

Numerical methods for matching for teams and Wasserstein barycenters

G. Carlier ^{*}, A. Oberman [†], E. Oudet [‡]

May 16, 2014

Abstract

Equilibrium multi-population matching (matching for teams) is a problem from mathematical economics which is related to multi-marginal optimal transport. A special but important case is the Wasserstein barycenter problem, which has applications in image processing and statistics. Two algorithms are presented: a linear programming algorithm and an efficient nonsmooth optimization algorithm, which applies in the case of the Wasserstein barycenters. The measures are approximated by discrete measures: convergence of the approximation is proved. Numerical results are presented which illustrate the efficiency of the algorithms.

Keywords: matching for teams, Wasserstein barycenters, duality, linear programming, numerical methods for nonsmooth convex minimization.

1 Introduction

Optimal transport theory has received a lot of attention in the last decades and is now recognized as a powerful tool in PDEs, geometry, and functional inequalities (for which we refer to the monographs of Villani [Vil03]-[Vil09]). Given two Borel probability measures μ_1, μ_2 , on metric spaces X_1 and X_2 , respectively, and a cost function $c \in C(X_1 \times X_2, \mathbb{R})$, the Monge-Kantorovich optimal transport problem consists in finding the cheapest way to transport μ_1 to μ_2 for the cost c :

$$W_c(\mu_1, \mu_2) := \inf_{\gamma \in \Pi(\mu_1, \mu_2)} \int_{X \times X} c(x_1, x_2) \gamma(dx_1, dx_2) \quad (\text{MK})$$

where $\Pi(\mu_1, \mu_2)$ denotes the set of transport plans between μ_1 and μ_2 , i.e. the set of probability measures on $X_1 \times X_2$ having μ_1, μ_2 , respectively as marginals.

^{*}CEREMADE, UMR CNRS 7534, Université Paris IX Dauphine, Pl. de Lattre de Tassigny, 75775 Paris Cedex 16, FRANCE, carlier@ceremade.dauphine.fr

[†]Department of Mathematics and Statistics, McGill University, 805 Sherbrooke Street West, Montreal, CANADA, adam.oberman@mcgill.ca

[‡]Laboratoire Jean Kuntzmann, Université Joseph Fourier, Tour IRMA, BP 53 51, rue des Mathématiques F-38041 Grenoble Cedex 9, FRANCE, edouard.oudet@imag.fr.

Since this problem is of linear programming type, under very mild assumptions (e.g. when X_1 and X_2 are compact), the least transport cost $W_c(\mu_1, \mu_2)$ admits a dual expression given by the Kantorovich duality formula:

$$W_c(\mu_1, \mu_2) := \sup_{\varphi \in C(X_1)} \left\{ \int_{X_1} \varphi(x_1) \mu_1(dx_1) + \int_{X_2} \varphi^c(x_2) \mu_2(dx_2) \right\} \quad (1.1)$$

where φ^c denotes the c -transform of φ :

$$\varphi^c(x_2) := \min_{x_1 \in X_1} \{c(x_1, x_2) - \varphi(x_1)\}.$$

A particularly important example is the quadratic case where $X_1 = X_2 = \mathbb{R}^d$, μ_1 and μ_2 have finite second moments, and $c(x_1, x_2) = |x_1 - x_2|^2$. This case was first solved by Brenier [Bre91], who proved that whenever μ_1 is absolutely continuous, there is a unique optimal transport plan that is given by the gradient of a convex potential. This important result relates optimal transport to Monge-Ampère equations. We refer to [BFO14] and the references therein for numerical methods for optimal transport based on the Monge-Ampère equation.

More generally, costs given by distances or convex power of distances are important because they lead to the so-called Wasserstein distances. More precisely, when $X_1 = X_2$ (a metric space with distance d) and $c(x_1, x_2) = d(x_1, x_2)^p$ for some $p \geq 1$, the value $W_c(\mu_1, \mu_2)$ in (MK) is the p -power of the so-called p -Wasserstein distance $W_p(\mu_1, \mu_2)$ between μ_1 and μ_2 :

$$W_p(\mu_1, \mu_2) := \left(\inf_{\gamma \in \Pi(\mu_1, \mu_2)} \int_{X \times X} d(x_1, x_2)^p \gamma(dx_1, dx_2) \right)^{1/p}.$$

In the present article, we are interested in solving numerically the following variant of the optimal transport problem which allows for more than two marginals. Given (compact metric, say) spaces X_1, \dots, X_I , equipped with Borel probability measures $(\mu_1, \dots, \mu_I) \in \mathcal{P}(X_1) \times \dots \times \mathcal{P}(X_I)$, a (compact metric) space Z , and cost functions $c_i \in C(X_i \times Z, \mathbb{R})$, we look for a probability measure ν on Z solving:

$$\inf_{\nu \in \mathcal{P}(Z)} J(\nu) := \sum_{i=1}^I W_{c_i}(\mu_i, \nu). \quad (1.2)$$

This problem was introduced in Carlier and Ekeland [CE10] in the framework of multi-population matching equilibrium; we will shortly recall in section 2 the economic interpretation of (1.2). Problem (1.2) is also a special case of multi-marginal optimal transport (a variant of (MK) where more than two marginals are prescribed). Multi-marginal optimal transport is currently an active research field: compared to (two marginals) optimal transport, there are fewer theoretical results, and the complexity of general multi-marginal optimal transport problems typically increases exponentially in the number of marginals. Regarding the rapidly developing theory of multi-marginal optimal transport, we refer the reader to the recent papers by Pass [Pas12a], [Pas12b], by Ghoussoub and

coauthors [GM13], [MG14] and the references therein for costs with special symmetry properties, motivated in particular by challenging computational issues in density functional theory in quantum physics.

We now discuss a special, but important case of (1.2) which has a clear geometric interpretation. Let all the X_i 's and Z coincide with \mathbb{R}^d , the measures μ_i have finite second moments, and the costs be quadratic (i.e. $c_i(x_i, z) := \lambda_i |x_i - z|^2$ for some weights $\lambda_i > 0$, summing to 1 without loss of generality). In this case (1.2) takes the form:

$$\inf_{\nu \in \mathcal{P}(\mathbb{R}^d)} J(\nu) := \sum_{i=1}^I \lambda_i W_2^2(\mu_i, \nu) \quad (1.3)$$

where W_2 denotes the 2-Wasserstein distance. In analogy with the Euclidean case, a solution to (1.3) will be called a Wasserstein barycenter of the measures μ_i with weights λ_i . Properties of Wasserstein barycenters were studied by Agueh and Carlier [AC11]. Wasserstein barycenters interpolate between the measures μ_i ; the idea of interpolating between points of a metric space by minimizing some weighted sum of squared distances goes back to the notion of Fréchet mean. The case $I = 2$ is well-known. Letting the weights $(\lambda, 1 - \lambda)$ vary, one obtains the classical notion of McCann's interpolation [McC97] between two probability measures. This interpolating curve is also a geodesic for W_2 , and in their seminal paper [BB00] on the dynamic formulation of optimal transport, Benamou and Brenier gave a numerical scheme to compute this geodesic. Finding barycenters between more than three measures is more complicated (barycenters are not associative as soon as $d \geq 2$). From a Partial Differential Equations problem, this problem requires to solve a system of Monge-Ampère equations, see (4.9)-(4.8) below. Interestingly, the Wasserstein barycenter problem recently found natural applications in image processing, see Peyré et al. [RPDB12] and statistics, see Bigot and Klein [BK12]. Of course, there are lots of variants of the interpolating scheme given by the W_2 -barycenter problem (1.3) and in particular one can replace W_2 by W_p for some $p \geq 1$ or even mix different powers of the distance. Slightly abusing the terminology, we will sometimes refer to barycenters even for these variants and even for the general form of the problem (1.2).

In the discrete setting, the transportation problem is classical. In fact, this problem motivated the historical development of optimization, by Kantorovich in 1939, working on Soviet railway transportation, and in the 1940's by Hitchcock [Sch03, Ch 21]. The "assignment problem" arises in case of integer values weights, it is a standard combinatorial optimization problem which can be solved by the Hungarian algorithm) [Sch03, Ch 21]. More generally, the "transportation problem" is a linear programming problem which arises when the weights are real-valued, it can be solved by the Hitchcock algorithm, [Sch03, Ch 21], or by modern commercial general linear programming software. Returning to the problem with continuous measures, it is natural to approximate the measure by weighted sums of delta measures. In theory, the resulting problem can be solved

using linear programming. However, the number of variables in the linear programming problem is quadratic in the number of variables used to represent the measures. In the discrete setting, current optimization algorithms are limited to approximately several thousands of variables for each of the measures. This problem size corresponds to a fairly coarse approximation of a two dimensional continuous measure. In special cases, or using specific approximations, improvements are available, see [PPO14] for quadratic costs, and for more references. For example, if each measure is represented by, for example, $40^2 = 1600$ variables, the linear program has $40^4 = 2\,560\,000$ which is near the limitations of linear programming algorithms (we performed experiments using CVX [GB10] and calling several academic and commercial optimization packages). Enlarging the resolution of the measures quickly overwhelms the capabilities of the algorithms.

The problem (1.2) is even more challenging, since it involves multiple marginals and an additional unknown measure. A reasonable goal is to resolve each of the measures in two dimensions, with (on the order of, say) 1600 variables. Since we wish to allow for measures which are separated, this may require a full grid of size up to 200^2 . Resolving the barycenter measure on the full grid generally leads to an intractable problem. Our main contributions regarding numerical schemes for the general problem (1.2) or the particular case of Wasserstein barycenters (1.3) are as follows:

- We give a simple linear programming reformulation of (1.2) in [subsection 2.3](#) whose size is proportional to the number of marginals. Together with a localization result that bounds the support of the unknown barycenter in [subsection 2.2](#), one then obtains a tractable problem. We discretize the problem to arrive at a finite dimensional linear programming problem in [subsection 2.4](#). We prove convergence, in the sense of weak convergence of measures, in [subsection 2.5](#).
- Numerical results are presented in [section 3](#). These illustrate the validity of this linear programming approach. Barycenter problems with different costs are solved, as well as a matching for teams problem.
- The second algorithm which is specialized to the case of Wasserstein barycenter measures (1.3), is described and implemented in [section 4](#). This problem uses the dual formulation of the problem explained in [section 4](#), and special features of the quadratic cost. The efficient nonsmooth optimization algorithm is described in [subsection 4.3](#). Large size computational examples are presented (on grids of size 200^2 , and for measures resolved with 15000 points). The examples include barycenter measures using up to five measures, and an example in texture synthesis in [subsection 4.4](#).

2 Matching for teams and approximation

In this section, following [CE10], we first derive the generalized barycenter problem (1.2) as an equivalent reformulation of an equilibrium problem for multi-population matching arising in economics. Next, we study localization of the barycenter measure. Then, we present an infinite dimensional linear programming reformulation of (2.5). This is followed by a discretization of the measures, which results in a finite dimensional linear programming problem that is tractable for moderate problem sizes. Finally, we address stability issues (in the sense of weak convergence of measures) when one approximates the measures μ_i by some (discrete) measures.

2.1 Variational characterization of matching equilibria

The model of Carlier and Ekeland [CE10] deals with the equilibrium of a market for a quality good (e.g. house, school, hospitals, ...). Producing the good requires assembling a team consisting of a buyer and a set of producers. For instance, in the case of houses, the producers could be a plumber, an electrician and a mason. The quality good has a range of feasible qualities (location, surface, number of rooms, facilities etc...), denoted by Z which we assume to be a compact metric space.

Each of the different populations (buyers, plumbers, electricians, masons...) is indexed by $i \in \{1, \dots, I\}$. The agents in each population are heterogeneous, characterized by a certain type which affects their (quality dependent) cost function. For example, some masons are used to working with lower quality bricks, while other work with luxury stone, some electricians leave quite far from the location of the house they work on, consumers differ in their tastes... To be precise, for each population i , we are given a compact metric space of types, X_i , and a continuous cost function $c_i \in C(X_i \times Z, \mathbb{R})$; with the interpretation that $c_i(x_i, z)$ is the cost for an agent of population i with type x_i to work in a team that produces good z . The distribution of type x_i in population i is known and given by some Borel probability measure $\mu_i \in \mathcal{P}(X_i)$.

The goal is to find an equilibrium production line $\nu \in \mathcal{P}(Z)$ (together with a price system) which clears both the quality good and the labor market. The equilibrium is described below, and as we shall see, it corresponds to the solution of the (generalized) barycenter measure problem (2.5). In this setting, one looks in particular for an equilibrium system of monetary transfers (paid by the buyer to the producers). A system of transfers is a collection of continuous functions $\varphi_1, \dots, \varphi_I: Z \rightarrow \mathbb{R}$ where $\varphi_i(z)$ is the amount paid to i by the other members of the team for producing z . An obvious equilibrium requirement is that teams are self-financed i.e.

$$\sum_{i=1}^I \varphi_i(z) = 0, \forall z \in Z. \quad (2.1)$$

Given transfers, $\varphi_1, \dots, \varphi_I$, an agent from population i with type $x_i \in X_i$, gets

a net minimal cost given by the so-called c_i -transform of φ_i :

$$\varphi_i^{c_i}(x_i) := \min_{z \in Z} \{c_i(x_i, z) - \varphi_i(z)\}. \quad (2.2)$$

By construction, $\varphi_i^{c_i}(x_i) + \varphi_i(z) \leq c_i(x_i, z)$, and since agents are rational, they choose cost minimizing qualities, i.e. a $z \in Z$ such that

$$\varphi_i^{c_i}(x_i) + \varphi_i(z) = c_i(x_i, z). \quad (2.3)$$

The final unknown is a collection of plans, $\gamma_i \in \mathcal{P}(X_i \times Z)$, such that $\gamma_i(A_i \times A)$ represents the probability that an agent in population i has a type in A_i , and belongs to a team that produces a quality in A . At equilibrium, the first marginal of γ_i should be μ_i (this is equilibrium on the i -th labor market) and the second marginal of γ_i should not depend on i (this is equilibrium on the quality good market), this common marginal represents the equilibrium quality line. An equilibrium can then be formally defined. It consists of a transfer system $(\varphi_1, \dots, \varphi_I) \in C(Z, \mathbb{R})^I$, probability measures $\gamma_i \in \mathcal{P}(X_i \times Z)$, and a probability measure $\nu \in \mathcal{P}(Z)$, such that

- teams are self-financed i.e. (2.1) holds,
- $\gamma_i \in \Pi(\mu_i, \nu)$ for $i = 1, \dots, I$ (equilibrium on the labor markets and on the good market),
- (2.3) holds on the support of γ_i for $i = 1, \dots, I$, (i.e. agents choose cost minimizing qualities).

If an equilibrium quality line, ν , was known, then clearly the last two conditions above would imply that the plan γ_i should be optimal for the Monge-Kantorovich problem:

$$W_{c_i}(\mu_i, \nu) := \inf_{\gamma \in \Pi(\mu_i, \nu)} \int_{X_i \times Z} c_i(x_i, z) \gamma(dx_i, dz). \quad (2.4)$$

In fact, it was proved in [CE10] that there is a purely variational characterization of equilibria, which is tightly related to the following convex problem

$$\inf_{\nu \in \mathcal{P}(Z)} J(\nu) := \sum_{i=1}^I W_{c_i}(\mu_i, \nu) \quad (2.5)$$

and its dual (concave maximization) formulation (see [CE10] or section 4 for details on this duality)

$$\sup \left\{ \sum_{i=1}^I \int_{X_i} \varphi_i^{c_i}(x_i) \mu_i(dx_i) : \sum_{i=1}^I \varphi_i = 0 \right\}. \quad (2.6)$$

Theorem 2.1. *$(\varphi_i, \gamma_i, \nu)$ is an equilibrium if and only if:*

- ν solves (2.5),
- the transfers $(\varphi_1, \dots, \varphi_I)$ solve (2.6),
- for $i = 1, \dots, I$, γ_i solves the Monge-Kantorovich problem $W_{c_i}(\mu_i, \nu)$.

2.2 Localization

As noted in [CE10], the minimization problem (2.5), which characterizes equilibrium quality lines, can be reformulated as an optimal transport problem with multi-marginal constraints, as follows. First define the cost

$$c(x) := \min_{z \in Z} \sum_{i=1}^I c_i(x_i, z), \quad (2.7)$$

where $x = (x_1, \dots, x_I)$. Let $T(x) \in Z$ be a measurable selection of the solution of the above minimization, meaning that $T(x) \in Z$ satisfies

$$\sum_{i=1}^I c_i(x_i, T(x)) = c(x).$$

Then consider the multi-marginal problem

$$\inf_{\gamma \in \Pi(\mu_1, \dots, \mu_I)} \int_{X_1 \times \dots \times X_I} c(x) \gamma(dx), \quad (2.8)$$

where $\Pi(\mu_1, \dots, \mu_I)$ denotes the set of probability measures on $X_1 \times \dots \times X_I$ having (μ_1, \dots, μ_I) as marginals. It is not difficult to see that if γ solves (2.8) then $\nu := T_{\#} \gamma$ solves (2.5) (where as usual $T_{\#} \gamma$ denotes the push forward of γ through T , i.e. $T_{\#} \gamma(B) := \gamma(T^{-1}(B))$ for every Borel B).

Conversely, one can relate the minimizers of (2.5) to those of (2.8). Indeed, let ν solve (2.5) and let $\gamma_i \in \Pi(\mu_i, \nu)$ be an optimal plan for $W_{c_i}(\mu_i, \nu)$. Disintegrating γ_i with respect to ν i.e. writing $\gamma_i = \gamma_i^z \otimes \nu$ and defining $\gamma \in \mathcal{P}(X_1 \times \dots \times X_I \times Z)$ by:

$$\gamma := \otimes_{i=1}^I \gamma_i^z \otimes \nu$$

and η as the marginal of γ on the variables (x_1, \dots, x_I) , one easily checks that

- $\eta \in \Pi(\mu_1, \dots, \mu_I)$ solves (2.8),
- on the support of γ , $\text{spt}(\gamma)$, one has

$$\sum_{i=1}^I c_i(x_i, z) = c(x),$$

- the previous relation, together with the fact that ν is the Z -marginal of γ and μ_i its X_i -marginal then imply a useful localization property: the support of the barycenter measure, $\text{spt}(\nu)$, is contained in the set of minimizers of the following problem.

$$\min_{z \in Z} \sum_{i=1}^I c_i(x_i, z) \text{ for some } x_i \in \text{spt } \mu_i, \quad i = 1, \dots, I. \quad (2.9)$$

Since the support of ν is unknown, which causes difficulties in practice, the localization property (2.9) gives a practical method for bounding the unknown support of the barycenter measure. The condition above results in a reduction of the dimensionality of discretized problems since it gives an a priori information on the support of the unknown measure, at the expense of solving an optimization problem. However this optimization problem is decoupled on the domain Z : each point (or neighborhood) can be tested by looping through points (or small neighborhoods) in the domain Z and choices of points in the support sets $\text{spt}(\mu_i)$.

In the case where X_i and Z coincide with some ball of \mathbb{R}^d , and the costs are powers of distance, $c_i(x_i, z) = \lambda_i |x_i - z|^p$ (with $\lambda_i > 0$ and $\sum \lambda_i = 1$, say) for some $p \geq 1$, one can easily derive an information on the unknown support. Indeed, using the optimality condition for the minimization problem (2.9), one deduces that $\text{spt}(\nu)$ is included in the convex hull of the supports of the μ_i 's. If we particularize further to the Wasserstein barycenter case, i.e. to the case $p = 2$, the solution of (2.9) is explicitly given by the barycenter $z = \sum_{i=1}^I \lambda_i x_i$ so that the localization property (2.9) gives the following estimate on the barycenter measure ν :

$$\text{spt}(\nu) \subset \sum_{i=1}^I \lambda_i \text{spt}(\mu_i). \quad (2.10)$$

2.3 Linear programming formulation

Multi-marginals optimal transport problems such as (2.8) are linear programs. For discrete marginals, such problems can in principle be solved exactly by the simplex method. In practice however, the number of variables explodes with the number of marginals, which makes the problem quickly untractable. We shall see below that one may take advantage of the fact that c is not any cost function but has the special structure (2.7). Interestingly, it was already proved by Pass [Pas12a] in the context of multi-marginal optimal transport that such costs are much more well-behaved than arbitrary costs of I variables.

To find a more tractable linear programming reformulation of the matching for teams problem, it is better to go back to the very definition of an equilibrium in terms of couplings and to reformulate problem (2.5) as

$$\inf_{(\gamma_1, \dots, \gamma_I) \in \Pi} \sum_{i=1}^I \int_{X_i \times Z} c_i(x_i, z) \gamma_i(dx_i, dz) \quad (2.11)$$

where Π consists of all measures $(\gamma_1, \dots, \gamma_I) \in \mathcal{P}(X_1 \times Z) \times \dots \times \mathcal{P}(X_I \times Z)$ such that

- the marginal of γ_i on the x_i variable is μ_i i.e.

$$\int_{X_i \times Z} \psi(x_i) \gamma(dx_i, dz) = \int_{X_i} \psi(x_i) \mu_i(dx_i), \quad \forall \psi \in C(X_i), \quad (2.12)$$

- the marginal of γ_i on the z variable does not depend on i :

$$\int_Z \varphi(z) \gamma_1(dx_1, dz) = \dots = \int_{X_I} \varphi(z) \gamma_I(dx_I, dz), \quad \forall \varphi \in C(Z). \quad (2.13)$$

Clearly, if the γ_i 's solve (2.11) then their common marginal $\nu \in \mathcal{P}(Z)$ solves (2.5) and the γ_i 's are optimal for the optimal transport problem $W_{c_i}(\mu_i, \nu)$. In other words, the γ_i 's are equilibrium couplings for the matching for teams problem.

The constraints above being linear, Π is a convex and weakly star compact subset of $\mathcal{P}(X_1 \times Z) \times \dots \times \mathcal{P}(X_I \times Z)$ so that (2.11) admits solutions. Moreover in the case of discrete μ_i 's and ν supported by N points, the number of variables in the linear program (2.11) is linear (and not exponential as in the case of the multi-marginal optimal transport problem) in the number of marginals.

2.4 Discretization

The (a priori) infinite dimensional linear programming problem (2.11) of subsection 2.3 can be discretized as follows. Let $\{S_j^i\}_{j=1}^{N_i}$ be a partition of $\text{spt}(\mu_i)$ and let $\{S_k^0\}_{k=1}^{N_0}$ be a partition of Z (or better, of the support set estimated by the method of subsection 2.2). Approximate the measures by weighted sums of atoms

$$\begin{aligned} \mu_i^A &= \sum_j \mu_i^j \delta_{x_j^i}, \quad \text{for } i = 1, \dots, I, \quad \text{with } \mu_i^j = \mu_i(S_j^i) \\ \nu^A &= \sum_k \nu^k \delta_{z_k^0}, \quad \text{with } \nu^k = \nu(S_k^0) \end{aligned}$$

where x_j^i and z_k^0 are representative points in the regions S_j^i, S_k^0 , respectively. It is well-known, that μ_i^A converges weakly $*$ to μ_i as the diameter of the partition $\{S_j^i\}_{j=1}^{N_i}$ tends to 0. More precisely, denoting by W_1 the 1-Wasserstein distance (which metrizes the weak $*$ topology on probability measures):

$$W_1(\mu_i^A, \mu_i) \leq \max_{j=1, \dots, N_i} \text{diam}(S_j^i). \quad (2.14)$$

Inserting the approximation defined above into the linear programming problem (2.11, 2.12, 2.13) results in the following finite dimensional linear programming problem

$$\text{minimize} \quad \sum_{i=1}^I \sum_{j,k} c_i(x_j^i, z_k^0) \gamma_{j,k}^i$$

subject to:

$$\begin{aligned} \sum_k \gamma_{j,k}^i &= \mu_i^j, & \text{for all } k = 1, \dots, N_i, \text{ and for } i = 1, \dots, I \\ \sum_j \gamma_{j,k}^1 &= \dots = \sum_j \gamma_{j,k}^I, & \text{for all } k = 1, \dots, N_0, \end{aligned} \quad (2.15)$$

along with the non-negativity constraints $\gamma_{j,k}^i \geq 0$. The (approximated) barycenter is then $\nu^A = \sum_k \nu^k \delta_{z_k^0}$ where the weight ν^k is given by the common value,

$$\nu^k = \sum_j \gamma_{j,k}^i, \quad \text{for any } i = 1, \dots, I.$$

The linear programming problem above can be implemented in standard software packages. The size of the problem above is as follows. The number of variables is $N_0 \times (N_1 + \dots + N_I)$ (or IN^2 if each $N_0 = N_1 = \dots = N_I = N$). The number of constraints is $(N_0 + 1) \times (N_1 + \dots + N_I) + IN_0$ (that is $I(N^2 + 2N)$ when $N_0 = N_1 = \dots = N_I = N$). The size of this linear programming problem thus scales linearly with the number of marginals, for a given, fixed value of N (contrary to the multi-marginal formulation (2.8)).

2.5 Approximation and convergence

Since in practice, one considers approximation by discrete measures just as in [subsection 2.4](#), we wish now to address the stability of the following convex problem when one replaces the measures μ_i by some discrete approximation

$$\inf_{\nu \in \mathcal{P}(Z)} J(\nu) := \sum_{i=1}^I W_{c_i}(\mu_i, \nu). \quad (2.16)$$

To do so, one has to control the dependence of $W_c(\mu, \nu)$ in its three arguments $(c, \mu, \nu) \in C(X \times Z) \times \mathcal{P}(X) \times \mathcal{P}(Z)$. We shall denote by d_X and d_Z the distances on X and Z , take now $(c, \mu, \nu) \in C(X \times Z) \times \mathcal{P}(X) \times \mathcal{P}(Z)$ and $(\tilde{c}, \tilde{\mu}, \tilde{\nu}) \in C(X \times Z) \times \mathcal{P}(X) \times \mathcal{P}(Z)$ and let ω_X and ω_Z be respectively a modulus of continuity of c and \tilde{c} with respect to x uniform in z and a modulus of continuity of c and \tilde{c} with respect to z uniform in x , that is

$$\max(|c(x, z) - c(x', z)|, |\tilde{c}(x, z) - \tilde{c}(x', z)|) \leq \omega_X(d_X(x, x')), \forall (x, x', z) \in X \times X \times Z$$

and

$$\max(|c(x, z) - c(x, z')|, |\tilde{c}(x, z) - \tilde{c}(x, z')|) \leq \omega_Z(d_Z(z, z')), \forall (x, z, z') \in X \times Z \times Z.$$

Obviously, one has

$$|W_c(\mu, \nu) - W_{\tilde{c}}(\mu, \nu)| \leq \|c - \tilde{c}\|_\infty. \quad (2.17)$$

Let $\varphi \in C(Z)$ be a solution in the Kantorovich dual of $W_{\tilde{c}}(\mu, \nu)$, that is

$$W_{\tilde{c}}(\mu, \nu) = \int_X \varphi^{\tilde{c}} d\mu + \int_Z \varphi d\nu$$

by the Kantorovich duality formula, we have

$$W_c(\tilde{\mu}, \nu) \geq \int_X \varphi^{\tilde{c}} d\tilde{\mu} + \int_Z \varphi d\nu.$$

Hence, for every $\theta \in \Pi(\mu, \tilde{\mu})$, we have

$$W_{\tilde{c}}(\mu, \nu) - W_{\tilde{c}}(\tilde{\mu}, \nu) \leq \int_X \varphi^{\tilde{c}} d(\mu - \tilde{\mu}) = \int_{X \times X} (\varphi^{\tilde{c}}(x) - \varphi^{\tilde{c}}(x')) \theta(dx, dx').$$

we then observe that $\varphi^{\tilde{c}}(x) - \varphi^{\tilde{c}}(x') \leq \omega_X(d_X(x, x'))$ so that

$$W_{\tilde{c}}(\mu, \nu) - W_{\tilde{c}}(\tilde{\mu}, \nu) \leq W_{\omega_X}(\mu, \tilde{\mu}) := \inf_{\theta \in \Pi(\mu, \tilde{\mu})} \int_{X \times X} \omega_X(d_X(x, x')) \theta(dx, dx'). \quad (2.18)$$

Similarly

$$W_{\tilde{c}}(\tilde{\mu}, \nu) - W_{\tilde{c}}(\tilde{\mu}, \tilde{\nu}) \leq W_{\omega_Z}(\nu, \tilde{\nu}) := \inf_{\eta \in \Pi(\nu, \tilde{\nu})} \int_{Z \times Z} \omega_Z(d_Z(z, z')) \eta(dz, dz'). \quad (2.19)$$

Putting everything together, we get

$$|W_c(\mu, \nu) - W_{\tilde{c}}(\tilde{\mu}, \tilde{\nu})| \leq \|c - \tilde{c}\|_{\infty} + W_{\omega_X}(\mu, \tilde{\mu}) + W_{\omega_Z}(\nu, \tilde{\nu}). \quad (2.20)$$

We then observe that if μ_n weakly $*$ converges to μ then $W_{\omega_X}(\mu, \mu_n) \rightarrow 0$. Indeed, it is known to imply that the 1-Wasserstein distance (corresponding to W_{ω_X} for $\omega_X(t) = t$) between μ_n and μ converges to 0, so that there is some $\theta_n \in \Pi(\mu, \mu_n)$ which (up to a non relabeled subsequence) weakly $*$ converges to some θ supported on the diagonal of $X \times X$, hence

$$W_{\omega_X}(\mu, \mu_n) \leq \int_{X \times X} \omega_X(d(x, x')) \theta_n(dx, dx') \rightarrow 0.$$

Getting back to the approximation of (2.5), take sequences $c_i^n \in C(X_i \times Z)$, $\mu_i^n \in \mathcal{P}(X_i)$, and $c_i \in C(X_i \times Z)$, $\mu_i \in \mathcal{P}(X_i)$, such that

$$\|c_i^n - c_i\|_{\infty} \rightarrow 0, \quad \mu_i^n \rightharpoonup^* \mu_i \quad (2.21)$$

and set:

$$J(\nu) := \sum_{i=1}^I W_{c_i}(\mu_i, \nu), \quad J_n(\nu) := \sum_{i=1}^I W_{c_i^n}(\mu_i^n, \nu), \quad \forall \nu \in \mathcal{P}(Z). \quad (2.22)$$

Denoting by $\omega_{X_i}^i$ and ω_Z^i common continuity modulus of the c_i^n (the first one in x_i uniformly in z and the second in z , uniformly in x_i just as above) we then have:

Proposition 2.2. *For every $(\nu, \nu_n) \in \mathcal{P}(Z) \times \mathcal{P}(Z)$:*

$$|J(\nu) - J_n(\nu_n)| \leq \sum_{i=1}^I \|c_i - c_i^n\|_{\infty} + \sum_{i=1}^I [W_{\omega_{X_i}^i}(\mu_i^n, \mu_i) + W_{\omega_Z^i}(\nu_n, \nu)] \quad (2.23)$$

this implies in particular

- $J_n(\nu_n) \rightarrow J(\nu)$ whenever $\nu_n \rightharpoonup^* \nu$,
- a quantitative estimate for the stability of values:

$$|\inf_{\mathcal{P}(Z)} J - \inf_{\mathcal{P}(Z)} J_n| \leq \sum_{i=1}^I \|c_i - c_i^n\|_\infty + \sum_{i=1}^I W_{\omega_{x_i}}(\mu_i^n, \mu_i) \quad (2.24)$$

- if ν_n minimizes J_n then, up to a subsequence, it weakly $*$ converges to a minimizer of J .

Proof. The statements directly follow from estimate (2.20) and the already observed fact that the right-hand side of (2.23) converges to 0 as soon as $\nu_n \rightharpoonup^* \nu$. \square

In the case where the cost functions $c_i = c_i^n$ are Lipschitz and the approximated measures μ_i^n satisfy (for the usual 1-Wasserstein distance) $W_1(\mu_i^n, \mu_i) \leq \frac{C}{n}$, (2.23) above just takes the form

$$|J(\nu) - J_n(\nu_n)| \leq \sum_{i=1}^I \text{Lip}(c_i) \left(\frac{C}{n} + W_1(\nu^n, \nu) \right).$$

3 Numerical simulations: Linear Programming

In this section, we present various numerical simulations using the Linear Programming approach of sections 2.3 and 2.4. The localization method of 2.2 is used to approximate the support of the barycenter measure. An alternate approach to approximating the support of the barycenter measure, which can be combined with localization is a two stage solution approach: the first stage, using a coarse grid, gives an approximation of the support of the barycenter, the second stage gives a more accurate representation of the barycenter using information on the support obtained in the first stage.

All computations in this section were performed in MATLAB on a Mid 2011 MacBook Air laptop. To solve (2.15) we use the software package CVX [GB10] [GB08] which is callable from MATLAB. The CVX language allows for a very concise description of the convex optimization problem, and allows for the use of multiple solver libraries (e.g. MOSEK, Gurobi). The numerical solution obtained is the correct up to tolerances near numerical precision.

3.1 Geodesic paths between measures in the plane

We considered two measures in the plane, and by varying the weights in the quadratic cost function, we computed three points on the geodesic path (or McCann interpolant) represented in figure 1. The computational time was less than a minute. The measures are illustrated by a circle centered on the atom (middle of the corresponding square) and a radius proportional to the weight.




Figure 1: Given measures in black. Figure left, centre, right: solution of the geodesic problem with weights .25, .5, .75, respectively.

Both the shape of the support (square, diamond) and the density of the measure are illustrated in the figure: the interpolated measures are influenced by both properties. Figure 2 illustrates the two-stage support refinement strategy.




Figure 2: Refinement for the geodesic problem. Left: solution on full grid. Right: solution on the implied support, but more resolved. The size of the problems is the same, but the resolution is increased by a factor of three.

3.2 Comparing different cost functions

For the next set of examples, we took two uniform measures, the first corresponding to a vertically oriented rectangle, and the second corresponding to an horizontal rectangle. These measures are shown in Figure 3. First, we compared the convergence of the solutions for different grid sizes in Figure 4. Notice that the general support of the computed measures seems stable, but there are oscillations in the density, for different resolutions.

Next, we computed the barycenter with various power cost functions $C(x, y) = |x - y|^p$, for $p = 1, 2, 3, 4$. The solutions we computed use grids of size 50^2 . Computational time was close to two minutes. A second run using grid size of 100^2 and a localization of the support, took 30 to 45 minutes. See figure 5. The densities are plotted using a grayscale which corresponds to the relative values,

however the grayscale is different for each figure. We also include another view of the density for $p = 1$ in Figure 6.

The solutions have a complicated geometry. For the cost with $p = 1$, the support of the barycenter is the entire convex hull of the supports of the measures, although the density is highly concentrated at the intersection of the measures. The density ranges from about 0.01 at the edges to 0.12 in the center.

For the case $p = 2$, the density is supported on a square, but wider than the width of the rectangle. The density has some oscillations, but it is strictly positive (taking values in the range $[.005, .004]$). For the case $p = 3$, the density is supported on a small octagonal shape, with zero density in the middle, and with larger oscillations. For the case $p = 4$, a much larger octagonal shape appears with a large zero density hole in the middle. The supports of the barycenter measures are close to the ones estimated by localization.




Figure 3: The measures m_1, m_2, m_3, m_4 used in the examples which follow.




Figure 4: Comparison of the numerical barycenter for measures m_1, m_2 using cost $C(x, y) = |x - y|^4$ on different grid sizes: $25^2, 50^2$ and 100^2 . Note the general shape of the solutions are similar, but the density has more oscillations at higher resolution.

Finally, we computed the barycenters using all four measures from figure 3, see figure 7.




Figure 5: Barycenters of two rectangles m_1, m_2 , with cost $C(x, y) = |x - y|^p$ for $p = 1, 2, 3$. Using grid size 100^2 .




Figure 6: Surface plot of the barycenter corresponding to $p = 1$.




Figure 7: Barycenters of the measures m_1, m_2, m_3, m_4 (four rotated rectangles) with cost $C(x, y) = |x - y|^p$ for $p = 1, 2, 3$. Using grid size 50^2 .

3.3 Matching for teams

We considered the matching for teams problem and used measures and costs as follows, also see Figure 8. Set the quality domain $Z = [0, 1]^2$ and write $z = (z_1, z_2)$ for points in Z . Set M_0, M_1, M_2 to be measures which have constant density on their support, and let their supports be $[1, 2]^2$, $[1.25, 1.75] \times [1, 2]$, and $[1, 2] \times [1.25, 1.75]$, respectively. The corresponding cost functions are

$$\begin{aligned} c_0(x, z) &= -5.5(x_1 z_1 + x_2 z_2) \\ c_1(x, z) &= c_2(x, z) = (x_1 + z_1)^2 + (x_2 + z_2)^2. \end{aligned}$$

The solution concentrates mass at the boundary, and especially at the corners of the domain.




Figure 8: Solution of the matching for teams problem. Top: the three measures, and the solution. Bottom: surface plot of the solution. The solution concentrates mass mostly on the corners with some mass on the edges.

4 Dual formulation

4.1 Duality and optimality conditions

Let us now explain why the variational problem (2.6) can be naturally be seen as a dual formulation of (2.11) (see [CE10] for more details on this duality). To that end, let us observe that $(\gamma_1, \dots, \gamma_I) \in \Pi$ if and only if (2.12) holds for every i (these are the fixed μ_i marginals constraints) and

$$\int_{X_I \times Z} \sum_{i=1}^I \varphi_i(z) \gamma_i(dx_i, dz) = 0, \text{ as soon as } \sum_{i=1}^I \varphi_i(z) dz = 0, \forall z \in Z. \quad (4.1)$$

Indeed, clearly if the γ_i 's have the same marginal on Z then (4.1) holds. Conversely assume (4.1), let $i \neq j$ and $\varphi \in C(Z)$ then applying (4.1) to the potentials

$\varphi_i = \varphi$, $\varphi_j = -\varphi$ and $\varphi_k = 0$ for $k \notin \{i, j\}$ we see that $\int_{X_i \times Z} \varphi(z) \gamma_i(dx_i, dz) = \int_{X_j \times Z} \varphi(z) \gamma_j(dx_j, dz)$. This proves that (4.1) characterizes the fact that the γ_i 's share the same marginal on Z . This enables us to rewrite (2.11) in inf-sup form:

$$\inf_{\gamma_i \geq 0} \sup \left\{ \mathcal{L}((\gamma_i)_i, (\psi_i)_i, (\varphi_i)_i) : \psi_i \in C(X_i), \varphi_i \in C(Z) : \sum_{i=1}^I \varphi_i = 0 \right\} \quad (4.2)$$

where the Lagrangian \mathcal{L} is given by

$$\begin{aligned} \mathcal{L}((\gamma_i)_i, (\psi_i)_i, (\varphi_i)_i) &:= \sum_{i=1}^I \int_{X_i \times Z} (c_i(x_i, z) - \psi_i(x_i) - \varphi_i(z)) \gamma_i(dx_i, dz) \\ &\quad + \sum_{i=1}^I \int_{X_i} \psi_i(x_i) \mu_i(dx_i). \end{aligned}$$

To obtain the desired dual formulation, we formally switch the inf and the sup (again, we refer to [CE10] for a rigorous derivation):

$$\sup_{(\psi_i, \varphi_i), \sum \varphi_i = 0} \inf_{\gamma_i \geq 0} \mathcal{L}((\gamma_i)_i, (\psi_i)_i, (\varphi_i)_i).$$

We next observe that

$$\begin{aligned} \inf_{\gamma_i \geq 0} \mathcal{L}((\gamma_i)_i, (\psi_i)_i, (\varphi_i)_i) &= \sum_{i=1}^I \int_{X_i} \psi_i(x_i) \mu_i(dx_i) \\ &\quad + \sum_{i=1}^I \inf_{\gamma_i \geq 0} \int_{X_i \times Z} (c_i(x_i, z) - \psi_i(x_i) - \varphi_i(z)) \gamma_i(dx_i, dz) \end{aligned}$$

and the latter infimum is 0 when

$$c_i(x_i, z) \geq \psi_i(x_i) + \varphi_i(z), \quad \forall (x_i, z) \in X_i \times Z \quad (4.3)$$

and $-\infty$ otherwise. The dual of (2.11) therefore consists in maximizing

$$\sum_{i=1}^I \int_{X_i} \psi_i(x_i) \mu_i(dx_i)$$

subject to the constraints (4.3) and $\sum_{i=1}^I \varphi_i = 0$. For fixed φ_i , the maximal ψ_i that satisfies (4.3) being $\psi_i = \varphi_i^{c_i}$, we see that the dual can be equivalently formulated as

$$\sup \left\{ \sum_{i=1}^I \int_{X_i} \varphi_i^{c_i}(x_i) \mu_i(dx_i) : \sum_{i=1}^I \varphi_i = 0 \right\} \quad (4.4)$$

which is exactly (2.6). For the existence of solutions and the equality between the infimum in (2.11) and the supremum in (4.4) (which is obtained by a slightly

different argument), we again refer to [CE10]. Now the optimality conditions for (2.11) and (4.4) are summarized by the equivalence between the following assertions:

- $(\gamma_i)_i \in \Pi$ solves (2.11) and $(\varphi_i)_i$ such that $\sum_{i=1}^I \varphi_i = 0$ solves (4.4),
- $((\gamma_i)_i, (\varphi_i^{c_i})_i, (\varphi_i)_i)$ is a saddle point of \mathcal{L} ,
- for every i , one has

$$\varphi_i^{c_i}(x_i) + \varphi_i(z) = c_i(x_i, z) \quad (4.5)$$

γ_i -almost everywhere on $X_i \times Z$ or, equivalently, by continuity, on the support of γ_i .

4.2 The case of Wasserstein barycenters

From now on, we restrict ourselves to the quadratic case where all the X_i 's and Z are some ball B (say) of \mathbb{R}^d and the costs c_i are quadratic:

$$c_i(x_i, z) := \frac{\lambda_i}{2} |x_i - z|^2$$

where the λ_i 's are positive coefficients which we normalize in such a way that they sum to 1. In this case, (2.5) corresponds to

$$\inf_{\nu \in \mathcal{P}(B)} \sum_{i=1}^I \lambda_i W_2^2(\mu_i, \nu) \quad (4.6)$$

where W_2^2 stands for the squared 2-Wasserstein distance. This problem has been studied in details in [AC11] where uniqueness (under the assumption that one of the measures does not give mass to small sets), characterization, L^p or L^∞ regularity results are established for Wasserstein barycenters as well as a close connection with the quadratic multimarginal optimal transport problem of Gangbo and Świąch [GS98]. Since Wasserstein barycenters constitute a natural way to interpolate between an arbitrary number of measures, they therefore also find applications in image processing [RPDB12] and statistics [BK12].

Let us now informally give the optimality conditions for the barycenter using once again the dual formulation (4.4) (see [AC11] for details). In the present quadratic setting, the formula for the c_i -transform takes the form

$$\varphi_i^{c_i}(x_i) = \inf_z \left\{ \frac{\lambda_i}{2} |x_i - z|^2 - \varphi_i(z) \right\},$$

which, defining

$$u_i(x_i) := \frac{1}{2} |x_i|^2 - \frac{\varphi_i^{c_i}(x_i)}{\lambda_i}$$

can conveniently be rewritten as

$$u_i = \left(\frac{1}{2} |\cdot|^2 - \frac{\varphi_i}{\lambda_i} \right)^*$$

(where $*$ denotes the standard Legendre transform). In particular, u_i is convex (hence differentiable outside of a small set). Moreover the optimal coupling γ_i is concentrated on the set where equality (4.5) holds which setting $v_i = u_i^*$ is equivalent to the relation $u_i(x_i) + v_i(z) = x_i \cdot z$ i.e. $z = \nabla u_i(x_i)$ (provided u_i is differentiable at x_i which is the case μ_i a.e. as soon as μ_i does not charge small sets...). This implies that the barycenter which is also the marginal ν that is common to all the γ_i 's is actually given by $\nu = \nabla u_{i\#} \mu_i$ for every i and ∇u_i is the optimal transport between μ_i and ν for W_2^2 . We then have

$$\frac{1}{2} |\cdot|^2 - \frac{\varphi_i}{\lambda_i} \geq \left(\frac{1}{2} |\cdot|^2 - \frac{\varphi_i}{\lambda_i} \right)^{**} = u_i^* = v_i. \quad (4.7)$$

But one can rewrite the fact that γ_i -almost everywhere equality (4.5) holds as the fact that for γ_i -almost every (x_i, z) , one has

$$\frac{1}{2} |z|^2 - \frac{\varphi_i(z)}{\lambda_i} = x_i \cdot z - u_i(x_i) \leq u_i^*(z) = v_i(z).$$

With (4.7), we then deduce that for ν -a.e. z , one has

$$\frac{1}{2} |z|^2 - \frac{\varphi_i(z)}{\lambda_i} = v_i(z).$$

Recalling that the φ_i have to sum to 0, we deduce that

$$\sum_{i=1}^I \lambda_i v_i(z) = \frac{|z|^2}{2} \quad (4.8)$$

on the support of ν . The optimality conditions for the barycenter $\nu = \nabla u_{i\#} \mu_i = \nabla v_{i\#} \mu_i$ therefore, at least formally take the form of the system of Monge-Ampère equations

$$\nu = \mu_i(\nabla v_i) \det(D^2 v_i), \quad i = 1, \dots, I \quad (4.9)$$

which is supplemented with equation (4.8) on the support of ν . We shall see in the next paragraph how to compute numerically in an efficient way the potentials φ_i .

4.3 An efficient algorithm for Wasserstein barycenters

Discretization of the dual problem. We assume in all this section that the sets X_i 's and Z are subsets of \mathbb{R}^d for some $d = 1, 2$ or 3 . As described in

the previous section, the computation of one Wasserstein quadratic barycenter is equivalent in its dual form to the maximization of

$$\sum_{i=1}^I \int_{X_i} \varphi_i^{c_i}(x_i) \mu_i(dx_i) \quad (4.10)$$

where

$$\varphi_i^{c_i}(x_i) = \inf_z \left\{ \frac{\lambda_i}{2} |x_i - z|^2 - \varphi_i(z) \right\},$$

under the linear constraint $\sum_{i=1}^I \varphi_i(z) = 0$ for all $z \in Z$. This formulation leads to the following natural discretization of Wasserstein quadratic barycenter computation. Suppose $(y_i^j, \nu_i^j)_{j=1, \dots, N_i} \subset X_i \times \mathbb{R}_+$ is a convergent quantization of the measures μ_i . More explicitly, we assume that for all $i = 1, \dots, I$

$$\lim_{N_i \rightarrow \infty} c(N_i) \sum_{j=1}^{N_i} \nu_i^j \delta_{y_i^j} = \mu_i$$

in the sense of the weak convergence of measures. In order to only consider probability measures, we set $c(N_i) = (\sum_{j=1}^{N_i} \nu_i^j)^{-1}$. Additionally, we suppose that (z_k) is a dense countable family of points of Z . Based on (4.10) and previous quantizations, for a given $N_k \in \mathbb{N}$, our discrete optimization problem of $I \times N_k$ variables reads

$$\Phi((\varphi_1^k), \dots, (\varphi_I^k)) = \sum_{i=1}^I c(N_i) \sum_{j=1}^{N_i} \nu_i^j \min_{k=1, \dots, N_k} \left\{ \frac{\lambda_i}{2} |y_i^j - z_k|^2 - \varphi_i^k \right\} \quad (4.11)$$

under the N_k pointwise linear constraints:

$$\sum_i \varphi_i^k = 0, \quad \forall k = 1, \dots, N_k. \quad (4.12)$$

This optimization problem in its dual form can be seen as a large scale non-smooth concave maximization problem. We discuss in the next paragraph alternatives that have been developed to solve numerically this type of problems.

Non smooth algorithms. Many different approaches have been introduced in the last decades to approximate optimal solution of non-smooth concave (or convex) problems, e.g. gradient sampling methods [BLO05] and bundle methods [LMfASA78]. These algorithms make use of a partial or complete description of superdifferentials in order to identify ascent directions (see next paragraph). Even though Proposition 4.1 describes explicitly the whole superdifferential, finding an effective ascent direction in practice is made difficult by the large dimension of some superdifferentials. Additionally, those approaches are essentially of order one and follow the singular parts of the graph of the cost function. These two facts could explain a slow rate of convergence when starting from an initial point far away from any optimal vector.

One surprisingly efficient alternative for minimizing non-smooth functions is the use of quasi-Newton methods. It is known [Pow76] that if the maximized function, Φ , is twice continuously differentiable and the suplevel set $\Phi \geq \Phi(x_0)$ is bounded, then the sequence of function values generated by the BFGS method with inexact Armijo-Wolfe line search, starting from x_0 converge to the maximal value of Φ . More recently, it has been pointed out by different authors [HUL96, LO09, LO08] that variable metric algorithms may produce in some cases sequences which converge to an optimal point in the sense of Clarke. The mathematical analysis of this good behavior has just been initiated in recent papers of Overton [LO09, LO08]. This efficiency could be explained heuristically by the fact that the approximated inverse of the Hessian matrix has a spectral decomposition in two subspaces which describe the two different local behaviors of the cost Φ : a subspace associated to the regular directions of the cost function Φ , and the subspace of eigenvectors whose eigenvalues are small in absolute value, which correspond to the singular directions of Φ .

It has been observed in simple situations that L-BFGS algorithms are sometimes able to converge to an optimal point. In more standard examples, where concentration can occur for instance, L-BFGS approach fails to converge. This expected behavior for strongly non-smooth functions illustrates the need for using more specific non-smooth techniques close from the optimal point. The costly, but reliable, bundle type algorithms have demonstrated their efficiency in this context.

We will not give here a detailed study of quasi Newton methods applied to optimal transportation which would be out of the scope of this paper. We only point out that the L-BFGS (low memory version of Broyden-Fletcher-Goldfarb-Shanno algorithm) algorithm combined with a bundle approach gives a rather efficient way to solve this type of problem. We refer to [HMM07] for a careful study and an efficient implementation of this kind of hybrid algorithm.

Gradient computation. The previous approach relies on the capability of providing at every iteration one supergradient vector of the current iterate. It is straightforward to obtain the following characterization of the supergradient of the discrete dual cost Φ :

Proposition 4.1. *Let $(\varphi_1^k), \dots, (\varphi_I^k) \in \mathbb{R}^{N_k \times I}$. Then*

$$((v_1^k), \dots, (v_I^k)) \in \partial\Phi((\varphi_1^k), \dots, (\varphi_I^k))$$

if and only if it is a convex combination of the finite set of extremal vectors defined in the following way. Let $\varphi_i^{k(i,j)}$ be any selection of minimizing values involved in the dual cost. That is $\forall i, j$*

$$\varphi_i^{k(i,j)*} \in \operatorname{argmin}_{k=1, \dots, N_k} \left\{ \frac{\lambda_i}{2} |y_i^j - z_k|^2 - \varphi_i^k \right\}. \quad (4.13)$$

Then, the set of extremal vectors is the finite collection of all vectors

$$(e^{k*}) = \left(\sum_j \sum_k -c(N_i) \nu_i^j \delta_{k(i,j)*}(k) \right) \quad (4.14)$$

for any selection $(k(i, j)^*)$.

A crucial observation that has been raised in [MO12] is the fact that the computation of a vector of the superdifferential does not require generically an order of $I \times N_k \times \sum N_i$ operations. Actually the following formulation makes it possible to use a special data structure called *kd-tree* which in most cases reduces the complexity of finding one vector of $\partial\Phi((\varphi_1^k), \dots, (\varphi_I^k))$.

Let i, j be given and assume we want to evaluate the minimal value

$$M = \min_k \left\{ \frac{\lambda_i}{2} |y_i^j - z_k|^2 - \varphi_i^k \right\}.$$

Let us then define $c_i = \max_k \varphi_i^k$ and

$$M = -c_i + \min_k \left\{ \frac{\lambda_i}{2} |y_i^j - z_k|^2 + c_i - \varphi_i^k \right\}.$$

Since the latter term is positive:

$$M = -c_i + \min_k \|P_i^j - Q_i^k\|^2$$

where $P_i^j = (\sqrt{\frac{\lambda_i}{2}} y_i^j, 0)$ is a fixed vector of \mathbb{R}^{d+1} , $Q_i^k = (\sqrt{\frac{\lambda_i}{2}} z_k, \sqrt{c_i - \varphi_i^k})$ and $\|\cdot\|$ stands for the standard euclidean norm of \mathbb{R}^{d+1} . Thus our supergradient request reduces to identify one closest point of P_i^j among points of $(Q_i^k)_k$. Observe that the family $(Q_i^k)_k$ does not depend of the parameter j . This task is a standard operation in computational geometry which can be performed efficiently with *kd-tree* structures. By using such tools, we can reduce the complexity of the supergradient request in the generic case to an order of $I \times \log N_k \times \sum N_i$ operations. Observe that if the N_i 's and N_k are of same order N , one request is generically of complexity of order $I^2 N \log N$.

Numerical quantization and localization. We suppose that all the measures μ_i are compactly supported. We use the discretization of Section 2.4. The support of the unknown barycenter measure is bounded using the results of Section 2.2, in particular, (2.10)

Reconstruction of the barycenter density. One additional difficulty associated to the dual formulation is the fact that optimal dual vectors only give an implicit description of the barycentric measure. In order to recover the support and the density of the barycentric measure, we introduce the following least square procedure.

Every optimal dual vector $(\varphi_i^k)_k$ must be associated to an optimal transport from μ_i to the barycentric measure. A crucial observation is the fact that every associated map transports the μ_i to the same measure. We exploit this optimality condition to recover the barycentric measure through the coefficients




Figure 9: Classical McCann interpolation between translated measures

$f_i^{j,k}$ which describe the mass transported from y_i^j to z_k . By optimality, some mass can be transported from y_i^j to z_k if and only if

$$G_{i,j}(\varphi_i^k) = \min_{k'=1,\dots,N_k} G_{i,j}(\varphi_i^{k'}). \quad (4.15)$$

where $G_{i,j}(\varphi_i^{k'}) = \frac{\lambda_i}{2} |y_i^j - z_k|^2 - \varphi_i^{k'}$. Let us fix some parameter $\varepsilon > 0$. Based on the previous observation, we only consider the unknown coefficients $f_i^{j,k}$ for which $G_{i,j}(\varphi_i^k)$ is less than the optimal value (4.15) plus ε . In order to recover the barycentric measure, we look for the set of coefficients which generate the same measures in an optimal least square sense. More precisely, we solve the sparse least square problem:

$$\min_{f_i^{j,k}} \sum_{l,m} \sum_k \left(\sum_j f_l^{j,k} \nu_l^j - \sum_{j'} f_m^{j',k} \nu_m^{j'} \right)^2 \quad (4.16)$$

among non negative coefficients less than one which satisfy the linear constraints

$$\forall i, j, \sum_k f_i^{j,k} = 1.$$

4.4 Numerical results

As detailed in the previous paragraphs, our approach relies first on a non smooth optimization step using an hybrid LBFGS/Bundle algorithm and a fast computation of supergradient vectors. In a second phase, a sparse least square problem is solved in order to recover an approximation of barycentric density. Let us point out that in all the following examples, the first optimization step was stopped after one hour of computation on a standard PC. This costly step could be dramatically sped up in using a straightforward parallelized cost evaluation.

We validate our approach by considering different test cases for which analytic descriptions of barycenters are available. The simplest situation is the case of the barycenters of a measure of density $\rho(\cdot)$ and a translated measure of density $\rho(\cdot + V)$ where V is some fixed vector. In this trivial case, the isobarycentric measure is of course the measure of density $\rho(\cdot + V/2)$. We display in figure 9, the barycentric measure obtained by our located approximation scheme applied to $\rho = \chi_c$ where χ_c is the characteristic function of a unit square. In this experiment, we used a grid of size 200×200 and a recovering parameter $\varepsilon = 10^{-5}$. In the least square optimization problem (4.16), we obtain an error of order 10^{-4} for every quadratic term.

Next we applied our approach to the case of Gaussian measures. A complete description of barycenters of Gaussian measures has been given in [AC11]: consider a family of Gaussian measures $\mu_i(m_i, S_i)$ of means $(m_i)_i$ and covariance matrices $(S_i)_i$. Then, the barycentric measure associated to the non-negative weights $(\lambda_i)_i$ is a Gaussian measure of mean the barycenter of the $(m_i, \lambda_i)_i$. Moreover, its covariance matrix is the only definite positive matrix S solution of the equation

$$\sum_i \lambda_i (S^{1/2} S_i S^{1/2})^{1/2} = S. \quad (4.17)$$




Figure 10: Isobarycenter computation of three gaussian measures by a global (first row) and a localized approach (second row).

We denote by $\mathcal{N}(m_i, \sigma_i)$ a Gaussian of mean m_i and of covariance matrix equal to $\sigma_i^2 Id$. We considered two different test cases and applied for both our global and localized approaches. In all the experiments the number of sampling points of the given measures and of the barycentric measure have been fixed for the global approach to $\forall i, N_i = N_z = 15 \times 10^3$. For the localized approach by Minkowski sum, we imposed $\forall i, N_i = N_z = 5 \times 10^3$. The first test case consists in approximating the isobarycenter of the three Gaussian measures of

random standard variations $\mathcal{N}((0.1, 0.8), 1/49.75)$, $\mathcal{N}((-0.9, -1), 1/35.89)$ and $\mathcal{N}((1, -0.9), 1/74.63)$. The resulting barycenter and the given measures are presented in figure 10. Our second test case is related to the approximation of the barycenter of the five gaussian measures $(\mathcal{N}(m_i, \sigma_i), \lambda_i)_{i=1,\dots,5}$ where the m_i are the vertices of a regular pentagon with $\sigma_i = 1/50$ and $\lambda_i = 1/7$ if i is odd and $\sigma_i = 1/100$ and $\lambda_i = 2/7$ otherwise. The resulting barycenter and the given measures are drawn in figure 11. Observe that in the localized illustrations, the support of the unknown measure is not anymore centered due to the loss of symmetry in Minkowski's sum. To conclude the study of those test cases, we provide in table 4.4, the errors between the theoretical and computed means and covariance matrices. As expected, even if the number of degree of freedom is smaller, the results obtained by the located approach are significantly better than the ones obtained by the first algorithm.

To conclude our numerical experiments, we provide large scale examples in which we interpolate three textures of images of size 150×150 . This type of applications have been first studied in the framework of optimal transportation in [RPDB12]. The texture mixing problem consists in synthesizing a texture from a family of given textures. The interest of using Wasserstein barycenters in this context is due to the spatial nature of Wasserstein distance which provides a more natural interpolation process than the naive pointwise means (see figures 12 and 13). We carried out similar experiments as the one depicted in [RPDB12]. Our contribution with respect to [RPDB12] lies in the fact that we do not replace the quadratic Wasserstein distance by the easier to handle so-called sliced Wasserstein distance (which is an average over directions of one dimensional Wasserstein distances). We obtained by our method an approximation of the original model up to an error of 10^{-3} for every quadratic term of (4.16).

Acknowledgements: The authors are grateful to the hospitality of BIRS where the present work was initiated at the occasion on a focused group meeting on Numerical methods for optimal transport. They are happy to thank Martial Agueh, Jean-David Benamou and Brendan Pass for many fruitful conversations. The first author gratefully acknowledges the support of the ANR,

	$ m_i^{th} - m_i $	$ \sigma_i^{th} - \sigma_i $
First test case	0.003	0.005
First test case localized	0.0002	0.0014
Second test case	0.002	0.07
Second test case localized	0.0002	0.0013

Table 1: Upper bounds of the gap between theoretical and computed means and covariance coefficients




Figure 11: Non uniform barycenter computation of five gaussian measures by a global (first row) and a localized approach (second row).




Figure 12: Isobarycenter of the three textures of the first row. The pointwise mean of the three textures corresponds to the left picture of the second row. Wasserstein barycenter is presented on the right picture of the second row.




Figure 13: Isobarycenter of the three textures of the first row. The pointwise mean of the three textures corresponds to the left picture of the second row. Wasserstein barycenter is presented on the right picture of the second row.

through the projects ISOTACE and OPTIFORM and INRIA through the "action exploratoire" MOKAPLAN.

References

- [AC11] Martial Agueh and Guillaume Carlier. Barycenters in the wasserstein space. *SIAM Journal on Mathematical Analysis*, 43(2):904–924, 2011.
- [BB00] Jean-David Benamou and Yann Brenier. A computational fluid mechanics solution to the Monge-Kantorovich mass transfer problem. *Numer. Math.*, 84(3):375–393, 2000.
- [BFO14] Jean-David Benamou, Brittany D Froese, and Adam M Oberman. Numerical solution of the optimal transportation problem using the monge-ampère equation. *Journal of Computational Physics*, 260:107–126, 2014.
- [BK12] Jérémie Bigot and Thierry Klein. Consistent estimation of a population barycenter in the wasserstein space. *arXiv preprint arXiv:1212.2562*, 2012.

- [BLO05] J.V. Burke, A.S. Lewis, and M.L. Overton. A robust gradient sampling algorithm for nonsmooth, nonconvex optimization. *SIAM Journal on Optimization*, 15(3):751–779, 2005.
- [Bre91] Yann Brenier. Polar factorization and monotone rearrangement of vector-valued functions. *Comm. Pure Appl. Math.*, 44(4):375–417, 1991.
- [CE10] Guillaume Carlier and Ivar Ekeland. Matching for teams. *Economic Theory*, 42(2):397–418, 2010.
- [GB08] Michael Grant and Stephen Boyd. Graph implementations for nonsmooth convex programs. In V. Blondel, S. Boyd, and H. Kimura, editors, *Recent Advances in Learning and Control*, Lecture Notes in Control and Information Sciences, pages 95–110. Springer-Verlag Limited, 2008. http://stanford.edu/~boyd/graph_dcp.html.
- [GB10] M. Grant and S. Boyd. CVX: Matlab software for disciplined convex programming, version 1.21. <http://cvxr.com/cvx>, May 2010.
- [GM13] Nassif Ghoussoub and Abbas Moameni. A self-dual polar factorization for vector fields. *Comm. Pure Appl. Math.*, 66(6):905–933, 2013.
- [GS98] Wilfrid Gangbo and Andrzej Swiech. Optimal maps for the multidimensional monge-kantorovich problem. *Communications on pure and applied mathematics*, 51(1):23–45, 1998.
- [HMM07] Napsu Haarala, Kaisa Miettinen, and Marko M Mäkelä. Globally convergent limited memory bundle method for large-scale nonsmooth optimization. *Mathematical programming*, 109(1):181–205, 2007.
- [HUL96] J.B. Hiriart-Urruty and C. Lemaréchal. *Convex analysis and minimization algorithms*, volume 1. Springer, 1996.
- [LMfASA78] C. Lemaréchal, R. Mifflin, and International Institute for Applied Systems Analysis. *Nonsmooth optimization*. Pergamon Press, 1978.
- [LO08] A.S. Lewis and M.L. Overton. Behavior of BFGS with an exact line search on nonsmooth examples. Technical report, Technical report, Optimization Online, 2008b. http://www.optimization-online.org/DB_FILE/2008/12/2173.pdf, submitted to SIAM J. Optimization, 2008.
- [LO09] A.S. Lewis and M.L. Overton. Nonsmooth optimization via BFGS. *SIAM Journal of Optimization*, submitted for publication, 2009.

- [McC97] Robert J. McCann. A convexity principle for interacting gases. *Adv. Math.*, 128(1):153–179, 1997.
- [MG14] Bernard Maurey and Nassif Ghoussoub. Remarks on multi-marginals symmetric monge-kantorovich problems. *Discrete Cont. Dyn. Syst.*, 34(4):1465–1480, 2014.
- [MO12] Quentin Mérigot and Édouard Oudet. Discrete optimal transport: complexity, geometry and applications. Preprint, 2012.
- [Pas12a] Brendan Pass. Multi-marginal optimal transport and multi-agent matching problems: uniqueness and structure of solutions. *arXiv preprint arXiv:1210.7372*, 2012.
- [Pas12b] Brendan Pass. On the local structure of optimal measures in the multi-marginal optimal transportation problem. *Calc. Var. Partial Differential Equations*, 43(3-4):529–536, 2012.
- [Pow76] M.J.D. Powell. Some global convergence properties of a variable metric algorithm for minimization without exact line searches. *Nonlinear programming*, 9:53–72, 1976.
- [PPO14] Nicolas Papadakis, Gabriel Peyré, and Edouard Oudet. Optimal transport with proximal splitting. *SIAM Journal on Imaging Sciences*, 7(1):212–238, 2014.
- [RPDB12] Julien Rabin, Gabriel Peyré, Julie Delon, and Marc Bernot. Wasserstein barycenter and its application to texture mixing. In *Scale Space and Variational Methods in Computer Vision*, pages 435–446. Springer, 2012.
- [Sch03] Alexander Schrijver. *Combinatorial optimization: polyhedra and efficiency*, volume 24. Springer, 2003.
- [Vil03] Cédric Villani. *Topics in optimal transportation*, volume 58. AMS Bookstore, 2003.
- [Vil09] Cédric Villani. *Optimal transport: old and new*, volume 338. Springer, 2009.