Optimal Transport: From moving soil to same-sex marriage

By Nassif Ghoussoub (UBC) and Brendan Pass (U. Alberta)

Optimal mass transport has come a long way since the 1781 "Mémoire sur la théorie des déblais et des remblais" of Gaspard Monge, who was looking for the most economical way of moving soil from one area to another. Mathematically, this amounts to minimizing the total cost $\int_X c(x, Tx)d\mu$ over all possible transport maps T that "push" the initial distribution μ of soil onto a final distribution ν, where $c(x, y)$ is the cost of moving x to y, which (for Monge) was proportional to the Euclidian distance $|x - y|$.

1 Monge-Kantorovich problems

Many years later, Kantorovich linearized and compactified the problem by enlarging the constraint set to contain all "transport plans"; that is, he allowed a piece of soil to be "split" between two or more destination points, hence multivalued mappings. This relaxed version of Monge's problem, which earned Kantorovich, together with T. Koopmans, the Nobel Prize in economics in 1975 for their work on optimum allocation of resources, consists of considering the following minimization problem

$$T_c(\mu, \nu) := \inf \left\{ \int_{X \times Y} c(x, y) d\gamma(x, y); \gamma \in \Pi(\mu, \nu) \right\}. \tag{1}$$

Here $\Pi(\mu, \nu)$ is the set of measures γ on the product space $X \times Y$ whose marginals are μ and ν. Kantorovich also defined a dual problem, allowing him to relate it to linear programming. It is instructive to think of a manufacturing company shipping resources (say iron) from a distribution $\mu(x)$ of mines on some landscape $X \subseteq \mathbb{R}^n$ to a distribution $\nu(y)$ of factories on a landscape $Y \subseteq \mathbb{R}^n$, where $c(x, y)$ is the cost of shipping one unit of iron from a mine at location x to a factory at location y; the goal is minimize the total shipping cost. A permissible γ represents a possible transport plan; heuristically $d\gamma(x, y)$ represents the amount of iron that should be shipped from mine x to factory y. For a good outline of the 2-marginal case and its applications, we refer to Villani [7].

1.1 Multi-marginal Monge-Kantorovich problems

Suppose now that a manufacturing company is just beginning business and has not yet built their factories. The company is making a certain product, requiring several resources, such as iron, aluminum, nickel, etc. There is a distribution of mines $\mu_i(x_i)$, supported on some set $X_i \subset \mathbb{R}^n$, producing each type of resource, and the cost to ship one unit of the ith resource from x_i to a location y is given by $c_i(x_i, y)$. The company then wants to build its factories in locations that minimize the total shipping costs of all the resources. That is, they want to build a distribution of factories $\nu(y)$ on Y in order to minimize

$$\sum_{i=1}^m T_{c_i}(\mu_i, \nu). \tag{2}$$

Another way to interpret this problem is to consider the function

$$c(x_1, ..., x_m) = \inf_{y \in Y} \sum_{i=1}^m c_i(x_i, y). \tag{3}$$

Assuming this infimum is always attained at a unique point $y(x_1, x_2, ...x_m)$, there is an equivalence between (2) and the problem of minimizing,

$$T_c(\mu_1, \mu_2, ..., \mu_m) := \inf \left\{ \int_{X_1 \times ..... \times X_m} c(x_1, x_2, ...x_m) d\gamma(x_1, x_2, ..., x_m); \gamma \in \Pi \right\}, \tag{MK}$$

over the set $\Pi := \Pi(\mu_1, \mu_2, ..., \mu_m)$ of measures γ on $X_1 \times ... \times X_m$ whose marginals are the μ_i. This is the *multi-marginal optimal transport (or Monge-Kantorovich) problem.* The case $m = 2$ is obviously the above mentioned classical optimal transport problem (1). Intuitively, $d\gamma(x_1, x_2, ..., x_m)$

represents the amount of resources that are shipped from locations $x_1, x_2, ..., x_m$ to a certain factory $y(x_1, x_2, ..., x_m)$. A fundamental problem (largely settled when $m = 2$) is to determine for which *cost functions* c, the infimum in (MK) is attained (uniquely!) by a measure supported by "a graph", meaning that

$$T_c(\mu_1, \mu_2, ..., \mu_m) = \int_{X_1} c(x, T_1 x, T_2 x, ... T_{m-1} x) d\mu_1(x),$$

for some maps $T_i : X_1 \to X_{i+1}$ that push the first marginal μ_1 onto μ_{i+1} for $i = 1, ..., m-1$. Recently, problems of this general type have begun to attract attention, due to surprisingly diverse applications. But unlike the classical case ($m = 2$), the structure of solutions to multi-marginal problems of form (MK) are not yet well understood. While there has been some progress on the uniqueness and structure of solutions to (MK) (see [6],[5] and the references therein), it has mostly been restricted to cost functions of the form (3), whereas many of these applications involve costs which are *not* of this form. Below, we outline several different applications of this problem.

1.2 Multi-agent matching problems in economics

Recent papers link (MK) to a matching problem in economics where agents' preferences depend on external contracts [1]. For example, consider a collection of consumers, parametrized by the set $X_1 \subseteq \mathbb{R}^n$, looking to buy custom built houses; imagine that the different components x_1^j of a consumer $x_1 = (x_1^1, x_1^2, ..., x_1^n) \in X_1$ represent characteristics which affect the consumers' preferences for different types of houses, for example, their income, family size, age, etc. Think of the probability measure $\mu_1(x_1)$ as representing the relative frequency of a consumer of type x_1. In order to build a house, a consumer must hire several (say $m-1$) tradespeople: for example, a carpenter, a plumber and an electrician. Imagine, for example, that X_2 parametrizes the set of carpenters available to be hired; the different components of $x_2 \in X_2$ may represent the age, years of experience and safety record, for example, of the carpenter x_2, and the measure $\mu_2(x_2)$ the relative frequency of carpenters of type x_2. The sets $X_3, ..., X_m$ will have similar interpretations in terms of plumbers, electricians, etc.

Now, suppose the set of houses that can feasibly be built is parameterised by $Y \subseteq \mathbb{R}^n$; the different components of a house $y \in Y$ may represent its size, location, etc. Of course, different consumers prefer different types of houses; let $f_1(x_1, y) \in \mathbb{R}$ represent the *utility* consumer x_1 would derive from owning a house of type y. Similarly, preferences differ among carpenters, plumbers and electricians as well; let $f_i(x_i, y)$ be the utility worker x_i would derive from building house y. Consumers want to buy houses which they like as much as possible, but also want to pay as little as possible for them. On the other hand, workers want to build houses making their utilities as high as possible, but they also want to be paid as high a wage as possible. Informally, if consumer x_1 hires carpenter x_2, plumber x_3, etc, to build some feasible house, then

$$b(x_1, x_2, ..., x_m) := \sup_{y \in Y} \sum_{i=1}^m f_i(x_i, y)$$

is the maximal total utility that can be obtained by this collection of agents. The link with (MK) is that finding an equilibrium in this market (in other words, an assignment of wages and agents to different types of houses so that no one has an incentive to change jobs) is equivalent to solving (MK), with cost function equal to $-b$.

2 Symmetric Monge-Kantorovich problems

Consider now problem (MK), but with the additional constraint that the measures γ in $\Pi(\mu_1, \mu_2, ..., \mu_m)$ should be invariant under the cyclic permutation $\sigma(x_1, x_2, ... x_m) = (x_2, x_3, ..., x_m, x_1)$; note that in this case, the marginals μ_i must all be equal to some common distribution μ. Here, the problem is to determine for which costs c, there exists an optimal measure that is supported on a graph of the form $x \to (x, Sx, S^2 x, ..., S^{m-1} x)$, where S is a μ-measure preserving m-involution, i.e. $S^m x = x$

a.e.

2.1 Monotone maps and polar factorizations of vector fields

When the cost function is given by $c(x_1, x_2, ..., x_m) = -\sum_{i=2}^{m}\langle u_i(x_1), x_i\rangle$, for a given family of bounded vector fields $(u_2, u_3, ..., u_m)$, the symmetric Monge-Kantorovich problem turns out to be instrumental in the proof of the following representation result for the u_i established by Ghoussoub-Moameni [5]: There exist a cyclically antisymmetric saddle function $H : X^m \to \mathbb{R}$ (i.e., $\sum_{i=0}^{m-1} H(\sigma^i(\cdot)) \equiv 0$ on X^m) and a measure preserving map $S : X \to X$ with $S^m = I$ such that

$$(u_2(x), u_3(x), ..., u_m(x)) = \nabla_{x_2,x_3,...x_m} H(x, Sx, S^2x...., S^mx) \text{ for a.e. } x \in X. \tag{4}$$

This extends an earlier decomposition for a single vector field ($m = 2$) by the same authors. Note that in the special case where the $u_i's$ are *jointly m-monotone*, one can take S to be the identity [4], which extends a well known theorem of Krauss for 2-monotone vector fields. In the other extreme, a classical result of Rockafellar yields that vector fields which are m-cyclically monotone for every m are essentially sub-differentials of convex functions. Taking $u_1 = u$, $u_i = 0$ for $i = 2, 3, ...m - 1$, the result of Ghoussoub-Moameni then yields that every bounded vector field u is m-monotone up to a measure preserving m-involution.

2.2 Matching and the Roommate problem

The economic literature has mostly modeled the marriage market as a bipartite matching game with transferable utility. Yet the bipartite assumption –even in the marriage market– is becoming restrictive in some contexts, where a match does not have to include exactly one individual from each of two exogenously given subpopulations, especially now that a growing number of countries have authorized same-sex unions in some form. This leads to problems with symmetry constraints in many types of matching problems [3]. Another example comes from a university housing office trying to assign students to dorm rooms, say three to a room. The problem of finding an assignment which maximizes some measure of overall compatibility between roommates can be formulated as (MK) but again restricted to measures γ which are invariant under any permutation of the arguments. Heuristically, the invariance arises because the population of interest is not a priori partitioned into disjoint subsets; one works with three copies of the original distribution of students. If it is optimal to couple a trio (x_1, x_2, x_3) of students together, it should also be optimal to couple the trio (x_2, x_3, x_1).

2.3 Density functional theory

A fundamental question in chemical physics is to determine the ground state energy of a system of m-interacting electrons (for example, an atom). A partitioning of this search leads to consideration of the Hohenberg-Kohn functional, which in the semi-classical limit, takes the form [2]:

$$F_{HK}[\mu] := \inf_{\gamma\in\Pi(\mu,...\mu)} \int_{\mathbb{R}^{dm}} \sum_{i\neq j} \frac{1}{|x_i - x_j|} d\gamma.$$

This is exactly problem (MK), where the cost function $\sum_{i\neq j} \frac{1}{|x_i - x_j|}$ represents the Coulombic interaction energy between the electrons. Note that in this case, the marginals, which represent the single particle densities of the electrons, are all the same, embodying the indistinguishability of the electrons. The measures γ in $\Pi(\mu, \mu, ..., \mu)$ represent potential N-particle densities of the system, each with single particle density μ. Heuristically, we can think of $F_{HK}[\mu]$ as representing the minimal (semi-classical) energy of all configurations of electrons, with single particle density μ. The problem in density functional theory is then to minimize $F_{HK}[\mu]$ (or $F_{HK}[\mu]$ plus an external potential) over all possible single particle densities μ. We should note that it is physically natural to impose that the measures γ are invariant under any permutation of the arguments. This does not effect the value of $F_{HK}[\mu]$, as symmetrizing the measure γ does not change the common marginal μ or the value $\int_{\mathbb{R}^{dm}} \sum_{i\neq j} \frac{1}{|x_i - x_j|} d\gamma$, due to the symmetry of the cost function. It is relevant, however, to questions about the structure and uniqueness of the optimal γ. In the case of two electrons (i.e., $m = 2$), it can be shown that the infimum $F_{HK}(\mu)$ is then attained at a measure $\bar{\gamma}$ which determines the co-motion function $x \to (x, Sx)$, with $S^2 = I$. The case when $m \geq 3$ is wide open.

References

[1] G. Carlier, I. Ekeland, *Matching for teams*, Econom. Theory 42 (2010) 397-418.

[2] Cotar, G. Friesecke, C. Klüppelberg. *Density functional theory and optimal transportation with Coulomb cost*, preprint available at arXiv:1104.0603 (2011).

[3] P-A Chiappori, A. Galichon, B. Salanié, *The Roommate Problem is More Stable than You Think*, preprint (2011).

[4] A. Galichon, N. Ghoussoub, *Variational representations for N-cyclically monotone vector fields*, arXiv:1207.2408v2 [math.OC] (2012).

[5] N. Ghoussoub, A. Moameni, *Symmetric Monge Kantorovich problems and polar decomposition of vector fields*, arXiv:1302.2886 (2013).

[6] B. Pass, *PhD thesis*, University of Toronto, 2011.

[7] C. Villani, *Topics in optimal transportation*, volume 58 of *Graduate Studies in Mathematics.* American Mathematical Society, Providence, 2003.