*, Punta Cana, Dominica.
- Cominetti, R. and Martín, J. S. (1994). Asymptotic analysis of the exponential penalty trajectory in linear programming. *Mathematical Programming*, 67(1):169–187.
- Conforti, G. and Tamanini, L. (2021). A formula for the time derivative of the entropic cost and applications. *Journal of Functional Analysis*, 280(11):108964.
- Cordero-Erausquin, D. (1999). Inégalité de Prékopa-Leindler sur la sphère. *Comptes Rendus de l'Académie des Sciences - Series I - Mathematics*, 329(9):789–792.
- Cordero-Erausquin, D. (2005). On Berndtsson’s generalization of Prékopa’s theorem. *Mathematische Zeitschrift*, 249(2):401–410.
- Cordero-Erausquin, D., McCann, R., and Schmuckenschlager, M. (2001). A Riemannian interpolation inequality a la Borell, Brascamp and Lieb. *Inventiones Mathematicae*, 146(2):219–257.
- Cordero-Erausquin, D., McCann, R. J., and Schmuckenschläger, M. (2006). Prékopa–leindler type inequalities on Riemannian manifolds, Jacobi fields, and optimal transport. *Annales de la Faculté des sciences de Toulouse : Mathématiques*, Ser. 6, 15(4):613–635.
- Cotar, C., Friesecke, G., and Klüppelberg, C. (2013). Density functional theory and optimal transportation with coulomb cost. *Communications on Pure and Applied Mathematics*, 66(4):548–599.
- Courty, N., Flamary, R., Habrard, A., and Rakotomamonjy, A. (2017). Joint distribution optimal transportation for domain adaptation. In *NIPS 2017*, Los Angeles, United States.
- Cullen, M., Norbury, J., and Purser, R. (1991). Generalised Lagrangian solutions for atmospheric and oceanic flows. *SIAM Journal on Applied Mathematics*, 51(1):20–31.
- Cuturi, M. (2013). Sinkhorn Distances: Lightspeed Computation of Optimal Transport. In *Proceedings of the 26th International Conference on Neural Information Processing Systems - Volume 2*, NIPS’13, pages 2292–2300.
- Cuturi, M. and Doucet, A. (2014). Fast Computation of Wasserstein Barycenters. In Xing, E. P. and Jebara, T., editors, *Proceedings of the 31st International Conference on Machine Learning*, volume 32(2) of *Proceedings of Machine Learning Research*, pages 685–693, Beijing, China. PMLR.

- Dalalyan, A. (2017). Further and stronger analogy between sampling and optimization: Langevin monte carlo and gradient descent. In Kale, S. and Shamir, O., editors, *Proceedings of the 2017 Conference on Learning Theory*, volume 65 of *Proceedings of Machine Learning Research*, pages 678–689. PMLR.
- Dantzig, G. B. (1949). Programming of interdependent activities: Ii mathematical model. *Econometrica*, 17(3/4):200–211.
- Dantzig, G. B. (1951). Application of the simplex method to a transportation problem. *Activity analysis and production and allocation*, pages 359–373.
- d’Aspremont, A., Ghaoui, L. E., Jordan, M. I., and Lanckriet, G. R. G. (2007). A Direct Formulation for Sparse PCA Using Semidefinite Programming. *SIAM Review*, 49(3):434–448.
- De Goes, F., Breeden, K., Ostromoukhov, V., and Desbrun, M. (2012). Blue noise through optimal transport. *ACM Transactions on Graphics (TOG)*, 31(6):171.
- De Philippis, G. and Figalli, A. (2014). The Monge–Ampère equation and its link to optimal transportation. *Bulletin of the American Mathematical Society*, 51(4):527–580.
- Delalande, A. (2022). Nearly Tight Convergence Bounds for Semi-discrete Entropic Optimal Transport. In Camps-Valls, G., Ruiz, F. J. R., and Valera, I., editors, *Proceedings of The 25th International Conference on Artificial Intelligence and Statistics*, volume 151 of *Proceedings of Machine Learning Research*, pages 1619–1642. PMLR.
- Delalande, A. and Mérigot, Q. (2021). Quantitative Stability of Optimal Transport Maps under Variations of the Target Measure. *Duke Mathematical Journal (to appear)*.
- Delon, J. (2006). Movie and video scale-time equalization application to flicker reduction. *IEEE Transactions on Image Processing*, 15(1):241–248.
- Divol, V. (2022). Measure estimation on manifolds: an optimal transport approach. *Probability Theory and Related Fields*, 183(1):581–647.
- Dognin, P., Melnyk, I., Mroueh, Y., Ross, J., Santos, C. D., and Sercu, T. (2019). Wasserstein Barycenter Model Ensembling. In *International Conference on Learning Representations*.
- Duchi, J., Shalev-Shwartz, S., Singer, Y., and Chandra, T. (2008). Efficient Projections onto the L1-Ball for Learning in High Dimensions. In *Proceedings of the 25th International Conference on Machine Learning, ICML ’08*, page 272–279, New York, NY, USA. Association for Computing Machinery.
- Duong, M. H., Laschos, V., and Renger, M. (2013). Wasserstein gradient flows from large deviations of many-particle limits. *ESAIM: Control, Optimisation and Calculus of Variations*, 19(4):1166–1188.
- Dvurechensky, P., Gasnikov, A., and Kroshnin, A. (2018). Computational Optimal Transport: Complexity by Accelerated Gradient Descent Is Better Than by Sinkhorn’s Algorithm. In Dy, J. and Krause, A., editors, *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pages 1367–1376, Stockholm, Sweden. PMLR.

- Ekeland, I. and Témam, R. (1999). *Convex Analysis and Variational Problems*. Society for Industrial and Applied Mathematics.
- Erbar, M., Maas, J., and Renger, M. (2015). From large deviations to Wasserstein gradient flows in multiple dimensions. *Electronic Communications in Probability*, 20(none):1 – 12.
- Evans, L. C. (2010). *Partial differential equations*. American Mathematical Society, Providence, R.I.
- Eymard, R., Gallouët, T., and Herbin, R. (2000). Finite volume methods. *Handbook of numerical analysis*, 7:713–1018.
- Fathi, A. and Figalli, A. (2010). Optimal transportation on non-compact manifolds. *Israel Journal of Mathematics*, 175(1):1–59.
- Feydy, J. (2020). *Analyse de données géométriques, au delà des convolutions*. Theses, Université Paris-Saclay.
- Flamary, R., Courty, N., Gramfort, A., Alaya, M. Z., Boisbunon, A., Chambon, S., Chapel, L., Corenflos, A., Fatras, K., Fournier, N., Gautheron, L., Gayraud, N. T., Janati, H., Rakotomamonjy, A., Redko, I., Rolet, A., Schutz, A., Seguy, V., Sutherland, D. J., Tavenard, R., Tong, A., and Vayer, T. (2021). POT: Python Optimal Transport. *Journal of Machine Learning Research*, 22(78):1–8.
- Flamary, R., Cuturi, M., Courty, N., and Rakotomamonjy, A. (2018). Wasserstein Discriminant Analysis. *Machine Learning*, 107(12):1923–1945.
- Fletcher, T. P., Lu, C., Pizer, S. M., and Joshi, S. (2004). Principal geodesic analysis for the study of nonlinear statistics of shape. *IEEE Transactions on Medical Imaging*, 23(8):995–1005.
- Föllmer, H. (1988). Random fields and diffusion processes. In Hennequin, P.-L., editor, *École d’Été de Probabilités de Saint-Flour XV–XVII, 1985–87*, pages 101–203, Berlin, Heidelberg. Springer Berlin Heidelberg.
- Ford, L. R. and Fulkerson, D. R. (1962). *Flows in Networks*. Princeton University Press.
- Fournier, N. and Guillin, A. (2015). On the rate of convergence in Wasserstein distance of the empirical measure. *Probability Theory and Related Fields*, 162(3):707–738.
- Friedland, S. and Nabben, R. (2000). On Cheeger-type inequalities for weighted graphs. *Journal of Graph Theory*, 41.
- Fukunaga, T. and Kasai, H. (2021). Wasserstein k-means with sparse simplex projection. In *2020 25th International Conference on Pattern Recognition (ICPR)*, pages 1627–1634.
- Galerie, B., Leclaire, A., and Rabin, J. (2017). Semi-Discrete Optimal Transport in Patch Space for Enriching Gaussian Textures. In *Geometric Science of Information*, volume 10589 of *Lecture Notes in Computer Science*, Paris, France.
- Galerie, B., Leclaire, A., and Rabin, J. (2018). A Texture Synthesis Model Based on Semi-discrete Optimal Transport in Patch Space. *SIAM Journal on Imaging Sciences*.

- Galichon, A. (2016). *Optimal Transport Methods in Economics*. Princeton University Press.
- Gangbo, W. and McCann, R. J. (1996). The geometry of optimal transportation. *Acta Mathematica*, 177(2):113–161.
- Gardner, R. (2001). The Brunn-Minkowski Inequality. *Bulletin of the American Mathematical Society*, 39.
- Genevay, A., Chizat, L., Bach, F., Cuturi, M., and Peyré, G. (2019). Sample Complexity of Sinkhorn Divergences. In Chaudhuri, K. and Sugiyama, M., editors, *Proceedings of the Twenty-Second International Conference on Artificial Intelligence and Statistics*, volume 89 of *Proceedings of Machine Learning Research*, pages 1574–1583. PMLR.
- Genevay, A., Cuturi, M., Peyré, G., and Bach, F. (2016). Stochastic optimization for large-scale optimal transport. In *Advances in neural information processing systems*, pages 3440–3448.
- Genevay, A., Peyre, G., and Cuturi, M. (2018). Learning Generative Models with Sinkhorn Divergences. In Storkey, A. and Perez-Cruz, F., editors, *Proceedings of the Twenty-First International Conference on Artificial Intelligence and Statistics*, volume 84 of *Proceedings of Machine Learning Research*, pages 1608–1617, Playa Blanca, Lanzarote, Canary Islands. PMLR.
- Gigli, N. (2011). On Hölder continuity-in-time of the optimal transport map towards measures along a curve. *Proceedings of the Edinburgh Mathematical Society*, 54(2):401–409.
- Goldberg, A. V. and Tarjan, R. E. (1989). Finding minimum-cost circulations by canceling negative cycles. *J. ACM*, 36(4):873–886.
- Gordaliza, P., Barrio, E. D., Fabrice, G., and Loubes, J.-M. (2019). Obtaining fairness using optimal transport theory. In Chaudhuri, K. and Salakhutdinov, R., editors, *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pages 2357–2365, Long Beach, California, USA. PMLR.
- Grünbaum, B. (1960). Partitions of mass-distributions and of convex bodies by hyperplanes. *Pacific Journal of Mathematics*, 10(4):1257 – 1261.
- Gunsilius, F. F. (2022). On the convergence rate of potentials of Brenier maps. *Economic Theory*, 38(2):381–417.
- Hadamard, J. (1902). Sur les problèmes aux dérivées partielles et leur signification physique. *Princeton University Bulletin*, 13:49–52.
- Haker, S., Zhu, L., Tannenbaum, A., and Angenent, S. (2004). Optimal Mass Transport for Registration and Warping. *International Journal of Computer Vision*, 60(3):225–240.
- Ho, N., Nguyen, X., Yurochkin, M., Bui, H. H., Huynh, V., and Phung, D. (2017). Multilevel clustering via Wasserstein means. In Precup, D. and Teh, Y. W., editors, *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, pages 1501–1509. PMLR.

- Hug, D. and Weil, W. (2020). *Lectures on convex geometry*. Graduate texts in mathematics. 286. Springer, Cham, Switzerland.
- Hütter, J.-C. and Rigollet, P. (2021). Minimax estimation of smooth optimal transport maps. *The Annals of Statistics*, 49(2):1166 – 1194.
- Jabin, P.-E., Mellet, A., and Molina-Fructuoso, M. (2021). Local regularity result for an optimal transportation problem with rough measures in the plane. *Journal of Functional Analysis*.
- Jordan, R., Kinderlehrer, D., and Otto, F. (1998). The variational formulation of the fokker–planck equation. *SIAM Journal on Mathematical Analysis*, 29(1):1–17.
- Jourani, A., Thibault, L., and Zagrodny, D. (2012). $C^{1,\omega(\cdot)}$ -regularity and Lipschitz-like properties of subdifferential. *Proceedings of the London Mathematical Society*, 105(1):189–223.
- Kantorovich, L. (1942). On the translocation of masses. *Dokl. Acad. Nauk. USSR*, 37(7-8).
- Kantorovich, L. and Rubinstein, G. S. (1958). On a space of totally additive functions. *Vestnik Leningrad. Univ*, 13:52–59.
- Kellerer, H. G. (1984). Duality theorems for marginal problems. *Zeitschrift für Wahrscheinlichkeitstheorie und Verwandte Gebiete*, 67(4):399–432.
- Kim, Y.-H. and Pass, B. (2017). Wasserstein barycenters over Riemannian manifolds. *Advances in Mathematics*, 307:640–683.
- Kitagawa, J., Mérigot, Q., and Thibert, B. (2019). Convergence of a Newton algorithm for semi-discrete optimal transport. *J. Eur. Math. Soc.*, 21(9):2603–2651.
- Klartag, B. (2013). Regularity through convexity in high dimensions. *Lecture notes* (<https://www.weizmann.ac.il/math/klartag/sites/math.klartag/files/uploads/klartag-notes-by-brazitikos.pdf>).
- Klartag, B. (2014). Logarithmically-concave moment measures I. In *Geometric aspects of functional analysis*, volume 2116 of *Lecture Notes in Math.*, pages 231–260. Springer, Cham.
- Knott, M. and Smith, C. S. (1984). On the optimal mapping of distributions. *Journal of Optimization Theory and Applications*, 43(1):39–49.
- Kolesnikov, A. V. and Milman, E. (2017). Brascamp–Lieb-type inequalities on weighted Riemannian manifolds with boundary. *The Journal of Geometric Analysis*, 27(2):1680–1702.
- Kolouri, S. and Rohde, G. K. (2015). Transport-based single frame super resolution of very low resolution face images. In *2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4876–4884.
- Kolouri, S., Tosun, A. B., Ozolek, J. A., and Rohde, G. K. (2016). A continuous linear optimal transport approach for pattern analysis in image datasets. *Pattern Recognition*, 51:453–462.

- Korba, A., Salim, A., Arbel, M., Luise, G., and Gretton, A. (2020). A non-asymptotic analysis for stein variational gradient descent. In Larochelle, H., Ranzato, M., Hadsell, R., Balcan, M., and Lin, H., editors, *Advances in Neural Information Processing Systems*, volume 33, pages 4672–4682. Curran Associates, Inc.
- Kosowsky, J. and Yuille, A. (1994). The invisible hand algorithm: Solving the assignment problem with statistical physics. *Neural Networks*, 7(3):477–490.
- Kuhn, D., Esfahani, P., Nguyen, V., and Shafieezadeh-Abadeh, S. (2019). Wasserstein Distributionally Robust Optimization: Theory and Applications in Machine Learning. *Operations Research & Management Science in the Age of Analytics*, pages 130–166.
- Le Gouic, T. and Loubes, J.-M. (2017). Existence and consistency of Wasserstein barycenters. *Probability Theory and Related Fields*, 168(3):901–917.
- Le Gouic, T., Paris, Q., Rigollet, P., and Stromme, A. (2022). Fast convergence of empirical barycenters in Alexandrov spaces and the Wasserstein space. *J. Eur. Math. Soc.*
- Le Peutrec, D. (2017). On Witten Laplacians and Brascamp–Lieb’s Inequality on Manifolds with Boundary. *Integral Equations and Operator Theory*, 87(3):411 – 434.
- Leclaire, A. and Rabin, J. (2020). A Stochastic Multi-layer Algorithm for Semi-Discrete Optimal Transport with Applications to Texture Synthesis and Style Transfer. *Journal of Mathematical Imaging and Vision*.
- Lee, D. D. and Seung, H. S. (1999). Learning the parts of objects by non-negative matrix factorization. *Nature*, 401(6755):788–791.
- Leindler, L. (1972). On a Certain Converse of Hölder’s Inequality. *Acta Scientiarum Mathematicarum*, 33:217–233.
- Lellmann, J., Lorenz, D. A., Schönlieb, C., and Valkonen, T. (2014). Imaging with kantorovich–rubinstein discrepancy. *SIAM Journal on Imaging Sciences*, 7(4):2833–2859.
- Leoni, G. (2009). *A First Course in Sobolev Spaces*. Graduate Studies in Mathematics. 105. American Mathematical Society (AMS).
- Lévy, B. (2015). A numerical algorithm for L2 semi-discrete optimal transport in 3D. *ESAIM: Mathematical Modelling and Numerical Analysis*, 49(6):1693–1715.
- Lian, X., Jain, K., Truszkowski, J., Poupart, P., and Yu, Y. (2020). Unsupervised Multilingual Alignment using Wasserstein Barycenter. In Bessiere, C., editor, *Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence, IJCAI-20*, pages 3702–3708. International Joint Conferences on Artificial Intelligence Organization. Main track.
- Lloyd, S. (1982). Least squares quantization in PCM. *IEEE Transactions on Information Theory*, 28(2):129–137.
- Luise, G., Salzo, S., Pontil, M., and Ciliberto, C. (2019). Sinkhorn Barycenters with Free Support via Frank-Wolfe Algorithm. In Wallach, H., Larochelle, H., Beygelzimer, A., d’Alché-Buc, F., Fox, E., and Garnett, R., editors, *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc.

- Lusternik, L. (1935). Die Brunn-Minkowskische Ungleichung für beliebige messbare Mengen. *Dokl. Akad. Nauk SSSR*, 3:55–58.
- Léonard, C. (2012). From the Schrödinger problem to the Monge-Kantorovich problem. *Journal of Functional Analysis*, 262.
- Léonard, C. (2014). A survey of the Schrödinger problem and some of its connections with optimal transport. *Discrete and Continuous Dynamical Systems*, 34(4):1533–1574.
- Mahajan, M., Nimbhorkar, P., and Varadarajan, K. (2012). The planar k-means problem is NP-hard. *Theoretical Computer Science*, 442:13–21. Special Issue on the Workshop on Algorithms and Computation (WALCOM 2009).
- Mallat, S. (2009). *A Wavelet Tour of Signal Processing (Third Edition)*. Academic Press, Boston, third edition edition.
- Matheron, G. (1978). La Formule De Steiner Pour Les Érosions. *Journal of Applied Probability*, 15(1):126–135.
- Maurey, B. (1991). Some deviation inequalities. *Geometric & Functional Analysis GAFA*, 1(2):188–197.
- Maurey, B. (2003-2004). Inégalité de Brunn-Minkowski-Lusternik, et autres inégalités géométriques et fonctionnelles. *Séminaire Bourbaki*, 46:95–114.
- Maury, B., Roudneff-Chupin, A., and Santambrogio, F. (2010). A macroscopic crowd motion model of gradient flow type. *Mathematical Models and Methods in Applied Sciences*, 20(10):1787–1821.
- McCann, R. J. (1994). *A convexity theory for interacting gases and equilibrium crystals*. PhD thesis, Princeton University.
- McCann, R. J. (1997). A Convexity Principle for Interacting Gases. *Advances in Mathematics*, 128(1):153–179.
- McCann, R. J. (2001). Polar factorization of maps on Riemannian manifolds. *Geometric & Functional Analysis GAFA*, 11(3):589–608.
- Mena, G. and Niles-Weed, J. (2019). Statistical bounds for entropic optimal transport: sample complexity and the central limit theorem. In Wallach, H., Larochelle, H., Beygelzimer, A., d'Alché-Buc, F., Fox, E., and Garnett, R., editors, *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc.
- Méridot, Q. (2011). A multiscale approach to optimal transport. In *Computer Graphics Forum*, volume 30(5), pages 1583–1592. Wiley Online Library.
- Méridot, Q., Delalande, A., and Chazal, F. (2020). Quantitative stability of optimal transport maps and linearization of the 2-Wasserstein space. In *Proceedings of the Twenty Third International Conference on Artificial Intelligence and Statistics*, volume 108, pages 3186–3196.
- Mikami, T. (2004). Monge’s problem with a quadratic cost by the zero-noise limit of h-path processes. *Probability Theory and Related Fields*, 129(2):245–260.
- Minkowski, H. (1896). *Geometrie der Zahlen*. Leipzig: Teubner.

- Mirebeau, J.-M. (2015). Discretization of the 3D Monge-Ampère operator, between wide stencils and power diagrams. *ESAIM: M2AN*, 49(5):1511–1523.
- Monge, G. (1781). Mémoire sur la théorie des déblais et des remblais. *Histoire de l'Académie Royale des Sciences de Paris, avec les Mémoires de Mathématique et de Physique pour la même année*, pages 666–704.
- Mérogot, Q. and Thibert, B. (2021). Chapter 2 - Optimal transport: discretization and algorithms. In Bonito, A. and Nocketto, R. H., editors, *Geometric Partial Differential Equations - Part II*, volume 22 of *Handbook of Numerical Analysis*, pages 133–212. Elsevier.
- Nesterov, Y. (2014). *Introductory Lectures on Convex Optimization: A Basic Course*. Springer Publishing Company, Incorporated, 1 edition.
- Nielsen, F. and Sun, K. (2019). *Clustering in Hilbert's Projective Geometry: The Case Studies of the Probability Simplex and the Elliptope of Correlation Matrices*, pages 297–331. Springer International Publishing, Cham.
- Nocketto, R. H. and Zhang, W. (2019). Pointwise rates of convergence for the Oliker–Prussner method for the Monge–Ampère equation. *Numerische Mathematik*, 141(1):253–288.
- Nutz, M. (2022). Introduction to Entropic Optimal Transport. *Lecture notes (http://www.math.columbia.edu/~mnutz/docs/EOT_lecture_notes.pdf)*.
- Nutz, M. and Wiesel, J. (2021). Entropic optimal transport: convergence of potentials. *Probability Theory and Related Fields*.
- Oliker, V. (2003). Mathematical aspects of design of beam shaping surfaces in geometrical optics. In Kirkilionis, M., Krömker, S., Rannacher, R., and Tomi, F., editors, *Trends in Nonlinear Analysis*, pages 193–224, Berlin, Heidelberg. Springer Berlin Heidelberg.
- Oliker, V. I. and Prussner, L. D. (1989). On the numerical solution of the equation $\frac{\partial^2 z}{\partial x^2} \frac{\partial^2 z}{\partial y^2} - \left(\frac{\partial^2 z}{\partial x \partial y} \right) = f$ and its discretizations, i. *Numerische Mathematik*, 54(3):271–293.
- Otto, F. (1998). Dynamics of Labyrinthine Pattern Formation in Magnetic Fluids: A Mean-Field Theory. *Archive for Rational Mechanics and Analysis*, 141(1):63–103.
- Otto, F. (2001). The geometry of dissipative evolution equations: the porous medium equation. *Communications in Partial Differential Equations*, 26:101–174.
- Pal, S. (2019). On the difference between entropic cost and the optimal transport cost. *arXiv: Probability*.
- Panaretos, V. M. and Zemel, Y. (2020). *An Invitation to Statistics in Wasserstein Space*. SpringerBriefs in Probability and Mathematical Statistics. Springer Cham.
- Park, S. and Thorpe, M. (2018). Representing and Learning High Dimensional Data with the Optimal Transport Map from a Probabilistic Viewpoint. In *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7864–7872.

- Pass, B. (2013). Optimal transportation with infinitely many marginals. *Journal of Functional Analysis*, 264(4):947–963.
- Peyré, G. and Cuturi, M. (2019). Computational Optimal Transport. *Foundations and Trends in Machine Learning*, 11(5-6):355–607.
- Pitié, F., Kokaram, A. C., and Dahyot, R. (2007). Automated colour grading using colour distribution transfer. *Computer Vision and Image Understanding*, 107(1):123–137. Special issue on color image processing.
- Prékopa, A. (1971). Logarithmic concave measures with application to stochastic programming. *Acta Scientiarum Mathematicarum*, 32:301–316.
- Prékopa, A. (1973). On logarithmic concave measures and functions. *Acta Scientiarum Mathematicarum*, 34:335–343.
- Rabin, J. and Papadakis, N. (2015). Convex Color Image Segmentation with Optimal Transport Distances. In *International Conference on Scale Space and Variational Methods in Computer Vision (SSVM’15)*, Image Processing, Computer Vision, Pattern Recognition, and Graphics, pages 256–269, Lège Cap Ferret, France.
- Rabin, J., Peyré, G., Delon, J., and Marc, B. (2011). Wasserstein Barycenter and its Application to Texture Mixing. In *SSVM’11*, pages 435–446, Israel. Springer.
- Ramdas, A., Garcia, N., and Cuturi, M. (2015). On Wasserstein Two Sample Testing and Related Families of Nonparametric Tests. *Entropy*, 19.
- Rao, R., Moscovich, A., and Singer, A. (2020). Wasserstein K-Means for Clustering Tomographic Projections.
- Rolet, A., Cuturi, M., and Peyré, G. (2016). Fast Dictionary Learning with a Smoothed Wasserstein Loss. In Gretton, A. and Robert, C. C., editors, *Proceedings of the 19th International Conference on Artificial Intelligence and Statistics*, volume 51 of *Proceedings of Machine Learning Research*, pages 630–638, Cadiz, Spain. PMLR.
- Rubner, Y., Tomasi, C., and Guibas, L. J. (2000). The Earth Mover’s Distance as a Metric for Image Retrieval. *International Journal of Computer Vision*, 40(2):99–121.
- Rüschendorf, L. and Rachev, S. (1990). A characterization of random variables with minimum l2-distance. *Journal of Multivariate Analysis*, 32(1):48–54.
- Samaria, F. and Harter, A. (1994). Parameterisation of a stochastic model for human face identification. In *Proceedings of 1994 IEEE Workshop on Applications of Computer Vision*, pages 138–142.
- Sandler, R. and Lindenbaum, M. (2011). Nonnegative Matrix Factorization with Earth Mover’s Distance Metric for Image Analysis. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 33(8):1590–1602.
- Santambrogio, F. (2015). Optimal transport for applied mathematicians. *Birkhäuser, NY*, 55:58–63.
- Santambrogio, F. (2017). {Euclidean, metric, and Wasserstein} gradient flows: an overview. *Bulletin of Mathematical Sciences*, 7(1):87–154.

- Santambrogio, F. and Wang, X.-J. (2016). Convexity of the support of the displacement interpolation: Counterexamples. *Applied Mathematics Letters*, 58:152–158.
- Schmitz, M. A., Heitz, M., Bonneel, N., Ngolè, F., Coeurjolly, D., Cuturi, M., Peyré, G., and Starck, J.-L. (2018). Wasserstein Dictionary Learning: Optimal Transport-Based Unsupervised Nonlinear Dictionary Learning. *SIAM Journal on Imaging Sciences*, 11(1):643–678.
- Schmitzer, B. (2019). Stabilized sparse scaling algorithms for entropy regularized transport problems. *SIAM J. Sci. Comput.*, 41(3):A1443–A1481.
- Schneider, R. (2013). *Convex Bodies: The Brunn–Minkowski Theory*. Encyclopedia of Mathematics and its Applications. Cambridge University Press, 2 edition.
- Schrödinger, E. (1931). Über die umkehrung der naturgesetze. *Angewandte Chemie*, 44(30):636–636.
- Smith, C. S. and Knott, M. (1987). Note on the optimal transportation of distributions. *Journal of Optimization Theory and Applications*, 52(2):323–329.
- Solomon, J., de Goes, F., Peyré, G., Cuturi, M., Butscher, A., Nguyen, A., Du, T., and Guibas, L. (2015). Convolutional Wasserstein Distances: Efficient Optimal Transportation on Geometric Domains. *ACM Trans. Graph.*, 34(4).
- Solomon, J., Rustamov, R., Guibas, L., and Butscher, A. (2014). Wasserstein propagation for semi-supervised learning. In Xing, E. P. and Jebara, T., editors, *Proceedings of the 31st International Conference on Machine Learning*, volume 32 of *Proceedings of Machine Learning Research*, pages 306–314, Beijing, China. PMLR.
- Srivastava, S., Li, C., and Dunson, D. B. (2018). Scalable Bayes via Barycenter in Wasserstein Space. *Journal of Machine Learning Research*, 19(8):1–35.
- Staib, M. and Jegelka, S. (2017). Wasserstein k-means++ for Cloud Regime Histogram Clustering. In *Climate Informatics*.
- Trudinger, N. S. and Wang, X.-J. (2001). On the Monge mass transfer problem. *Calculus of Variations and Partial Differential Equations*, 13(1):19–31.
- Varadarajan, V. S. (1958). On the convergence of sample probability distributions. *Sankhyā: The Indian Journal of Statistics (1933-1960)*, 19(1/2):23–26.
- Vershynin, R. (2018). *High-Dimensional Probability: An Introduction with Applications in Data Science*. Cambridge Series in Statistical and Probabilistic Mathematics. Cambridge University Press.
- Villani, C. (2003). *Topics in Optimal Transportation*. Graduate studies in mathematics. American Mathematical Society.
- Villani, C. (2008). *Optimal transport: old and new*, volume 338. Springer Science & Business Media.
- Virtanen, P., Gommers, R., Oliphant, T. E., Haberland, M., Reddy, T., Cournapeau, D., Burovski, E., Peterson, P., Weckesser, W., Bright, J., van der Walt, S. J., Brett, M., Wilson, J., Millman, K. J., Mayorov, N., Nelson, A. R. J., Jones, E., Kern, R.,

- Larson, E., Carey, C. J., Polat, İ., Feng, Y., Moore, E. W., VanderPlas, J., Laxalde, D., Perktold, J., Cimrman, R., Henriksen, I., Quintero, E. A., Harris, C. R., Archibald, A. M., Ribeiro, A. H., Pedregosa, F., van Mulbregt, P., and SciPy 1.0 Contributors (2020). SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python. *Nature Methods*, 17:261–272.
- Wang, W., Slepčev, D., Basu, S., Ozolek, J. A., and Rohde, G. K. (2013). A linear optimal transportation framework for quantifying and visualizing variations in sets of images. *Int. J. Comput. Vision*, 101(2):254–269.
- Weed, J. (2018). An explicit analysis of the entropic penalty in linear programming. In Bubeck, S., Perchet, V., and Rigollet, P., editors, *Conference On Learning Theory, COLT 2018, Stockholm, Sweden, 6-9 July 2018*, volume 75 of *Proceedings of Machine Learning Research*, pages 1841–1855. PMLR.
- Weed, J. and Bach, F. (2019). Sharp asymptotic and finite-sample rates of convergence of empirical measures in Wasserstein distance. *Bernoulli*, 25(4 A):2620–2648.
- Weed, J. and Berthet, Q. (2019). Estimation of smooth densities in Wasserstein distance. In Beygelzimer, A. and Hsu, D., editors, *Proceedings of the Thirty-Second Conference on Learning Theory*, volume 99 of *Proceedings of Machine Learning Research*, pages 3118–3119. PMLR.
- Yan, M. (2014). Extension of Convex Function. *Journal of Convex Analysis*, 21(4):965–987.
- Yang, C., Delcher, C., Shenkman, E., and Ranka, S. (2018). Clustering Inter-Arrival Time of Health Care Encounters for High Utilizers. In *2018 IEEE 20th International Conference on e-Health Networking, Applications and Services (Healthcom)*, pages 1–6.
- Ye, J., Wu, P., Wang, J. Z., and Li, J. (2017). Fast Discrete Distribution Clustering Using Wasserstein Barycenter With Sparse Support. *IEEE Transactions on Signal Processing*, 65(9):2317–2332.

Titre: Stabilité quantitative en transport optimal quadratique

Mots clés: Transport optimal, Stabilité quantitative, Analyse convexe, Optimisation, Géométrie métrique, Apprentissage automatique

Résumé: L'étude du problème de transport optimal permet de définir des métriques sur des espaces de mesures de probabilité qui ont de fortes interprétations physiques et géométriques. Ce problème et ses métriques (appelées distances de Wasserstein) ont trouvé de nombreuses applications, allant de la physique quantique, la mécanique des fluides, la conception optique, l'économie et les statistiques, à l'apprentissage automatique. Le profond ancrage physique du problème de transport optimal appelle à l'étude de son caractère bien posé. Si l'existence et l'unicité des solutions du problème de transport optimal ont été largement étudiées dans des travaux antérieurs, l'étude de sa stabilité est moins avancée. Des résultats de stabilité généraux et abstraits garantissent que les solutions de transport optimal dépendent continûment des données du problème qui les définit. Cependant, ces résultats ne sont presque jamais quantitatifs, ce qui est problématique dans les applications, où les données sont souvent disponibles de manière approximative. Cette thèse vise à combler cette lacune. En se concentrant dans un contexte euclidien sur le problème

de transport optimal quadratique, nous donnons dans une première partie des estimations de la forte convexité du problème dual de Kantorovich en nous appuyant sur des inégalités géométriques et fonctionnelles bien connues. Nous rassemblons ensuite dans une deuxième partie des estimations quantitatives de la stabilité des applications de transport optimal par rapport à leur mesure cible, des barycentres de Wasserstein par rapport à leurs mesures marginales et des potentiels de Schrödinger par rapport au paramètre de température dans le transport optimal entropique semi-discret. Ces estimations suggèrent toutes des applications naturelles en transport optimal numérique et statistique. Elles donnent également de nouvelles indications sur la plongeabilité de l'espace Wasserstein-2 dans un espace de Hilbert avec une distorsion de la métrique contrôlée. Dans une dernière partie, nous exploitons cette dernière idée dans des applications d'apprentissage automatique et proposons des approches pour résoudre approximativement des problèmes de K -moyennes et d'apprentissage de dictionnaire dans l'espace Wasserstein-2.

Title: Quantitative Stability in Quadratic Optimal Transport

Keywords: Optimal Transport, Quantitative Stability, Convex Analysis, Optimization, Metric Geometry, Machine Learning.

Abstract: The study of the optimal transport problem allows to define metrics on spaces of probability measures that come with strong geometrical and physical interpretations. This problem and its metrics (called Wasserstein distances) have found many applications, ranging from quantum physics, fluid dynamics, optics design, economics and statistics, to machine learning. The profound physical rooting of the optimal transport problem calls for the study of its well-posedness. While the existence and uniqueness of solutions to the optimal transport problem have been extensively studied in previous works, the investigation of its stability is less advanced. General and abstract stability results ensure that optimal transport solutions change continuously with the problem data. Yet these results are hardly ever quantitative, which is problematic in applications, where the data is often approximated. This thesis works towards closing this gap. Focus-

ing in a Euclidean context on the quadratic optimal transport problem, we derive in a first part strong convexity estimates for the Kantorovich dual problem relying on well-known geometric and functional inequalities. We then collect in a second part quantitative stability estimates for optimal transport maps with respect to their target measure, for Wasserstein barycenters with respect to their marginals and for Schrödinger potentials with respect to the temperature parameter in semi-discrete entropic optimal transport. These estimates all suggest natural applications in computational and statistical optimal transport. They also give new insights on the embeddability of the 2-Wasserstein space in a Hilbert space with a controlled distorsion. In a last part, we leverage this last idea in machine learning applications and propose approaches to approximately solve K -means and dictionary learning problems in the 2-Wasserstein space.