A Minmax algorithm for Optimal Transportation

Alexandre Milewski *

Abstract

In Optimal Transportation one seeks to construct an optimal mapping between two given distributions. In this paper, we consider a variant of this problem, suited for change detection. From this standpoint samples drawn from both distributions are given, yet the underlying distribution is unknown. We detail a numerical scheme for calculating the Optimal map given a distance squared cost.

1 Introduction

Systems often change in ways which tend to minimize some kind of cost. It might, therefore, stand to reason that this change could be linked to optimal transport and that, given the correct cost function, solving the optimal transport map between the state of a system before and after change may serve to detect what change has occurred. This raises a need for a numerical method to solve optimal transport problems, in particular, 'statistical' problems in which the initial and final states of the system are known through samples rather than their distributions. In this paper we detail a method given in [Essid et al., 2018] and how it may be applied to change detection.

The seminal problem to the theory of Optimal Transport was formulated by Gaspard Monge; given a 'sand-pile' and a 'hole' of the same volume, find a map which transports the sand into the hole, which minimizes the transportation cost.

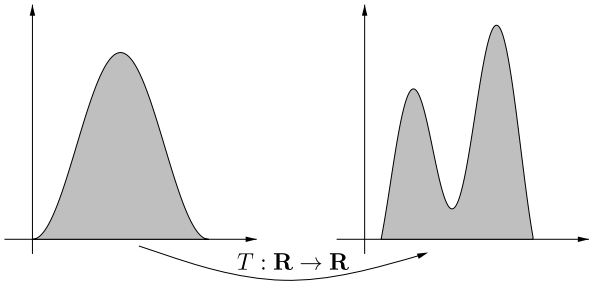

*Supervisors: Prof. Esteban Tabak, Debra Laefer

Monge's problem can be expressed more rigorously as follows

Problem 1 (Monge)**.** *Given two normalized mass distributions on* $\mathbf{R}^d$*, say* P *and* Q*, we seek a transport map* $T : \mathbf{R}^d \to \mathbf{R}^d$ *satisfying*

$$\int_A dQ = \int_{T^{-1}(A)} dP \text{ for any Borel subset } A \subset \mathbf{R}^d \tag{1}$$

which minimizes the quantity

$$\int_{\mathbf{R}^d} \phi(x, T(x))\, dP \tag{2}$$

for some given cost function $\phi : \mathbf{R}^d \times \mathbf{R}^d \to \mathbf{R}$

Intuitively speaking, the map T is a blueprint, instructing where the mass at each point should be redistributed to, and condition (1) is essentially a mass conservation condition. $\phi(x, y)$ is the transportation cost of moving mass from the point x to y. When (1) is satisfied, we say T is a pushforward from P to Q, denoted $T_{\#}P = Q$.

To modify this problem to detect change in a system, say two photos taken at two instances in time, we represent the initial and final states of the system by two probability distributions. These take the place of the 'sand-pile' and 'hole' in Monge's problem. Under a suitable choice of cost function, the optimal map between the two distributions would serve as an indicator of where change had occurred. In this paper, we chose the distance squared norm $\phi(x, y) = \|x - y\|^2$. In our new framework, Monge's problem becomes

Problem 2. *Given two probability distributions on* $\mathbf{R}^d$*, say* P *and* Q*, find the map* $T : \mathbf{R}^d \to \mathbf{R}^d$ *such that*

$$T = \arg\min_F \left\{ \int \|x - F(x)\|^2 \, dP \;\middle|\; F_{\#}P = Q \right\}$$

[Essid et al., 2018] derives the reformulation of problem 2 into a minmax problem, and [Essid et al., 2019] provides a natural algorithm for solving the minmax problem and hence finding the Optimal map.

2 Brenier's Theorem

A priori, there is no guarantee that problem 2 is well posed; such a T might not exist. Brenier's theorem, stated without proof, rectifies this by guaranteeing that, given assumptions on P and Q, such a pushforward T exists and is unique.

Theorem 2.1 (Brenier's Theorem)**.** *Let* P*,* Q *be probability distributions over* $\mathbf{R}^d$*, such that* P *and* Q *don't assign any mass to sets of Hausdorff dimension* $d - 1$ *or less. Then there exists a unique solution to problem 2 of the form* $T = \nabla\varphi$ *where* $\varphi : \mathbf{R}^d \to \mathbf{R}$ *is some convex function. Furthermore,* φ *is the only convex function (neglecting the constant term) such that* $\nabla\varphi_{\#}P = Q$*.*

Proof. A proof can be found in [Villani, 2003] □

The second part of Brenier's theorem implies that the pushforward condition admits a unique transport map T such that T is the gradient of a convex scalar function. The first part of Brenier's theorem then implies that such a map must also minimize the cost and therefore be the solution to problem 2. Hence problem 2 has the following equivalent formulation.

Problem 3. *Find a convex function φ whose gradient satisfies the pushforward condition*

$$\nabla\varphi_{\#}P = Q \tag{3}$$

3 A Minmax reformulation

In order to construct an algorithm to solve problem 3, one must search the set of convex scalar functions for the single function which satisfies (3). This equality constraint is impractical to enforce. Instead, [Essid et al., 2018] introduces an indicator, the Kullback-Leibler divergence, to measure the 'difference' between $\nabla\varphi_{\#}P$ and Q. When the Kullback-Leibler divergence is minimized, (3) holds. Therefore, the minimization of this quantity can be used to direct the search in the space of convex functions.

3.1 Kullback-Leibler Divergence

Definition 3.1. *For any two given probability distributions P and Q, We define the Kullback-Leibler Divergence, denoted $D_{KL}(P\|Q)$ to be*

$$D_{KL}(P\|Q) := \int \ln\left(\frac{dP}{dQ}\right)\, dP$$

Lemma 3.1. *The Kullback-Leibler Divergence $D_{KL}(P\|Q)$ is non-negative and is zero if and only if $P = Q$.*

Proof.

1. If P is not absolutely continuous with respect to Q, then $D_{\mathrm{KL}}(P\|Q) = +\infty > 0$.

2. If P is absolutely continuous with respect to Q, then dP/dQ is integrable with respect to Q. Therefore,

$$\begin{aligned} D_{\mathrm{KL}}(P\|Q) &= \int \ln\left(\frac{dP}{dQ}\right)\, dP \\ &= \int \frac{dP}{dQ} \ln\left(\frac{dP}{dQ}\right)\, dQ \\ &\geq \int \frac{dP}{dQ}\, dQ\ \ln\left(\int \frac{dP}{dQ}\, dQ\right) = 0 \end{aligned} \tag{4}$$

by Jensen's inequality, since $x \to x\ln(x)$ is convex. Equality of equation (4) holds if and only if $dP/dQ = \text{Const.}$ since P and Q have total mass 1, this is equivalent to $P = Q$.

□

It follows from this lemma that minimizing the Kullback-Leibler divergence is equivalent to satisfying the pushforward condition. This provides an alternative to formulation of problem 3.

Theorem 3.1. *φ is a solution to problem 3 if and only if φ and satisfies*

$$\varphi = \arg\min_{\psi} \{D_{KL}(\nabla\psi_{\#}P\|Q) \mid \psi \text{ is convex}\} \tag{5}$$

Proof. Suppose φ is the unique (up to constants) solution to problem 3. Hence φ is convex and we have $\nabla\varphi_{\#}P = Q$. Substituting into the Kullback-Leibler divergence, we obtain

$$D_{\mathrm{KL}}(\nabla\varphi_{\#}P\|Q) = \int \ln(1)\ dP = 0$$

Since $D_{\mathrm{KL}}(P\|Q)$ is always non-negative, $\nabla\varphi$ indeed satisfies (5). For the other half of the proof, we note that

$$0 \leq \inf_{\psi} \{D_{\mathrm{KL}}(\nabla\psi_{\#}P\|Q) \mid \psi \text{ is convex}\} \leq D_{\mathrm{KL}}(\nabla\varphi_{\#}P\|Q) = 0$$

Thus, any convex function ξ satisfying (5) must have $D_{\mathrm{KL}}(\nabla\xi_{\#}P\|Q) = 0$. Hence, by Lemma 3.1, $\nabla\xi_{\#}P = Q$. Finally, by uniqueness of solutions to problem 3, $\xi = \varphi$ (up to a constant). □

3.2 Variational Formulation

The theorem above provides a means to solve problem 3 in the case when P and Q are known. However, in the problem we seek to solve, P and Q are known only through samples. Therefore, it becomes impractical to calculate $D_{\mathrm{KL}}(\nabla\psi_{\#}P\|Q)$. Nevertheless, the Kullback-Leibler divergence has a variational formulation which can be written as the maximization of the difference of two expectations, which is naturally adaptable to the sample based problem.

Theorem 3.2.

$$D_{KL}(P\|Q) = 1 + \max_{g}\left\{\int g\ dP - \int e^{g}\ dQ \;\middle|\; g \text{ Borel measurable}\right\} \tag{6}$$

Proof.

1. If P is not absolutely continuous with respect to Q, then there exists some set A such that $P(A) > 0$ and $Q(A) = 0$. Consider the family of indicator functions

$$k\chi_A := \begin{cases} 1 \text{ if } x \in A \\ 0 \text{ if } x \notin A \end{cases}$$

where $k \in \mathbf{R}$. It follows that

$$kP(A) \leq \max_g \left\{ \int g \, dP - \int e^g \, dQ \;\middle|\; g \text{ Borel measurable} \right\}$$

Taking the limit as $k \to +\infty$, we obtain

$$\max_g \left\{ \int g \, dP - \int e^g \, dQ \;\middle|\; g \text{ Borel measurable} \right\} = +\infty = D_{\mathrm{KL}}(P\|Q)$$

2. If P is not absolutely continuous with respect to Q, the function

$$g \to g\frac{dP}{dQ} - e^g$$

is concave in g, and is uniquely maximized at $g = \ln(dP/dQ)$. From this we have that for every $x \in \mathbf{R}^d$,

$$g(x)\frac{dP}{dQ}(x) - e^{g(x)} \leq \frac{dP}{dQ}(x)\left(\ln\left(\frac{dP}{dQ}(x)\right) - 1\right)$$

Integrating with respect to Q yields

$$1 + \int g \, dP - \int e^g \, dQ \leq \int \ln\left(\frac{dP}{dQ}\right) \, dP = D_{\mathrm{KL}}(P\|Q)$$

Taking the maximum over all g yields the final result.

□

Assuming that $X \sim P$ and $Y \sim Q$, we can replace the integral with expectations; (3.2) becomes

$$D_{\mathrm{KL}}(P\|Q) = 1 + \max_g \left\{ \mathbb{E}\left[g(X)\right] - \mathbb{E}\left[e^{g(Y)}\right] \;\middle|\; g \text{ Borel measurable} \right\}$$

Suppose $\{x_j\}$ is a sample drawn from P and $\{y_j\}$ is drawn from Q. Replacing the expectations by empirical means yields the approximation

$$D_{\mathrm{KL}}(P\|Q) \approx \max_g \left\{ \frac{1}{n}\sum_i g(x_i) - \frac{1}{m}\sum_j e^{g(y_j)} \right\}$$

This yields a clean way to estimate the value of the Kullback-Leibler divergence from samples of P and Q, without knowing their underlying distributions.

3.3 Minmax

We now construct a theoretical minmax problem equivalent to problem 3. Suppose that $X \sim P$ and $Y \sim Q$. For any map T, it is known that $X \sim P \Rightarrow T(X) \sim Q$ (see [Villani, 2003] for proof). From Theorem 3.2 we have

$$D_{\mathrm{KL}}(\nabla\psi_{\#}P\|Q) = 1 + \max_{g}\left\{\mathbb{E}[g(\nabla\psi(X))] - \mathbb{E}\left[e^{g(Y)}\right] \;\middle|\; g \text{ Borel measurable}\right\}$$

By applying Theorem 3.1, φ is a solution to problem 3 if and only if φ satisfies

$$\varphi = \arg\min_{\psi}\{D_{\mathrm{KL}}(\nabla\psi_{\#}P\|Q) \mid \psi \text{ is convex}\}$$

$$= \arg\min_{\psi}\left\{\max_{g}\left\{\mathbb{E}[g(\nabla\psi(X))] - \mathbb{E}\left[e^{g(Y)}\right] \;\middle|\; g \text{ Borel measurable}\right\} \;\middle|\; \psi \text{ is convex}\right\}$$

If $\{x_j\}$ is a sample drawn from P and $\{y_j\}$ is drawn from Q, replacing the expectations with empirical means gives us

$$\varphi \approx \arg\min_{\psi}\left\{\max_{g}\left\{\frac{1}{n}\sum_{i} g(\nabla\psi(x_i)) - \frac{1}{m}\sum_{j} e^{g(y_j)}\right\}\right\} \tag{7}$$

4 Algorithm

4.1 Minimax Algorithm

We now take the theoretical problem (7) and develop an algorithm to solve it given two samples $\{x_j\}$ and $\{y_j\}$. One naive approach to solve a minmax problem would be to preform a 'twisted' gradient descent, i.e. ascent in g-space and gradient descent in ψ-space. Parametrising ψ-space by α and g-space by β, we can define the objective function f by

$$f[\alpha,\beta] = \frac{1}{n}\sum_{i} g(\nabla\psi(x_i;\alpha);\beta) - \frac{1}{m}\sum_{j} e^{g(y_j;\beta)} \tag{8}$$

the iterative scheme would be

$$\begin{bmatrix} \alpha_{n+1} \\ \beta_{n+1} \end{bmatrix} = \begin{bmatrix} \alpha_n \\ \beta_n \end{bmatrix} - \Delta t \cdot KG_n$$

Where

$$G_n := \begin{bmatrix} \nabla_\alpha f \\ \nabla_\beta f \end{bmatrix}_{[\alpha_n,\beta_n]}, \qquad K := \begin{bmatrix} I_\alpha & 0 \\ 0 & -I_\beta \end{bmatrix} \tag{9}$$

I_v is the square identity matrix, with dimensions equal to the length of the vector v. Whilst effective in some cases, this algorithm is prone to be caught in

oscillatory solutions. To rectify this we may instead apply an implicit 'twisted' gradient descent. The iterative scheme would instead be

$$\begin{bmatrix} \alpha_{n+1} \\ \beta_{n+1} \end{bmatrix} = \begin{bmatrix} \alpha_n \\ \beta_n \end{bmatrix} - \Delta t \cdot K G_{n+1} \tag{10}$$

Using a first order expansion of G_{n+1} about (α_n, β_n), we have

$$G_{n+1} \approx G_n + H_n \left(\begin{bmatrix} \alpha_{n+1} \\ \beta_{n+1} \end{bmatrix} - \begin{bmatrix} \alpha_n \\ \beta_n \end{bmatrix} \right)$$

Where

$$H_n = \begin{bmatrix} \nabla_{\alpha\alpha} f & \nabla_{\alpha\beta} f \\ \nabla_{\alpha\beta} f & \nabla_{\beta\beta} f \end{bmatrix}_{[\alpha_n, \beta_n]} \tag{11}$$

Substituting this into (10), we obtain a linear problem for $(\alpha_{n+1}, \beta_{n+1})$ in terms of $(\alpha_{n+1}, \beta_{n+1})$.

$$(I + \Delta t \cdot K H_n) \begin{bmatrix} \alpha_{n+1} \\ \beta_{n+1} \end{bmatrix} = \begin{bmatrix} \alpha_n \\ \beta_n \end{bmatrix} - \Delta t \cdot K \left(G_n - H_n \begin{bmatrix} \alpha_n \\ \beta_n \end{bmatrix} \right) \tag{12}$$

Solving for $[\alpha_{n+1}, \beta_{n+1}]$ yields an iterative scheme for an implicit twisted gradient descent algorithm.

Algorithm 1. *(Minmax)*
Input$(x, y, \alpha_0, \beta_0)$
Define: G, H, K as in (9), and (11)
$G_0 := G(\alpha_0, \beta_0)$
$H_0 := H(\alpha_0, \beta_0)$
$n = 0$
while $\|G_n\| > E$
  $v =$ *linsolve(*$(I + (\Delta t) K H_n) v = (\alpha_n, \beta_n) - (\Delta t) K (G_n - H_n \cdot (\alpha_n, \beta_n))$*)*
  $\alpha_{n+1} := v(1)$
  $\beta_{n+1} := v(2)$
  $G_{n+1} := G(\alpha_{n+1}, \beta_{n+1})$
  $H_{n+1} := H(\alpha_{n+1}, \beta_{n+1})$
  $n = n + 1$
end
Output$(\alpha_{n+1}, \beta_{n+1})$

Notice that the in limit as $\Delta t \to +\infty$, (12) becomes Newton's method. It is possible to accelerate convergence by dynamically modifying the value of Δt, further details can be found in [Essid et al., 2019].

4.2 Global Algorithm

Implicit twisted gradient descent is not guaranteed to converge to the correct solution, unless the measures P and Q are close in proximity. One way to circumvent this problem is to break the global problem into N local optimal

transport problems, composing them to obtain the global map. By Brenier's theorem, this composition will only be optimal if it is itself the gradient of a convex scalar function which satisfies the push forward condition.

Lemma 4.1. *For any measure P,*

$$(\varphi \circ \psi)_{\#} P = \varphi_{\#}(\psi_{\#} P)$$

Proof.

$$\begin{aligned}(\varphi \circ \psi)_{\#} P &= P((\varphi \circ \psi)^{-1}(A)) \\ &= P\left(\psi^{-1}\left(\varphi^{-1}(A)\right)\right) \\ &= (\psi_{\#} P)\left(\varphi^{-1}(A)\right) \\ &= \varphi_{\#}\left(\psi_{\#} P\right)\end{aligned}$$

for any Borel subset $A \subset \mathbf{R}^d$ □

Suppose P and Q are probability measures on $\mathbf{R}^d$. By Brenier's theorem, there exists a unique convex scalar function $\psi : \mathbf{R}^d \to \mathbf{R}$ which satisfies $\nabla\psi_{\#} P = Q$. For some $N \in \mathbf{N}$, define $Z_n = \left(\frac{n}{N}\nabla\psi + (1 - \frac{n}{N})\mathbb{1}\right)_{\#} P$ for $0 \leq n \leq N$, where $\mathbb{1}$ denotes the identity map. These distributions are called McCann interpolants. For a sufficiently large N, the distribution $P = Z_0$ will be close enough to Z_1 such that the minmax algorithm will converge to the local optimal map between them. For every $1 \leq n \leq N$, let φ_n be the convex scalar function such that $(\nabla\varphi_n)_{\#} Z_{n-1} = Z_n$. Denote $\circ_{n=1}^{k} \nabla\varphi_n = \nabla\varphi_k \circ \nabla\varphi_{k-1} \circ \cdots \circ \nabla\varphi_1$; It follows from Lemma 4.1 that $T := \circ_{n=1}^{N} \nabla\varphi_n$ satisfies the pushforward condition. It remains to be shown that T is the gradient of a convex scalar function.

Notice that by Lemma 4.1,

$$\left(\bigcirc_{n=1}^{k} \nabla\varphi_n\right)_{\#} Z_0 = Z_k = \left(\frac{k}{N} T + \left(1 - \frac{k}{N}\right)\mathbb{1}\right)_{\#} Z_0$$

Hence

$$\bigcirc_{n=1}^{k} \nabla\varphi_n = \frac{k}{N} T + \left(1 - \frac{k}{N}\right)\mathbb{1} \text{ almost everywhere, for all } 1 \leq k \leq N \tag{13}$$

It turns out that this condition is sufficient to prove that T is convex.

Lemma 4.2. *Suppose $\psi, \varphi : \mathbf{R^d} \to \mathbf{R}$ are convex scalar functions. Define $T = \nabla\psi \circ \nabla\varphi$. If*

$$\nabla\varphi = \alpha T + \beta\mathbb{1} \quad \text{for any } \alpha, \beta \in \mathbf{R} \tag{14}$$

Then, there exists a convex scalar function $\tau : \mathbf{R}^d \to \mathbf{R}$ such that $T = \nabla\tau$

Proof. Rearranging (14) we obtain for any $x \in \mathbf{R}^d$

$$\begin{aligned}T(x) &= \frac{1}{\alpha}\Big(\nabla\varphi(x) - \beta\mathbb{1}(x)\Big) \\ &= \frac{1}{\alpha}\nabla\left(\varphi(x) - \beta\frac{\|x\|^2}{2}\right)\end{aligned}$$

We have $T(x) = \nabla\tau(x)$, where $\tau : \mathbf{R}^d \to \mathbf{R}$ is defined by $\left(2\varphi(x) - \beta\|x\|^2\right)/2\alpha$. It remains to be shown that τ is convex. Let $H(\cdot)$ denote the Hessian. Since $\nabla\tau = \nabla\psi \circ \nabla\varphi$, we have

$$H(\tau)\big|_x = H(\psi)\big|_{\varphi(x)} H(\varphi)\big|_x$$

But $H(\cdot)$ is a symmetric matrix, so it follows that

$$H(\psi)\big|_{\varphi(x)} H(\varphi)\big|_x = H(\varphi)\big|_x H(\psi)\big|_{\varphi(x)}$$

Commuting symmetric matrices share a basis of orthonormal eigenvectors (see [Lax, 2007]). For $H(\varphi)$ and $H(\psi)$ we denote this basis by $\{v_i\}$, with corresponding eigenvalues λ_i^φ and λ_i^ψ respectively. Since φ and ψ are convex, $H(\varphi)$ and $H(\psi)$ are positive semi-definite and therefore $\lambda_i^\varphi, \lambda_i^\psi \geq 0$. For any $w \in \mathbf{R}^d$, we have the expansion

$$w = \sum_{i=1}^{d} \alpha_i v_i$$

which implies

$$w^T H(\tau) w = \sum_{i=1}^{d} \lambda_i^\varphi \lambda_i^\psi \alpha_i^2 v_i^2 \geq 0$$

Hence $H(\tau)$ is positive semi-definite, verifying that τ is convex $\square$

From (13) we obtain the interative relation

$$\overset{k-1}{\underset{n=1}{\circ}} \nabla\varphi_n = \left(\frac{k-1}{k}\right) \overset{k}{\underset{n=1}{\circ}} \nabla\varphi_n + \left(\frac{N+1-k}{Nk}\right) \mathbb{1}$$

For all $1 < k \leq N$. Then, using Lemma 4.2 we may inductively show that $\circ_{n=1}^{k} \nabla\varphi_n$ is a gradient of a convex scalar function for all $1 < k \leq N$. Hence T is indeed the optimal solution to the transport problem.

In practice, however, in order to obtain the the McCann interpolants Z_n, we require a priori knowledge of T. To address this issue, we generate an initial T^0, which we use to find approximate McCann interpolants Z_n^0. Then, using the minmax algorithm described above, we can compute the locally optimal maps φ_n^0, which we then compose to construct a better approximation

$$T^1 := \overset{N}{\underset{n=1}{\circ}} \nabla\varphi_n^0$$

of T. Iterating this process we obtain a sequence of transport maps T^k which, if the converge, approximate the actual solution T. We obtain the following algorithm.

Algorithm 2. *[Theoretical Global Optimal Transport]*
Input(Z_0,Z_N)
construct an initial T^0
$k = 0$
while $\|T^{k-1} - T^k\| > E$
  for $1 \leq n < N$
    $Z_n^k = (\frac{n}{N}T^k + (1 - \frac{n}{N})\mathbb{1})_{\#}P$
    φ_n^k = *minmax*(Z_{n-1},Z_n)
  end
  $T^{k+1} = \circ_{n=1}^{N} \nabla\varphi_n^k$
  $k = k + 1$
end

4.3 Sample based algorithm

In the sample based algorithm, we are not afforded the luxury of knowing the initial and final distributions exactly, rather we are given samples $\{x_j\}$ and $\{y_j\}$, which capture their essence. For each intermediate distribution Z_n, we seek to obtain a sample which is drawn from it. Since $T_{\#}P = Q$, we have $T(X) \sim Q$. Hence the set $(\frac{n}{N}T + (1 - \frac{n}{N})\mathbb{1})(\{x_j\})$ is a sample drawn from Z_n. We modify the algorithm accordingly

Algorithm 3. *[Sample Based Global Optimal Transport]*
Input(x,y)
construct an initial T^0
$\left[s^0\right]_j = T^0([x]_j)$
$k = 0$
while $k == 0$ *or* $\|s^{k-1} - s^k\| > E$
  for $1 \leq n \leq N$
    Define $[z^n]_j := \left[\frac{n}{N}s^k + (1 - \frac{n}{N})x\right]_j$
  end
  for $1 \leq n \leq N$
    φ_n^k = *minmax*(z^{n-1},z^n)
    $[z^n]_j = \nabla\varphi_n^k(\left[z^{n-1}\right]_j)$
  end
  $\left[s^{k+1}\right]_j = \left[z^N\right]_j$
  $k = k + 1$
end

5 Example

We now apply algorithm 3 to solve a sample based optimal transport problem: Consider the space $\mathbf{R}^2$, and let $\{x_j\}$ and $\{y_j\}$ be samples drawn from unknown

distributions on $\mathbf{R}^2$ (see figure 1). Before proceeding, we must address several ambiguities glossed over when describing the algorithm.

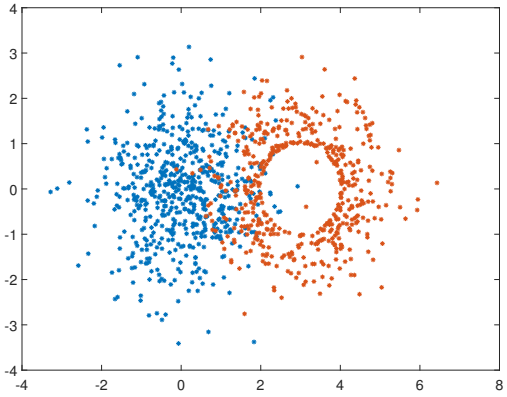

Figure 1: The blue sample, with 600 points is drawn from a Gaussian distribution; the orange sample with 500 points is drawn from a Gaussian with a hole punched through it.

5.1 Choosing T_0

In algorithms 2 and 3, it is left unspecified how to generate an initial guess T_0. In the theoretical algorithm, provided that $(T_0)_{\#}P = Q$ we have $(T_n)_{\#}P = Q$ for all $n \in \mathbf{N}$. On the other hand, if P and Q are unknown, it becomes unclear how to find a suitable T_0. If the samples $\{x_j\}$ and $\{y_j\}$ are of the same size, then any one to one correspondence $T_0 : \{x_j\} \to \{y_j\}$ would satisfy the pushforward condition. In our example however, it is not assumed that both samples are of the same length.

In such a case, we fabricate the initial (when $k = 0$) interpolating samples $\{z_j^0\}$ directly, in such a way as to make the transition from $\{x_j\}$ to $\{y_j\}$ is as smooth as possible. We then define T_0 to be the composition $\circ_{n=1}^{N} \nabla\varphi_n$ of all the local optimal maps $\nabla\varphi_n$ from $\{z_j^{n-1}\}$ to $\{z_j^n\}$. Even if these fabricated samples are not McCann interpolants, in subsequent iterations ($k \geq 1$) of the algorithm, z_j^n are, by definition, McCann interpolants. T_0 is also guaranteed to satisfy the pushforward condition $T_{\#}P = Q$, provided the transition from $\{x_j\}$ to $\{y_j\}$ through $\{z_j^n\}$ is smooth enough such that the local maps converge to the correct solution.

Given two Gaussian distributions, $P = \mathcal{N}(\bar{x}, \Sigma_x)$ and $Q = \mathcal{N}(\bar{y}, \Sigma_y)$ (where $\bar{x}, \bar{y}$ are their respective means and Σ_x, Σ_y their respective covariant matrices), the optimal map from P to Q has an explicit form, and is given by

$$F_x(x) := \bar{y} + \Sigma_x^{1/2}(\Sigma_x^{-1/2}\Sigma_y\Sigma_x^{-1/2})^{1/2}\Sigma_x^{1/2}(x - \bar{x})$$

The optimal map from Q to P is

$$F_y(y) := \bar{x} + \Sigma_y^{1/2}(\Sigma_y^{-1/2}\Sigma_x\Sigma_y^{-1/2})^{1/2}\Sigma_y^{1/2}(y - \bar{y})$$

For each $1 \leq n \leq N$, $\{z_j^n\}$ is defined to be a 'blend' of the sets $\{\frac{n}{N}x_j + (1 - \frac{n}{N})F_x(x_j)\}$ and $\{\frac{n}{N}F_y(y_j) + (1 - \frac{n}{N})y_j\}$; $\lfloor \frac{n}{N}\#(\{x_j\})\rfloor$ and $\lfloor(1 - \frac{n}{N})\#(\{x_j\})\rfloor$ elements are chosen from each set respectively (where $\#(A)$ is the cardinality of A). See figure 2 to see this method used on our example.

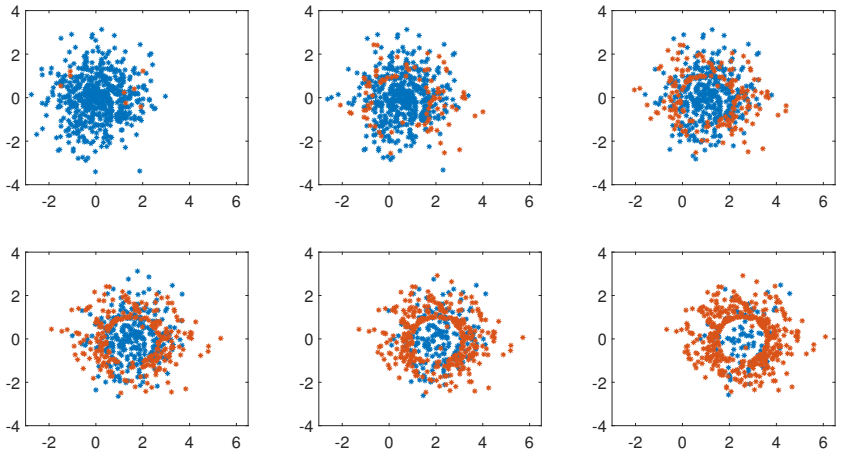

Figure 2: 6 different initial samples $\{z_j^n\}$ obtained by the method described above. The blue points belong to the set $\{\frac{n}{N}x_j + (1 - \frac{n}{N})F_x(x_j)\}$ and the orange points belong to the set $\{\frac{n}{N}F_y(y_j) + (1 - \frac{n}{N})y_j\}$

5.2 Parametrisation of the function spaces

The minmax algorithm requires parametrisations of both ψ-space and g-space. Given two between nearby distributions, the optimal map $\nabla\varphi$ between them will be a small perturbation of the identity map. One can argue that $\nabla\varphi$ therefore belongs to a low dimensional function space, which contains Gaussian perturbations from the set of linear functions. In particular, we parametrise ψ-space by

$$\psi(x;\alpha) = x^T \begin{bmatrix} \alpha_1 & \alpha_2 \\ \alpha_2 & \alpha_3 \end{bmatrix} x + \begin{bmatrix} \alpha_4 & \alpha_5 \end{bmatrix} x + \sum_{k=0}^{m-1} \alpha_{1+3k} e^{-\frac{\|x-(\alpha_{2+3k},\alpha_{3+3k})\|^2}{\sigma^2}}$$

Where σ is a constant, which, intuitively speaking, is the resolution of the perturbations. We chose g-space to be similar to ψ-space

$$g(y;\beta) = y^T \begin{bmatrix} \beta_1 & \beta_2 \\ \beta_2 & \beta_3 \end{bmatrix} y + \begin{bmatrix} \beta_4 & \beta_5 \end{bmatrix} y + \beta_5 + \sum_{k=0}^{m-1} \beta_{1+3k} e^{-\frac{\|y-(\beta_{2+3k},\beta_{3+3k})\|^2}{\sigma^2}}$$

5.3 Choosing α_0 and β_0

The minmax algorithm requires initial conditions α_0 and β_0. Since the minmax algorithm is used to find the the optimal map between nearby distributions, $\nabla\psi(x;\alpha_0) = \mathbb{1}$ is a suitable initial condition. From (7) the corresponding β_0 should be chosen such that

$$\beta_0 = \arg\max_{\beta} \left\{ \frac{1}{n}\sum_i g(x_i;\beta) - \frac{1}{m}\sum_j e^{g(y_j,\beta)} \right\}$$

5.4 Result

Inputting our example into algorithm 3, we obtain the optimal map between both samples, shown in figures 3.

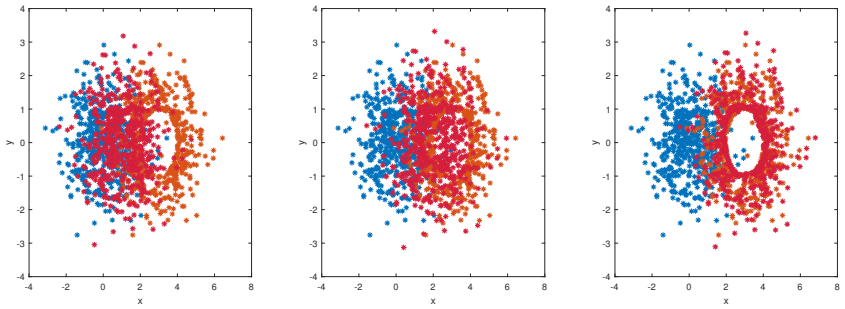

Figure 3: The blue and orange points are the initial and target samples respectively. The red points are some of the McCann interpolants of final T_k returned by algorithm 3, the rightmost distribution being the the image of $\{x_j\}$ under T_k.

6 Conclusion

In the example presented, it is possible to obtain an approximation of the underlying optimal map. The current iteration of the code is not robust enough to be suitable for applications in change detection. This is in part due to the lack of complexity of the function space. Improve this algorithm, one might want to consider a parametrisation of the entirety of ψ-space and g-space to improve the complexity of local maps. Furthermore one might introduce a penalty term to enforce convexity. In this algorithm, no such penalty was needed, since N was large enough to ensure that the interpolants were sufficiently close. This however, came at cost of computational efficiency.

References

[Essid et al., 2018] Essid, M., Laefer, D., and Tabak, E. G. (2018). Adaptive optimal transport. *arXiv preprint arXiv:1807.00393*.

[Essid et al., 2019] Essid, M., Tabak, E., and Trigila, G. (2019). An implicit gradient-descent procedure for minimax problems. *arXiv preprint arXiv:1906.00233*.

[Galichon, 2018] Galichon, A. (2018). *Optimal transport methods in economics*. Princeton University Press.

[Kuang and Tabak, 2019] Kuang, M. and Tabak, E. G. (2019). Sample-based optimal transport and barycenter problems. *Communications on Pure and Applied Mathematics*, 72(8):1581–1630.

[Lax, 2007] Lax, P. (2007). *Linear Algebra and Its Applications*. Pure and Applied Mathematics: A Wiley Series of Texts, Monographs and Tracts. Wiley.

[McCann, 1997] McCann, R. J. (1997). A convexity principle for interacting gases. *Advances in mathematics*, 128(1):153–179.

[Villani, 2003] Villani, C. (2003). *Topics in optimal transportation*. Number 58. American Mathematical Soc.