

Solving Monge Problem by Hilbert Space Embeddings of Probability Measures

Takafumi Saito^a and Yumiharu Nakano^b

Department of Mathematical and Computing Science, Institute of Science Tokyo, Tokyo, Japan

Keywords: Monge Problem, Maximum Mean Discrepancy, Deep Learning.

Abstract: We propose deep learning methods for classical Monge's optimal mass transportation problems, where the distribution constraint is treated as a penalty term defined by the maximum mean discrepancy in the theory of Hilbert space embeddings of probability measures. We prove that the transport maps given by the proposed methods converge to optimal transport maps in the problem with L^2 cost. Several numerical experiments validate our methods. In particular, we show that our methods are applicable to large-scale Monge problems.

1 INTRODUCTION

Our aim in this paper is to propose numerical methods for Monge's mass transportation problem, described as follows: given two Borel probability measures μ, ν on \mathbb{R}^d and a cost function $c : \mathbb{R}^d \times \mathbb{R}^d \rightarrow [0, \infty]$, to minimize Our problem is to minimize

$$M(T) := \int_{\mathbb{R}^d} c(x, T(x)) d\mu(x)$$

over all Borel measurable mapping $T : \mathbb{R}^d \rightarrow \mathbb{R}^d$ such that $\mu \circ T^{-1} = \nu$. Here, $\mu \circ T^{-1}$ denotes the pushforward of μ with T , i.e., $\mu \circ T^{-1}(A) = \mu(T \in A)$ for any Borel set A .

We shall briefly describe the background of the Monge problem. Monge problem was proposed by Gaspagnol Monge in 1781 (Monge, 1781). In the 20th century, this problem was expanded by Kantorovich to make it mathematically easier to handle, and is now called the Monge-Kantorovich problem (Kantorovich, 1942), (Kantorovich, 1948).

The Monge problem, as extended by Kantorovich, is called the Monge-Kantorovich problem. The first existence and uniqueness result was established by Brein, (Breiner, 1987), (Breiner, 1991). Gangbo and McCann further developed the generalized problem (Gangbo and McCann, 1996). Mikami also provided a probabilistic proof of Breiner's result (Mikami, 2004).

Algorithms for solving the Monge-Kantorovich problem can be traced back nearly 100 years (Tolstoi,

1930). Since the advent of mathematical programming, these algorithms have been a field of significant interest (Dantzig, 1951). This is largely because Dantzig's initial motivation for developing mathematical programming was related to solving transportation problems (Dantzig, 1949), and it was later discovered that optimal transportation problems and minimum cost flow problems are equivalent (Korte and Vygen, 2012). Research has advanced in the development of solvers that apply mathematical programming to machine learning (Bauschke and Combettes, 2011), as well as in solvers used in dynamic optimal transport (OT) (Papadakis et al., 2014). Today, it remains one of the most actively researched fields.

We describe recent solvers developed to address the Monge-Kantorovich problem. In recent years, the Monge-Kantorovich problem has been actively applied in machine learning. One prominent approach involves solving the entropy optimization problem by introducing an entropy regularization term, using the Sinkhorn algorithm (Peyré and Cuturi, 2019). The Python library POT (Flamary et al., 2021) provides a solver for various optimal transport problems, including entropy-regularized optimal transport.

When solving the optimal transport problem, computations can scale with the cube of the input data size. Therefore, it is crucial to support large-scale computations and to leverage GPUs for numerical calculations. Optimal transport solvers designed for large-scale computations include the Python library Geomloss (Feydy et al., 2019). Another tool, the Python library KeOps (Charlier et al., 2021), opti-

^a <https://orcid.org/0009-0007-0598-0238>

^b <https://orcid.org/0000-0002-5723-4444>

mizes the reduction of large arrays using neural networks and kernel formulas.

Additionally, the OTT library (Cuturi et al., 2022), based on the high-performance JAX library for numerical calculations, offers implementations that can solve problems such as the Monge-Kantorovich problem.

Our aim is to derive a numerical solution for the basic mathematical analysis of the original Monge problem, rather than the Monge-Kantorovich problem. In particular, the goal is to develop algorithms capable of performing GPU-based numerical computations and handling large-scale calculations. The method for deriving the numerical solution primarily utilizes the embedding theory of probability measures, which was previously applied by Nakano (Nakano, 2024) to obtain a numerical solution for the Schrödinger bridge problem. In this study, we also apply this theory to derive a numerical solution for the Monge problem. The penalty method is employed to find the optimal solution, with the use of Maximum Mean Discrepancy (MMD) as the penalty function being a novel approach. Unlike existing methods, this method is independent of input data size. We also confirm this through numerical experiments.

This paper is structured as follows. In the next section, we review some basic results on the theory of Hilbert space embedding of probability measures and describe a numerical method that describes a theoretical approximation solution to the Monge problem with L^2 cost. Section 3 gives numerical experiments.

2 PENALTY METHOD

2.1 Hilbert Space Embeddings of Probability Measures

We shall give a quick review of theory of Hilbert space embeddings of probability measures, as developed in Sriperumbudur et al. (Sriperumbudur et al., 2010). Denote by $\mathcal{P}(\mathbb{R}^d)$ the set of all Borel probability measures on \mathbb{R}^d . Let K be a symmetric and strictly positive definite kernel on \mathbb{R}^d , i.e., $K(x, y) = K(y, x)$ for $x, y \in \mathbb{R}^d$ and for any positive distinct $x_1, \dots, x_N \in \mathbb{R}^d$ and $\alpha = (\alpha_1, \dots, \alpha_N)^T \in \mathbb{R}^N \setminus \{0\}$,

$$\sum_{j, \ell=1}^N \alpha_j \alpha_\ell K(x_j, x_\ell) > 0.$$

Assume further that K is bounded and continuous on $\mathbb{R}^d \times \mathbb{R}^d$. Then, there exists a unique Hilbert space \mathcal{H} such that K is a reproducing kernel on \mathcal{H} with norm $\|\cdot\|$ (Wendland, 2010). Then consider the *maximum*

mean discrepancy (MMD) γ_K defined by, for $\mu_0, \mu_1 \in \mathcal{P}(\mathbb{R}^d)$,

$$\begin{aligned} \gamma_K(\mu_0, \mu_1) &:= \sup_{f \in \mathcal{H}, \|f\| \leq 1} \left| \int_{\mathbb{R}^d} f d\mu_0 - \int_{\mathbb{R}^d} f d\mu_1 \right| \\ &= \left\| \int_{\mathbb{R}^d} K(\cdot, x) \mu_0(dx) \right. \\ &\quad \left. - \int_{\mathbb{R}^d} K(\cdot, x) \mu_1(dx) \right\|_{\mathcal{H}} \end{aligned}$$

It is known that if K is an integrally strictly positive definite then γ_K defines a metric on $\mathcal{P}(\mathbb{R}^d)$. Examples of integrally strictly positive definite kernels include the Gaussian kernel $K(x, y) = e^{-\alpha|x-y|^2}$, $x, y \in \mathbb{R}^d$, where $\alpha > 0$ is a constant, and the Matérn kernel $K(x, y) = K_\alpha(|x-y|)$, $x, y \in \mathbb{R}^d$, where K_α is the modified Bessel function of order $\alpha > 0$. It is also known that Gaussian kernel as well as Matérn kernel metrize the weak topology on $\mathcal{P}(\mathbb{R}^d)$ (Sriperumbudur et al., 2010), (Nakano, 2024).

Define

$$\begin{aligned} K_1(x, y) &:= K(x, y) - \int_{\mathbb{R}^d} K(x, y') \mu_1(dy') \\ &\quad - \int_{\mathbb{R}^d} K(x', y) \mu_1(dx'). \end{aligned}$$

Then,

$$\begin{aligned} \gamma_K(\mu_0, \mu_1)^2 &= \int_{\mathbb{R}^d} \int_{\mathbb{R}^d} K_1(x, y) \mu_0(dx) \mu_0(dy) \\ &\quad + \int_{\mathbb{R}^d \times \mathbb{R}^d} K(x', y') \mu_1(dx') \mu_1(dy'). \end{aligned} \quad (1)$$

2.2 Proposed Methods

Let $\gamma = \gamma_K$ be as in Section 2.1. We proposed a penalty method for Monge problem by Hilbert space embeddings of probability measures, described as follows:

$$\inf_T M_\lambda(T) := \inf_T \{M(T) + \lambda \gamma^2(\mu \circ T^{-1}, \nu)\}.$$

Note that the second term penalizes the distance between the laws of $T(x)$ and ν . Moreover, the second term in the above formula can be expressed discretely as follows: given IID samples $X_1, \dots, X_M \sim \mu_0$ and $Y_1, \dots, Y_M \sim \mu_1$, an unbiased estimator of $\gamma_K(\mu_0, \mu_1)^2$ is given by

$$\begin{aligned} \tilde{\gamma}_K(\mu_0, \mu_1)^2 &:= \frac{1}{M(M-1)} \sum_i \sum_{j \neq i} K(X_i, X_j) \\ &\quad - \frac{2}{M^2} \sum_{i, j} K(X_i, Y_j) \\ &\quad + \frac{1}{M(M-1)} \sum_i \sum_{j \neq i} K(Y_i, Y_j) \end{aligned} \quad (2)$$

(Gretton et al., 2006). Then, we approximate T by a class $\{T_\theta\}_{\theta \in \Theta}$ of deep neural networks. Each T_θ can be given by a multilayer perception with input layer g_0 , $L-1$ hidden layer g_1, \dots, g_{L-1} , and output layer g_L , where $L \geq 1$ and for $\xi \in \mathbb{R}^{1+m}$, $g_0(\xi) = \xi$, $g_\ell(\xi) = \phi_{\ell-1}(w_\ell g_{\ell-1}(\xi) + \beta_\ell) \in \mathbb{R}^{m_\ell}$, $\ell = 1, \dots, L$, for some matrices w_ℓ and vectors β_ℓ , $\ell = 1, \dots, L$. Here m_ℓ denotes the number of units in the layer ℓ , and $\phi_{\ell-1}$ is an activation function. Then the parameter θ is described by $\theta = (w_\ell, \beta_\ell)_{\ell=1, \dots, L}$ and $T_\theta(x) = g_L(x) \in \mathbb{R}^d$. For $\lambda > 0$, the integral term in (1) is replaced by (2) and $T_\theta(x)$ as follows: by Subsetcion 3.2. in (Nakano, 2024),

$$F_1(\theta) = \frac{1}{\lambda M} \sum_{i=1}^M C(X_i, T_\theta(X_i)) + \frac{1}{M(M-1)} \sum_i \sum_{j \neq i} K(T_\theta(X_i), T_\theta(X_j)) - \frac{2}{M^2} K(T_\theta(X_i), Y_j) \quad (3)$$

The algorithm is described below and we test our one algorithm thorough a numerical experiment.

Algorithm 1: Deep learning algorithm with empirical MMD.

Data: The number n of the iterations, the batch size M , weight parameter $\lambda > 0$
Result: transport map T
 initialization
 $X_1, \dots, X_M \leftarrow$ IID samples from μ ,
 $Y_1, \dots, Y_M \leftarrow$ IID samples from ν .
for $k=1, 2, \dots, n$ **do**
 Compute $F_1(\theta)$ in (3) using $\{X_j, Y_j\}_{j=1}^M$.
 Take the gradient step on $\nabla_\theta F_1(\theta)$.

2.3 Theoretical Result

For given $\lambda > 0$ consider the minimization problem

$$M_\lambda(T) := M(T) + \lambda \gamma^2(\mu \circ T^{-1}, \nu)$$

over all Borel measurable mappings T . Take arbitrary positive sequences $\{\varepsilon_n\}_{n=1}^\infty$ and $\{\lambda_n\}_{n=1}^\infty$ such that

$$\varepsilon_n \rightarrow 0, \quad \lambda_n \rightarrow +\infty \quad (n \rightarrow \infty).$$

Then take $T_n : \mathbb{R}^d \rightarrow \mathbb{R}^d$ such that

$$M_{\lambda_n}(T_n) - \varepsilon_n \leq M_{\lambda_n}^* := \inf_T M_{\lambda_n}(T).$$

Then we have the following:

Theorem 2.1. Let $c(x, y) = |x - y|^2$. Suppose that μ is absolutely continuous with respect to the Lebesgue measure and that

$$\int_{\mathbb{R}^d} |x|^2 \mu(dx) + \int_{\mathbb{R}^d} |y|^2 \nu(dy) < \infty.$$

Suppose moreover that γ metrizes the weak topology on $\mathcal{P}(\mathbb{R}^d)$. Then,

$$\lim_{n \rightarrow \infty} \sqrt{\lambda_n} \gamma(\mu \circ T_n^{-1}, \nu) = 0, \quad (4)$$

$$\lim_{n \rightarrow \infty} M_{\lambda_n}(T_n) = M(T^*), \quad (5)$$

where T^* is the unique optimal transport map. In particular, $\{T_n\}_{n=1}^\infty$ converges to T^* in law under μ .

Proof. First, note that under our assumption, an optimal transport map does exist uniquely (see, e.g., Theorem 2.12 in (Villani, 2021) and Theorem 1.25 in (Santambrogio, 2015)).

Step (i). We will show (4). This claim can be proved by almost the same argument as that given in the proof of Theorem 3.1 in (Nakano, 2024), but we shall give a proof for reader's convenience. Assume contrary that

$$\limsup_{n \rightarrow \infty} \lambda_n \gamma^2(\mu \circ T_n^{-1}, \nu) = 3\delta$$

for some $\delta > 0$. Then there exists a subsequence $\{n_k\}$ such that

$$\lim_{k \rightarrow \infty} \lambda_{n_k} \gamma^2(\mu \circ T_{n_k}^{-1}, \nu) = 3\delta.$$

Since γ is a metric, we have $\gamma(\mu \circ (T^*)^{-1}, \nu) = 0$, whence $M_\lambda^* \leq M(T^*)$ for any λ . This means

$$M_{\lambda_n}(T_n) \leq M_{\lambda_n}^* + \varepsilon_n \leq M(T^*) + \varepsilon_n.$$

Thus, the sequence $\{M_{\lambda_n}(T_n)\}_{n=1}^\infty$ is bounded, and so there exists a further subsequence $\{n_{k_m}\}$ such that

$$\lim_{m \rightarrow \infty} M_{\tilde{\lambda}_m}(\tilde{T}_m) = \kappa := \sup_k M_{\lambda_{n_k}}(T_{n_k}) < \infty,$$

where $\tilde{\lambda}_m = \lambda_{n_{k_m}}$ and $\tilde{T}_m = T_{n_{k_m}}$. Now choose m_0 and m_1 such that

$$\kappa < M_{\tilde{\lambda}_{m_0}}(\tilde{T}_{m_0}) + \delta, \quad \tilde{\lambda}_{m_1} > 4\tilde{\lambda}_{m_0}, \quad 2\delta < \tilde{\lambda}_{m_1} \tilde{\gamma}_{m_1}^2 < 4\delta,$$

where $\tilde{\gamma}_m = \gamma(\mu \circ T_{n_{k_m}}, \nu)$. With these choices it follows that

$$\begin{aligned} \kappa &< M_{\tilde{\lambda}_{m_0}}(\tilde{T}_{m_0}) + \delta \leq M_{\tilde{\lambda}_{m_0}}^* + \delta \leq M_{\tilde{\lambda}_{m_0}}(\tilde{T}_{m_1}) + \delta \\ &= M(\tilde{T}_{m_1}) + \frac{\tilde{\lambda}_{m_0}}{\tilde{\lambda}_{m_1}} \tilde{\lambda}_{m_1} \tilde{\gamma}_{m_1}^2 + \delta \\ &< M(\tilde{T}_{m_1}) + \frac{1}{4} \tilde{\lambda}_{m_1} \tilde{\gamma}_{m_1}^2 + \delta < M(\tilde{T}_{m_1}) + 2\delta \\ &< M(\tilde{T}_{m_1}) + \tilde{\lambda}_{m_1} \tilde{\gamma}_{m_1}^2 = M_{\tilde{\lambda}_{m_1}}(\tilde{T}_{m_1}) \leq \kappa, \end{aligned}$$

which is impossible.

Step (ii). Next we will show (5). Let $\varepsilon' > 0$. For

$R > 0$ we have

$$\begin{aligned} \mu \circ T_n(|x| > R) &\leq \frac{1}{R^2} \int_{\mathbb{R}^d} |T_n(x)|^2 \mu(dx) \\ &\leq \frac{2}{R^2} \int_{\mathbb{R}^d} |x - T_n(x)|^2 \mu(dx) \\ &\quad + \frac{2}{R^2} \int_{\mathbb{R}^d} |x|^2 \mu(dx) \\ &\leq \frac{2}{R^2} (M(T^*) + \varepsilon_n) + \frac{C}{R^2} \end{aligned}$$

for some constant $C > 0$. Thus we can take a sufficiently large R' such that $\mu_n \circ T_n(|x| > R') \leq \varepsilon'$. This means that $\{\mu \circ T_n^{-1}\}_{n=1}^\infty$ is tight, whence there exists a subsequence $\{n_k\}$ such that $\mu \circ T_{n_k}^{-1}$ converges weakly to some μ^* . Since we have assumed that γ metrizes the weak topology, we get

$$\lim_{k \rightarrow \infty} \gamma(\mu \circ T_{n_k}^{-1}, \mu^*) = 0.$$

This together with the step (i) yields

$$\gamma(\mu^*, \nu) \leq \gamma(\mu^*, \mu \circ T_{n_k}^{-1}) + \gamma(\mu \circ T_{n_k}^{-1}, \nu) \rightarrow 0, \quad k \rightarrow \infty,$$

whence $\mu^* = \nu$. Hence we have shown that each subsequence $\{\mu \circ T_{n_j}^{-1}\}$ contain a further subsequence $\{\mu \circ T_{n_{jm}}^{-1}\}$ that converges weakly to ν . Then, by Theorem 2.6 in Billingsley (Billingsley, 2013), we deduce that $\{\mu \circ T_n^{-1}\}$ converges weakly to ν . Denote by $W_2(\mu_1, \mu_2)$ the 2-Wasserstein distance between μ_1 and μ_2 . Then we have $W_2^2(\mu, \nu) = M(T^*)$ and the duality formula

$$W_2^2(\mu, \nu) = \sup \left\{ \int \varphi d\mu + \int \psi d\nu \right\},$$

where the supremum is taken over all bounded continuous functions φ and ψ such that $\varphi(x) + \psi(y) \leq |x - y|^2$, $x, y \in \mathbb{R}^d$. See, e.g., Proposition 5.3 in Carmona and Delarue (Carmona and Delarue, 2018). Let $\varepsilon'' > 0$ be arbitrary. Then there exist φ' and ψ' such that

$$\begin{aligned} W_2^2(\mu, \nu) &\leq \int_{\mathbb{R}^d} \varphi'(x) d\mu(x) + \int_{\mathbb{R}^d} \psi'(y) d\nu(y) + \varepsilon''. \end{aligned}$$

Further, since ψ' is bounded and continuous, there exists $n_0 \in \mathbb{N}$ such that $\varphi'(x) + \psi'(y) \leq |x - y|^2$ and

$$\int_{\mathbb{R}^d} \psi'(T^*(x)) d\mu(x) \leq \int_{\mathbb{R}^d} \psi'(T_n(x)) d\mu(x) + \varepsilon''$$

for $n \geq n_0$. With these choices it follows that for $n \geq$

n_0 ,

$$\begin{aligned} M(T^*) &= W_2^2(\mu, \nu) \\ &\leq \int_{\mathbb{R}^d} \varphi'(x) d\mu(x) + \int_{\mathbb{R}^d} \psi'(T^*(x)) d\mu(x) + \varepsilon'' \\ &\leq \int_{\mathbb{R}^d} (\varphi'(x) + \psi'(T_n(x))) d\mu(x) + 2\varepsilon'' \\ &\leq \int_{\mathbb{R}^d} |x - T_n(x)|^2 d\mu(x) + 2\varepsilon'' \\ &\leq M_{\lambda_n}(T_n) + 2\varepsilon''. \end{aligned}$$

Then letting $n \rightarrow \infty$ we get

$$M(T^*) \leq \liminf_{n \rightarrow \infty} M_{\lambda_n}(T_n) + 2\varepsilon''.$$

Since ε'' is arbitrary, we deduce $M(T^*) \leq \liminf_{n \rightarrow \infty} M_{\lambda_n}(T_n)$. On the other hand, (4) immediately leads to $\limsup_{n \rightarrow \infty} M_{\lambda_n}(T_n) \leq M(T^*)$. Therefore, $\lim_{n \rightarrow \infty} M_{\lambda_n}(T_n) = M(T^*)$, as wanted. \square

3 NUMERICAL EXPERIMENTS

Here we test our two algorithms through several numerical experiments.

3.1 Interpolation of Synthetic Datasets

All of numerical examples below are implemented in PyTorch on a Core(TM) i7-13700H with 32GB memory in this subsection. In these experiments, we describe three experiments on synthetic datasets. Date size is set to be 500 for each experiment. The Gaussian kernel $K(x, y) = e^{-|x-y|^2}$, cost function $c(x, y) = |x - y|^2$ and the Adam optimizer with learning rate 0.0001 is used. Here, the function $T(x)$ is described by a multi-layer perceptron with 1 hidden layer. These result that obtained after about 3000 epochs. Penalty parameter λ defined by $1/\lambda = 0.000001$.

3.1.1 From Moon to Circle

In this experiment, the initial distribution is the well-known synthetic dataset generated by two ‘‘moons’’ (Figure 1), and the target distribution is the one generated by two ‘‘circles’’ (Figure 2).




Figure 1: Initial distribution.




Figure 2: Target distribution (blue) and generated samples (orange).




Figure 3: Loss curve.

We can see from Figure 2 that the proposed method faithfully generates the target distribution. Figure 3 shows the change in loss, and the loss converges in the first 500 epochs.

3.1.2 From Normal Distribution to Moon

In this experiment, the initial distribution is a two-dimensional uncorrelated normal distribution with mean 0 and variance 1 (Figure 4), and the final distribution is the synthetic dataset generated by the two moons as in Section 3.1.1 (Figure 5).




Figure 4: Initial distribution.




Figure 5: Target distribution (blue) and generated samples (orange).




Figure 6: Loss curve.

We can see from Figure 5 that the proposed method again generates the target distribution correctly with a small variance. Figure 6 shows the change in loss. We can see that the loss converges in the first 500 epochs.

3.1.3 From Normal Distribution to Normal Distribution

In this experiment, the initial distribution is a two-dimensional uncorrelated normal distribution with mean 0 and variance 1, and the target distribution is a two-dimensional uncorrelated normal distribution with mean 5 and variance 1.




Figure 7: Initial distribution.




Figure 8: Target distribution (blue) and generated samples (orange).




Figure 9: Loss curve.

For normal distributions, we see from Figures 7-9 that stable generation is achieved as in Sections 3.1.1 and 3.1.2.

3.2 Comparison with POT

We compared the performance with POT, the existing Python library as mentioned in Section 1. Here, the initial distribution is a two-dimensional uncorrelated normal distribution with mean 0 and variance 1, and the target distribution is a two-dimensional uncorrelated normal distribution with mean 5 and variance 1. The Gaussian kernel $K(x, y) = e^{-|x-y|^2}$, cost function $c(x, y) = |x - y|^2$ and the Adam optimizer with learning rate 0.0001 is used. Here, the function $T(x)$ is described by a multi-layer perceptron with 2 hidden layer. Penalty parameter λ defined by $1/\lambda = 0.000001$. In this experiment, we compare the number of points that can be calculated, the number of times the calculation can be run, and the accuracy achieved when using an AMD EPYC 9654 with 768GB of memory. The experimental result is described in Table 1, where SD stands for standard deviation. Here, we set the number of batch size to

Table 1: The proposed method with CPU.

Data size	epochs	expectation, SD
600000	3	5.03, 0.99

be 10000 and so the number of iterations is 60. Table 1 shows that the proposed method is able to compute accurately even for large data sizes of 600000.

Next, we perform a similar experiment on NVIDIA H100. The experimental result is described in Table 2, where SD again stands for standard deviation. Here we set the number of batch size to be

Table 2: The proposed method with GPU.

Data size	epochs	expectation, SD
600000	10	4.97, 1.01

10000 and so the number of iterations is 60. We can see from Table 2 that our method is able to compute accurately even for large data sizes of 600000 as in the CPU case.

Then, we use the solver `ot.sinkhorn()` in POT to compare the performance of POT with that of our algorithm on a Core(TM) i7-13700H with 32GB memory. The computational complexity of this solver is known to be $O(n^2)$, where n is the input data size. Table 3 shows that a stable result is obtained when

Table 3: Python Optimal Transport.

Data size	expectation, SD
1000	5.00, 0.97

the data size is 1000. When the data size exceeded 2000, the solver is unstable due to the computational limitations mentioned above.

Numerical experiments in this subsection show that our proposed method is a promising option for solving large-scale Monge problems.

4 CONCLUSION

In this paper, we derive an approximate solution to the Monge problem by using the embeddings of probability measures and a deep learning algorithm. Through several numerical experiments, we confirmed that our method produces accurate approximate solutions and is efficiently computable on a GPU. In future work, we aim to extend our research to handle larger datasets, explore the Monge problem with alternative cost functions, and investigate numerical solutions for multi-marginal transport problems.

REFERENCES

- Bauschke, H. H. and Combettes, P. L. (2011). *Convex analysis and monotone operator theory in Hilbert spaces*. Springer-Verlag.
- Billingsley, P. (2013). *Convergence of probability measures*. John Wiley & Sons.
- Breiner, Y. (1987). Décomposition polaire et réarrangement monotone des champs de vecteurs. *C. R. Acad. Sci. Paris Série I*, 305(19):805–808.
- Breiner, Y. (1991). Polar factorization and monotone rearrangement of vector-valued functions. *Comm. Pure Appl. Math.*, 44(4):375–417.
- Carmona, R. and Delarue, F. (2018). *Probabilistic theory of mean field games with applications I*. Springer.
- Charlier, B., Feydy, J., Glaunès, J. A., Collin, F. D., and Durif, G. (2021). Kernel operations on the GPU, with autodiff, without memory overflows. *Journal of Machine Learning Research*, 22(74):1–6.
- Cuturi, M., Meng-Papaxanthos, L., Tian, Y., Bunne, C., Davis, G., and Teboul, O. (2022). Optimal transport tools (ott): A jax toolbox for all things wasserstein. *arXiv:2201.12324*.
- Dantzig, G. B. (1949). Programming of interdependent activities: II mathematical model. *Econometrica*, 17(3/4):200–211.
- Dantzig, G. B. (1951). Application of the simplex method to a transportation problem. *Activity Analysis of Production and Allocation*, 13:359–373.
- Feydy, J., Séjourné, T., Vialard, F. X., Amari, S., Trounev, A., and Peyré, G. (2019). Interpolating between Optimal Transport and MMD using Sinkhorn Divergences. In *The 22nd International Conference on Artificial Intelligence and Statistics*, pages 2681–2690.
- Flamary, R., Courty, N., Gramfort, A., Alaya, M. Z., Boissunon, A., Chambon, S., Chapel, L., Corenflos, A., Fatras, K., Fournier, N., Gautheron, L., Gayraud, N.

- T. H., Janati, H., Rakotomamonjy, A., Redko, I., Rolet, A., Schutz, A., Seguy, V., Sutherland, D. J., Tavenard, R., Tong, A., and Vayer, T. (2021). POT: Python Optimal Transport. *Journal of Machine Learning Research*, 22(78):1–8.
- Gangbo, W. and McCann, R. J. (1996). The geometry of optimal transportation. *Acta Math.*, 177:113–161.
- Gretton, A., Borgwardt, K., Rasch, M., Schölkopf, B., and Smola, A. (2006). A kernel method for the two-sample-problem. *Advances in neural information processing systems*, 19.
- Kantorovich, L. V. (1942). On mass transportation. (*Russian*) Reprinted from *C. R. (Doklady) Acad.Sci. URSS (N. S.)*, 37:199–201.
- Kantorovich, L. V. (1948). On a problem of monge. (*Russian*) Reprinted from *C. R. (Doklady) Acad.Sci. URSS (N. S.)*, 3:225–226.
- Korte, B. and Vygen, J. (2012). *Combinatorial Optimization*. Springer.
- Mikami, T. (2004). Monge’s problem with a quadratic cost by the zero-noise limit of h-path process. *Probab. Theory Related Fields*, 129:245–260.
- Monge, G. (1781). *Mémoires sur la théorie des déblais et des remblais*. De l’Imprimerie Royale.
- Nakano, Y. (2024). A kernel-based method for Schrödinger bridges. arXiv:2310.14522.
- Papadakis, N., Peyré, G., and Oudet, E. (2014). Optimal transport with proximal splitting. *SIAM Journal on Imaging Sciences*, 7:212–238.
- Peyré, G. and Cuturi, M. (2019). *Computational Optimal Transport*. Now Publishers.
- Santambrogio, F. (2015). *Optimal Transport for Applied Mathematicians Calculus of Variations, PDEs, and Modeling*. Springer.
- Sriperumbudur, B. K., Gretton, A., Fukumizu, K., Schölkopf, B., and Lanckriet, G. R. G. (2010). Hilbert space embeddings and metrics on probability measures. *J. Mach. Learn. Res.*, 11:1517–1561.
- Tolstoi, A. N. (1930). Metody nakhozhdeniya naimen’ shego summovogo kilome-trazha pri planirovanii perevozok v prostranstve (russian; methods of finding the minimal total kilometrage in cargo transportation planning in space). *TransPress of the National Commissariat of Transportation*, pages 23–55.
- Villani, C. (2021). *Topics in optimal transportation*. American Mathematical Soc.
- Wendland, H. (2010). *Scattered data approximation*. Cambridge University Press, Cambridge.