Vector and Matrix Optimal Mass Transport: Theory, Algorithm, and Applications

Ernest K. Ryu, Yongxin Chen, Wuchen Li, and Stanley Osher

December 29, 2017

Abstract

In many applications such as color image processing, data has more than one piece of information associated with each spatial coordinate, and in such cases the classical optimal mass transport (OMT) must be generalized to handle vector-valued or matrix-valued densities. In this paper, we discuss the vector and matrix optimal mass transport and present three contributions. We first present a rigorous mathematical formulation for these setups and provide analytical results including existence of solutions and strong duality. Next, we present a simple, scalable, and parallelizable methods to solve the vector and matrix-OMT problems. Finally, we implement the proposed methods on a CUDA GPU and present experiments and applications.

1   Introduction

Optimal mass transport (OMT) is a subject with a long history. Started by Monge [38] and developed by many great mathematicians [32, 8, 26, 37, 31, 5, 43], the subject now has incredibly rich theory and applications. It has found numerous applications in different areas such as partial differential equations, probability theory, physics, economics, image processing, and control [24, 29, 52, 39, 15, 16, 11]. See also [47, 53, 2] and references therein.

However, in many applications such as color image processing, there is more than one piece of information associated with each spatial coordinate, and such data can be interpreted as vector-valued or matrix-valued densities. As the classical optimal mass transport works with scalar probability densities, such applications require a new notion of mass transport. For this purpose, Chen et al. [18, 19, 14, 12] recently developed a framework for vector-valued and matrix-valued optimal mass transport. See also [41, 40, 42, 25, 45, 55] for other different frameworks. For vector-valued OMT, potential applications include color image processing, multi-modality medical imaging, and image processing involving textures. For matrix-valued OMT, we have diffusion tensor imaging, multivariate spectral analysis, and stress tensor analysis.

Several mathematical aspects of vector and matrix-valued OMT were not addressed in the previous work [18, 19, 14]. As the first contribution of this paper, we present duality and existence results of the continuous vector and matrix-valued OMT problems along with rigorous problem formulations.

Although the classical theory of OMT is very rich, only recently has there been much attention to numerical methods to compute the OMT. Several recent work proposed algorithms to solve the L^2 OMT [3, 22, 7, 28, 6, 13, 17, 27] and the L^1 OMT [34, 50]. As the second contribution of this paper, we present first-order primal-dual methods to solve the vector and matrix-valued OMT problems. The methods simultaneously solve for both the primal and dual solutions (hence a primal-dual method) and are scalable as they are first-order methods. We also discuss the convergence of the methods.

As the third contribution of this paper, we implement the proposed method on a CUDA GPU and present several applications. The proposed algorithms' simple structure allows us to effectively utilize the computing capability of the CUDA architecture, and we demonstrate this through our experiments. We furthermore release the code for scientific reproducibility.

The rest of the paper is structured as follows. In Section 2 we give a quick review of the classic OMT theory, which allows us to present the later sections in an analogous manner and thereby outline the similarities and differences. In Section 3 and Section 4 we present the vector and matrix-valued OMT problems

and state a few theoretical results. In Section 5, we present and prove the analytical results. In Section 7, we present the algorithm. In Section 8, we present the experiments and applications.

2 Optimal mass transport

Let $\Omega \subset \mathbb{R}^d$ be a closed, convex, compact domain. Let λ^0 and λ^1 be nonnegative densities supported on Ω with unit mass, i.e., $\int_\Omega \lambda^0(\mathbf{x}) \, d\mathbf{x} = \int_\Omega \lambda^1(\mathbf{x}) \, d\mathbf{x} = 1$. Let $\| \cdot \|$ denote any norm on $\mathbb{R}^d$.

In 1781, Monge posed the optimal mass transport (OMT) problem, which solves

$$\underset{T}{\text{minimize}} \quad \int_\Omega \|\mathbf{x} - T(\mathbf{x})\| \lambda^0(\mathbf{x}) \, d\mathbf{x}. \tag{1}$$

The optimization variable $T : \Omega \to \Omega$ is smooth, one-to-one, and transfers $\lambda^0(\mathbf{x})$ to $\lambda^1(\mathbf{x})$. The optimization problem (1) is nonlinear and nonconvex. In 1940, Kantorovich relaxed (1) into a linear (convex) optimization problem:

$$S(\lambda^0, \lambda^1) = \begin{pmatrix} \underset{\pi}{\text{minimize}} & \int_{\Omega \times \Omega} \|\mathbf{x} - \mathbf{y}\| \pi(\mathbf{x}, \mathbf{y}) \, d\mathbf{x} d\mathbf{y} \\ \text{subject to} & \pi(\mathbf{x}, \mathbf{y}) \geq 0 \\ & \int_\Omega \pi(\mathbf{x}, \mathbf{y}) \, d\mathbf{y} = \lambda^0(\mathbf{x}) \\ & \int_\Omega \pi(\mathbf{x}, \mathbf{y}) \, d\mathbf{x} = \lambda^1(\mathbf{y}). \end{pmatrix} \tag{2}$$

The optimization variable π is a joint nonnegative measure on $\Omega \times \Omega$ having $\lambda^0(\mathbf{x})$ and $\lambda^1(\mathbf{y})$ as marginals. To clarify, $S(\lambda^0, \lambda^1)$ denotes the optimal value of (2).

2.1 Scalar optimal mass transport

The theory of optimal transport [26, 53, 54] remarkably points out that (2) is equivalent to the following flux minimization problem:

$$S(\lambda^0, \lambda^1) = \begin{pmatrix} \underset{\mathbf{u}}{\text{minimize}} & \int_\Omega \|\mathbf{u}(\mathbf{x})\| \, d\mathbf{x} \\ \text{subject to} & \text{div}_\mathbf{x}(\mathbf{u})(\mathbf{x}) = \lambda^0(\mathbf{x}) - \lambda^1(\mathbf{x}) \\ & \mathbf{u}(\mathbf{x})^T \mathbf{n}(\mathbf{x}) = 0, \text{ for all } \begin{cases} \mathbf{x} \in \partial\Omega, \\ \mathbf{n}(\mathbf{x}) \text{ normal to } \partial\Omega, \end{cases} \end{pmatrix} \tag{3}$$

where $\mathbf{u} = (u_1, \dots, u_d) : \Omega \to \mathbb{R}^d$ is the optimization variable and $\text{div}_\mathbf{x}$ denote the (spatial) divergence operator. Although (3) and (2) are mathematically equivalent, (3) is much more computationally effective as its optimization variable $\mathbf{u}$ is much smaller when discretized.

It is worth mentioning that OMT in formulation (3) is very close to the problems in compressed sensing. Its objective function is homogeneous degree one and the constraint is linear. It can be observed that for characterizing the OMT, the gradient operator in (3) and divergence operator in (3) play the key roles. Later on, we extend the definition of L_1 OMT problem by extending these differential operators into a general meaning.

The optimization problem (3) has the following dual problem:

$$S(\lambda^0, \lambda^1) = \begin{pmatrix} \underset{\phi}{\text{maximize}} & \int_\Omega \phi(\mathbf{x})(\lambda^1(\mathbf{x}) - \lambda^0(\mathbf{x})) \, d\mathbf{x} \\ \text{subject to} & \|\nabla_\mathbf{x}\phi(\mathbf{x})\|_* \leq 1 \quad \text{for all } \mathbf{x} \in \Omega, \end{pmatrix} \tag{4}$$

where the optimization and $\phi : \Omega \to \mathbb{R}$ is a function. We write $\| \cdot \|_*$ for the dual norm of $\| \cdot \|$.

It is well-known that, strong duality holds between (3) and (4) in the sense that the minimized and maximized objective values are equal [53]. Therefore, we take either (3) or (4) as the definition of S.

Rigorous definitions of the optimization problems (3) or (4) are somewhat technical. We skip this discussion, as scalar optimal mass transport is standard. In Section 5, we rigorously discuss the vector-OMT problems, so any rigorous discussion of the scalar-OMT problems can be inferred as a special case.

Figure 1: Example graph with $k = 4$ nodes and $\ell = 5$ edges. To make c_j the cost of traversing edge j, the edge weight is defined to be $1/c_j^2$ for $j = 1, \ldots, \ell$.

2.2 Theoretical properties

The algorithm we present in Section 7 is a primal-dual algorithm and, as such, finds solutions to both Problems (3) and (4). This is well-defined as both the primal and dual problems have solutions [53].

Write $\mathbb{R}_+$ for the set of nonnegative real numbers. Write $\mathcal{P}(\Omega, \mathbb{R})$ for the space of nonnegative densities supported on Ω with unit mass. We can use $S(\lambda^0, \lambda^1)$ as a distance measure between $\lambda^0, \lambda^1 \in \mathcal{P}(\Omega, \mathbb{R})$. The value $S : \mathcal{P}(\Omega, \mathbb{R}) \times \mathcal{P}(\Omega, \mathbb{R}) \to \mathbb{R}_+$ defines a metric on $\mathcal{P}(\Omega, \mathbb{R})$ [53].

3 Vector optimal mass transport

Next we discuss the vector-valued optimal transport, proposed recently in [19]. The basic idea is to combine scalar optimal mass transport with network flow problems [1].

3.1 Gradient and divergence on graphs

Consider a connected, positively weighted, undirected graph $\mathcal{G}$ with k nodes and ℓ edges. To define an incidence matrix for $\mathcal{G}$, we say an edge $\{i, j\}$ points from i to j, i.e., $i \to j$, if $i < j$. This choice is arbitrary and does not affect the final result. With this edge orientation, the incidence matrix $D \in \mathbb{R}^{k \times \ell}$ is

$$D_{ie} = \begin{cases} +1 & \text{if edge } e = \{i, j\} \text{ for some node } j > i \\ -1 & \text{if edge } e = \{j, i\} \text{ for some node } j < i \\ 0 & \text{otherwise.} \end{cases}$$

For example, the incidence matrix of the graph of Figure 1 is

$$D = \begin{bmatrix} 1 & 1 & 0 & 1 & 0 \\ -1 & 0 & 1 & 0 & 0 \\ 0 & -1 & -1 & 0 & 1 \\ 0 & 0 & 0 & -1 & -1 \end{bmatrix}.$$

Write $\Delta_{\mathcal{G}} = -D \operatorname{diag}\{1/c_1^2, \cdots, 1/c_\ell^2\} D^T$ for the (negative) graph Laplacian, where $1/c_1^2, \cdots, 1/c_\ell^2$ are the edge weights. The edge weights are defined so that and c_j represents the cost of traversing edge j for $j = 1, \ldots, \ell$.

We define the gradient operator on $\mathcal{G}$ as $\nabla_{\mathcal{G}} x = \operatorname{diag}\{1/c_1, \cdots, 1/c_\ell\} D^T x$ and the divergence operator as $\operatorname{div}_{\mathcal{G}} y = -D \operatorname{diag}\{1/c_1, \cdots, 1/c_\ell\} y$. So the Laplacian can be rewritten as $\Delta_{\mathcal{G}} = \operatorname{div}_{\mathcal{G}} \nabla_{\mathcal{G}}$. Note that $\operatorname{div}_{\mathcal{G}} = -\nabla_{\mathcal{G}}^*$, where $\nabla_{\mathcal{G}}^*$ is the adjoint of $\nabla_{\mathcal{G}}$. This is in analogy with the usual spatial gradient and divergence operators [36, 20, 21, 19].

3.2 Vector optimal mass transport

We say $\vec{\lambda} : \Omega \to \mathbb{R}_+^k$ is a nonnegative vector-valued density with unit mass if

$$\vec{\lambda}(\mathbf{x}) = \begin{bmatrix} \lambda_1(\mathbf{x}) \\ \vdots \\ \lambda_k(\mathbf{x}) \end{bmatrix}, \qquad \int_\Omega \sum_{i=1}^k \lambda_i(\mathbf{x}) \, d\mathbf{x} = 1.$$

3

Assume $\vec{\lambda}^0$ and $\vec{\lambda}^1$ are nonnegative vector-valued densities supported on Ω with unit mass.

We define the optimal mass transport between vector-valued densities as

$$V(\vec{\lambda}^0, \vec{\lambda}^1) = \left(\begin{array}{ll} \underset{\vec{\mathbf{u}}, \vec{w}}{\text{minimize}} & \int_\Omega \|\vec{\mathbf{u}}(\mathbf{x})\|_u + \alpha\|\vec{w}(\mathbf{x})\|_w \, d\mathbf{x} \\ \text{subject to} & \text{div}_\mathbf{x}(\vec{\mathbf{u}})(\mathbf{x}) + \text{div}_\mathcal{G}(\vec{w}(\mathbf{x})) = \vec{\lambda}^0(\mathbf{x}) - \vec{\lambda}^1(\mathbf{x}) \\ & \vec{\mathbf{u}} \text{ satisfies zero-flux b.c.} \end{array} \right) \tag{5}$$

where $\vec{\mathbf{u}} : \Omega \to \mathbb{R}^{k \times d}$ and $\vec{w} : \Omega \to \mathbb{R}^\ell$ are the optimization variables, $\alpha > 0$ is a parameter, and $\|\cdot\|_u$ is a norm on $\mathbb{R}^{k \times d}$ and $\|\cdot\|_w$ is a norm on $\mathbb{R}^\ell$. The parameter α represents the relative importance of the two flux terms $\vec{\mathbf{u}}$ and $\vec{w}$. We write

$$\vec{\mathbf{u}} = \begin{bmatrix} \mathbf{u}_1^T \\ \vdots \\ \mathbf{u}_k^T \end{bmatrix} \qquad \vec{w} = \begin{bmatrix} w_1 \\ \vdots \\ w_\ell \end{bmatrix} \qquad \text{div}_\mathbf{x}(\vec{\mathbf{u}}) = \begin{bmatrix} \text{div}_\mathbf{x}(\mathbf{u}_1) \\ \vdots \\ \text{div}_\mathbf{x}(\mathbf{u}_k) \end{bmatrix}.$$

We call $\text{div}_\mathbf{x}$ the spatial divergence operator. The zero-flux boundary condition is

$$\mathbf{u}_i(\mathbf{x})^T \mathbf{n}(\mathbf{x}) = 0, \text{ for all } \begin{cases} \mathbf{x} \in \partial\Omega, \\ \mathbf{n}(\mathbf{x}) \text{ normal to } \partial\Omega. \end{cases}$$

for $i = 1, \ldots, k$. Note that $\vec{w}$ has no boundary conditions.

The optimization problem (5) has the following dual problem:

$$V(\vec{\lambda}^0, \vec{\lambda}^1) = \left(\begin{array}{ll} \underset{\vec{\phi}}{\text{maximize}} & \int_\Omega \langle \vec{\phi}(\mathbf{x}), \vec{\lambda}^1(\mathbf{x}) - \vec{\lambda}^0(\mathbf{x}) \rangle \, d\mathbf{x} \\ \text{subject to} & \|\nabla_\mathbf{x}\vec{\phi}(\mathbf{x})\|_{u*} \leq 1 \\ & \|\nabla_\mathcal{G}\vec{\phi}(\mathbf{x})\|_{w*} \leq \alpha \quad \text{for all } \mathbf{x} \in \Omega, \end{array} \right) \tag{6}$$

where the optimization variable $\vec{\phi} : \Omega \to \mathbb{R}^k$ is a function. We write $\|\cdot\|_{u*}$ and $\|\cdot\|_{w*}$ for the dual norms of $\|\cdot\|_u$ and $\|\cdot\|_w$, respectively.

As stated in Theorem 1, strong duality holds between (5) and (6) in the sense that the minimized and maximized objective values are equal. Therefore, we take either (5) or (6) as the definition of V. In Section 5, we rigorously define the primal and dual problems and prove Theorem 1.

3.3 Theoretical properties

The algorithm we present in Section 7 is a primal-dual algorithm and, as such, finds solutions to both Problems (5) and (6). This is well-defined as both the primal and dual problems have solutions.

Theorem 1. *The (infinite dimensional) primal and dual optimization problems* (5) *and* (6) *have solutions, and their optimal values are the same, i.e., strong duality holds.*

Write $\mathcal{P}(\Omega, \mathbb{R}^k)$ for the space of nonnegative vector-valued densities supported on Ω with unit mass. We can use $V(\vec{\lambda}^0, \vec{\lambda}^1)$ as a distance measure between $\vec{\lambda}^0, \vec{\lambda}^1 \in \mathcal{P}(\Omega, \mathbb{R}^k)$.

Theorem 2. $V : \mathcal{P}(\Omega, \mathbb{R}^k) \times \mathcal{P}(\Omega, \mathbb{R}^k) \to \mathbb{R}_+$ *defines a metric on* $\mathcal{P}(\Omega, \mathbb{R}^k)$.

4 Quantum gradient operator and matrix optimal mass transport

We closely follow the treatment in [18]. In particular, we define a notion of gradient on the space of Hermitian matrices and its dual, i.e., the (negative) divergence.

Some applications of matrix-OMT, such as diffusion tensor imaging, have real-valued data while some applications, such as multivariate spectral analysis, have complex-valued data [51]. To accommodate the wide range of applications, we develop the matrix-OMT with complex-valued matrices.

Write $\mathcal{C}$, $\mathcal{H}$, and $\mathcal{S}$ for the set of $k \times k$ complex, Hermitian, and skew-Hermitian matrices respectively. We write $\mathcal{H}_+$ for the set of $k \times k$ positive semidefinite Hermitian matrices, i.e., $M \in \mathcal{H}_+$ if $v^* M v \geq 0$ for all $v \in \mathbb{C}^k$. Write tr for the trace, i.e. for any $M \in \mathcal{H}$, we have $\mathrm{tr}(M) = \sum_{i=1}^k M_{ii}$.

Write $\mathcal{C}^N$ for the block-column concatenation of N matrices in $\mathcal{C}$, i.e., $\mathbf{Z} \in \mathcal{C}^N$ if $\mathbf{Z} = [Z_1^* \cdots Z_N^*]^*$ and $Z_1, \ldots, Z_N \in \mathcal{C}$. Define $\mathcal{H}^N$ and $\mathcal{S}^N$ likewise. For $X, Y \in \mathcal{C}$, we use the Hilbert-Schmidt inner product

$$\langle X, Y \rangle = \mathrm{Re}\,\mathrm{tr}(XY^*) = \sum_{i=1}^k \sum_{j=1}^k (\mathrm{Re}X_{ij}\mathrm{Re}Y_{ij} + \mathrm{Im}X_{ij}\mathrm{Im}Y_{ij}).$$

(This is the standard inner product when we view $\mathcal{C}$ as the real vector space $\mathbb{R}^{2n^2}$.) For $X \in \mathcal{C}$, we use the norm $\|X\|_2 = (\langle X, X \rangle)^{1/2}$. For $\mathbf{X}, \mathbf{Y} \in \mathcal{C}^N$, we use the inner product $\langle \mathbf{X}, \mathbf{Y} \rangle = \sum_{s=1}^N \langle X_s, Y_s \rangle$.

4.1 Quantum gradient and divergence operators

We define the gradient operator, given a $\mathbf{L} = [L_1, \cdots, L_\ell]^* \in \mathcal{H}^\ell$, as

$$\nabla_{\mathbf{L}} : \mathcal{H} \to \mathcal{S}^\ell, \quad X \mapsto \begin{bmatrix} L_1 X - X L_1 \\ \vdots \\ L_\ell X - X L_\ell \end{bmatrix}.$$

Define the divergence operator as

$$\mathrm{div}_{\mathbf{L}} : \mathcal{S}^\ell \to \mathcal{H}, \quad Z = \begin{bmatrix} Z_1 \\ \vdots \\ Z_\ell \end{bmatrix} \mapsto \sum_{s=1}^\ell -L_s Z_s + Z_s L_s.$$

Note that $\mathrm{div}_{\mathbf{L}} = -\nabla_{\mathbf{L}}^*$, where $\nabla_{\mathbf{L}}^*$ is the adjoint of $\nabla_{\mathbf{L}}$. This is in analogy with the usual spatial gradient and divergence operators. Write $\Delta_{\mathbf{L}} = \mathrm{div}_{\mathbf{L}} \nabla_{\mathbf{L}}$.

This notion of gradient and divergence operators is motivated by the Lindblad equation in Quantum mechanics [18]. The choice of $\mathbf{L}$ affects $\nabla_{\mathbf{L}}$. There is no standard way of choosing $\mathbf{L}$. A standing assumption throughout, is that the null space of $\nabla_{\mathbf{L}}$, denoted by $\ker(\nabla_{\mathbf{L}})$, contains only scalar multiples of the identity matrix I.

4.2 Matrix optimal mass transport

We say $\Lambda : \Omega \to \mathcal{H}_+$ is a nonnegative matrix-valued density with unit mass if

$$\int_\Omega \mathrm{tr}(\Lambda(\mathbf{x}))\, d\mathbf{x} = 1.$$

Assume Λ^0 and Λ^1 are nonnegative matrix-valued densities supported on Ω with unit mass.

We define the optimal mass transport between matrix-valued densities as

$$M(\Lambda^0, \Lambda^1) = \left(\begin{array}{ll} \underset{\mathbf{U},\mathbf{W}}{\mathrm{minimize}} & \int_\Omega \|\mathbf{U}(\mathbf{x})\|_u + \alpha \|\mathbf{W}(\mathbf{x})\|_w\, d\mathbf{x} \\ \mathrm{subject\ to} & \mathrm{div}_{\mathbf{x}}(\mathbf{U})(\mathbf{x}) + \mathrm{div}_{\mathbf{L}}(\mathbf{W}(\mathbf{x})) = \Lambda^0(\mathbf{x}) - \Lambda^1(\mathbf{x}) \\ & \mathbf{U} \text{ satisfies zero-flux b.c.} \end{array} \right) \tag{7}$$

where $\mathbf{U} : \Omega \to \mathcal{H}^d$ and $\mathbf{W} : \Omega \to \mathcal{S}^\ell$ are the optimization variables, $\alpha > 0$ is a parameter, and $\|\cdot\|_u$ is a norm on $\mathcal{H}^d$ and $\|\cdot\|_w$ is a norm on $\mathcal{S}^\ell$. The parameter α represents the relative importance of the two flux terms $\mathbf{U}$ and $\mathbf{W}$. We write

$$\mathbf{U} = \begin{bmatrix} U_1 \\ \vdots \\ U_d \end{bmatrix} \quad \mathbf{W} = \begin{bmatrix} W_1 \\ \vdots \\ W_\ell \end{bmatrix} \quad \mathbf{u}_{ij} = \begin{bmatrix} (U_1)_{ij} \\ \vdots \\ (U_d)_{ij} \end{bmatrix}.$$

5

We define the spatial divergence as

$$\text{div}_{\mathbf{x}}(\mathbf{U}) = \begin{bmatrix} \text{div}_{\mathbf{x}}(\mathbf{u}_{11}) & \text{div}_{\mathbf{x}}(\mathbf{u}_{12}) & \cdots & \text{div}_{\mathbf{x}}(\mathbf{u}_{1k}) \\ \text{div}_{\mathbf{x}}(\overline{\mathbf{u}}_{12}) & \ddots & & \vdots \\ \vdots & & \ddots & \vdots \\ \text{div}_{\mathbf{x}}(\overline{\mathbf{u}}_{1k}) & \text{div}_{\mathbf{x}}(\overline{\mathbf{u}}_{2k}) & \cdots & \text{div}_{\mathbf{x}}(\mathbf{u}_{k,k}) \end{bmatrix}.$$

The zero-flux boundary condition is

$$\mathbf{u}_{ij}(\mathbf{x})^T \mathbf{n}(\mathbf{x}) = 0, \text{ for all } \begin{cases} \mathbf{x} \in \partial\Omega, \\ \mathbf{n}(\mathbf{x}) \text{ normal to } \partial\Omega. \end{cases}$$

for $i, j = 1, \ldots, k$. Note that $\mathbf{W}$ has no boundary conditions [14].

The optimization problem (7) has the following dual problem:

$$M(\Lambda^0, \Lambda^1) = \begin{pmatrix} \underset{\Phi}{\text{maximize}} & \int_\Omega \langle \Phi(\mathbf{x}), \Lambda^1(\mathbf{x}) - \Lambda^0(\mathbf{x}) \rangle \, d\mathbf{x} \\ \text{subject to} & \|\nabla_{\mathbf{x}}\Phi(\mathbf{x})\|_{u*} \leq 1 \\ & \|\nabla_{\mathbf{L}}\Phi(\mathbf{x})\|_{w*} \leq \alpha \quad \text{for all } \mathbf{x} \in \Omega, \end{pmatrix} \tag{8}$$

where the optimization variable $\Phi : \Omega \to \mathcal{H}$ is a function. We write $\|\cdot\|_{u*}$ and $\|\cdot\|_{w*}$ for the dual norms of $\|\cdot\|_u$ and $\|\cdot\|_w$, respectively.

As stated in Theorem 3, strong duality holds between (7) and (8) in the sense that the minimized and maximized objective values are equal. Therefore we take either (7) or (8) as the definition of M.

To avoid repeating the same argument with different notation, we simply point out that the rigorous definitions of the matrix optimal mass transport problems are analogous to those of the vector setup. So the precise definitions of (7) or (8) can be inferred from the discussion of Section 5.

4.3 Theoretical properties

The algorithm we present in Section 7 is a primal-dual algorithm and, as such, finds solutions to both Problems (7) and (8). This is well-defined as both the primal and dual problems have solutions.

Theorem 3. *The (infinite dimensional) primal and dual optimization problems* (7) *and* (8) *have solutions, and their optimal values are the same, i.e., strong duality holds.*

Write $\mathcal{P}(\Omega, \mathcal{H}_+)$ for the space of nonnegative matrix-valued densities supported on Ω with unit mass. We can use $M(\Lambda^0, \Lambda^1)$ as a distance measure between $\Lambda^0, \Lambda^1 \in \mathcal{P}(\Omega, , \mathcal{H}_+)$.

Theorem 4. $M : \mathcal{P}(\Omega, \mathcal{H}_+) \times \mathcal{P}(\Omega, \mathcal{H}_+) \to \mathbb{R}_+$ *defines a metric on* $\mathcal{P}(\Omega, \mathcal{H}_+)$.

5 Duality proof

In this section, we establish the theoretical results. For notational simplicity, we only prove the results for the vector-OMT primal and dual problems (5) and (6). Analogous results for the matrix-OMT primal and dual problems (7) and (8) follow from the same logic.

Although the classical scalar-OMT literature is very rich, standard techniques for proving scalar-OMT duality do not simply apply to our setup. For example, Villani's proof of strong duality, presented as Theorem 1.3 of [53], relies on and works with the linear optimization formulation (2). However, our vector and matrix-OMT formulations directly generalize the flux formulation (3) and do not have formulations analogous to (2). We need a direct approach to analyze duality between the flux formulation and the dual (with one function variable), and we provide this in this section.

We further assume $\Omega \subset \mathbb{R}^d$ has a piecewise smooth boundary. Write Ω° and $\partial\Omega$ for the interior and boundary of Ω. For simplicity, assume Ω as full affine dimensions, i.e., $\overline{\Omega^\circ} = \Omega$.

The rigorous form the dual problem (6) is

$$
\begin{aligned}
&\underset{\vec{\phi}\in W^{1,\infty}(\Omega,\mathbb{R}^k)}{\text{maximize}} && \int_\Omega \langle \vec{\phi}(\mathbf{x}), \vec{\lambda}^1(\mathbf{x}) - \vec{\lambda}^0(\mathbf{x}) \rangle \, d\mathbf{x} \\
&\text{subject to} && \operatorname{ess\,sup}_{\mathbf{x}\in\Omega} \|\nabla_{\mathbf{x}}\vec{\phi}(\mathbf{x})\|_{u*} \leq 1 \\
& && \sup_{\mathbf{x}\in\Omega} \|\nabla_{\mathcal{G}}\vec{\phi}(\mathbf{x})\|_{w*} \leq \alpha,
\end{aligned}
\tag{9}
$$

where $W^{1,\infty}(\Omega,\mathbb{R}^k)$ is the standard Sobolev space of functions from Ω to $\mathbb{R}^k$ with bounded weak gradients. That (9) has a solution directly follows from the Arzelà-Ascoli Theorem.

To rigorously define the primal problem (5) requires more definitions, and we do so later as (11).

5.1 Fenchel-Rockafellar duality

Let $L : X \to Y$ be a continuous linear map between locally convex topological vector spaces X and Y, and let $f : X \to \mathbb{R} \cup \{\infty\}$ and $g : Y \to \mathbb{R} \cup \{\infty\}$ are lower-semicontinuous convex functions. Write

$$
d^\star = \sup_{x\in X} \{-f(x) - g(Lx)\} \qquad p^\star = \inf_{y^*\in Y^*} \{f^*(L^* y^*) + g^*(-y^*)\},
$$

where

$$
f^*(x^*) = \sup_{x\in X} \{\langle x^*, x \rangle - f(x)\} \qquad g^*(y^*) = \sup_{y\in Y} \{\langle y^*, y \rangle - g(y)\}.
$$

In this framework of Fenchel-Rockafellar duality, $d^\star \leq p^\star$, i.e., weak duality, holds unconditionally, and this is not difficult to prove. To establish $d^\star = p^\star$, i.e., strong duality, requires additional assumptions and is more difficult to prove. The following theorem does this with a condition we can use.

Theorem 5. *[Theorem 17 and 18 [48]] If there is an $x \in X$ such that $f(x) < \infty$ and g is bounded above in a neighborhood of Lx, then $p^\star = d^\star$. Furthermore, if $p^\star = d^\star < \infty$, the infimum of $\inf_{y^*\in Y^*}\{f^*(L^* y^*) + g^*(-y^*)\}$ is attained.*

5.2 Spaces

Throughout this section, $\|\cdot\|_1, \|\cdot\|_2, \dots$ denote unspecified finite dimensional norms. As all finite dimensional norms are equivalent, we do not bother to precisely specify which norms they are.

Define

$$
C(\Omega,\mathbb{R}^k) = \left\{ \vec{\phi} : \Omega \to \mathbb{R}^k \;\middle|\; \vec{\phi} \text{ is continuous}, \|\vec{\phi}\|_\infty = \max_{\mathbf{x}\in\Omega} \|\vec{\phi}(\mathbf{x})\|_1 < \infty \right\}.
$$

Then $C(\Omega,\mathbb{R}^k)$ is a Banach space equipped with the norm $\|\cdot\|_\infty$. We define $C(\Omega,\mathbb{R}^{k\times d})$ likewise. If $\vec{\phi} \in C(\Omega,\mathbb{R}^k)$ is continuously differentiable, $\nabla_{\mathbf{x}}\vec{\phi}$ is defined on Ω°. We say $\nabla_{\mathbf{x}}\vec{\phi}$ has a continuous extension to Ω, if there is a $g \in C(\Omega,\mathbb{R}^{k\times d})$ such that $g|_{\Omega^\circ} = \nabla_{\mathbf{x}}\vec{\phi}$. Define

$$
\begin{aligned}
C^1(\Omega,\mathbb{R}^k) = \Big\{ \vec{\phi} : \mathbb{R}^d \to \mathbb{R}^k \;\Big|\; & \vec{\phi} \text{ is continuously differentiable on } \Omega^\circ, \\
& \nabla_{\mathbf{x}}\vec{\phi} \text{ has a continuous extension to } \Omega, \\
& \|\vec{\phi}\|_{\infty,\infty} = \max_{\mathbf{x}\in\Omega} \|\vec{\phi}(\mathbf{x})\|_2 + \sup_{\mathbf{x}\in\Omega} \|\nabla_{\mathbf{x}}\vec{\phi}(\mathbf{x})\|_3 < \infty \Big\}.
\end{aligned}
$$

Then $C^1(\Omega,\mathbb{R}^k)$ is a Banach space equipped with the norm $\|\cdot\|_{\infty,\infty}$.

Write $\mathcal{M}(\Omega,\mathbb{R}^k)$ for the space of $\mathbb{R}^k$-valued signed finite Borel measures on Ω, and define $\mathcal{M}(\Omega,\mathbb{R}^{k\times d})$ likewise. Write $(C(\Omega,\mathbb{R}^{k\times d}))^*$, $(C^1(\Omega,\mathbb{R}^k))^*$ for the topological dual of $C(\Omega,\mathbb{R}^{k\times d})$, $C^1(\Omega,\mathbb{R}^k)$, respectively. The standard Riesz-Markov theorem tells us that $(C(\Omega,\mathbb{R}^{k\times d}))^* = \mathcal{M}(\Omega,\mathbb{R}^{k\times d})$.

Fully characterizing $(C^1(\Omega,\mathbb{R}^k))^*$ is hard, but we do not need to do so. Instead, we only use the following simple fact. Any $g \in \mathcal{M}(\Omega,\mathbb{R}^k)$ defines the bounded linear map $\vec{\phi} \mapsto \int_\Omega \langle \vec{\phi}(\mathbf{x}), g(d\mathbf{x}) \rangle$ for any $\vec{\phi} \in C^1(\Omega,\mathbb{R}^k)$. In other words, $\mathcal{M}(\Omega,\mathbb{R}^{k\times d}) \subset (C^1(\Omega,\mathbb{R}^k))^*$ with the appropriate identification.

5.3 Operators

We redefine $\nabla_{\mathbf{x}} : C^1(\Omega, \mathbb{R}^k) \to C(\Omega, \mathbb{R}^{k \times d})$ so that $\nabla_{\mathbf{x}} \vec{\phi}$ is the continuous exntension of the usual $\nabla_{\mathbf{x}} \vec{\phi}$ to all of Ω. This makes $\nabla_{\mathbf{x}}$ a bounded linear operator. Define the dual (adjoint) operator $\nabla_{\mathbf{x}}^* : \mathcal{M}(\Omega, \mathbb{R}^{k \times d}) \to (C^1(\Omega, \mathbb{R}^k))^*$ by

$$\int_\Omega \langle \vec{\phi}(\mathbf{x}), (\nabla_{\mathbf{x}}^* \vec{\mathbf{u}})(d\mathbf{x}) \rangle = \int_\Omega \langle (\nabla_{\mathbf{x}} \vec{\phi})(\mathbf{x}), \vec{\mathbf{u}}(d\mathbf{x}) \rangle$$

for any $\vec{\phi} \in C^1(\Omega, \mathbb{R}^k)$ and $\vec{\mathbf{u}} \in \mathcal{M}(\Omega, \mathbb{R}^{k \times d})$.

Define the $\nabla_{\mathcal{G}}$ (which is simply a multiplication by a $\mathbb{R}^{k \times \ell}$ matrix) as $\nabla_{\mathcal{G}} : C^1(\Omega, \mathbb{R}^k) \to C(\Omega, \mathbb{R}^k)$. Since $C^1(\Omega, \mathbb{R}^k) \subset C(\Omega, \mathbb{R}^k)$, there is nothing wrong with defining the range of $\nabla_{\mathcal{G}}$ to be $C(\Omega, \mathbb{R}^k)$, and this still makes $\nabla_{\mathcal{G}}$ a bounded linear operator. Define the dual (adjoint) operator $\nabla_{\mathcal{G}}^* : \mathcal{M}(\Omega, \mathbb{R}^k) \to (C^1(\Omega, \mathbb{R}^k))^*$ by identifying $\nabla_{\mathcal{G}}^*$ with the transpose of the matrix that defines $\nabla_{\mathcal{G}}$. Since $\nabla_{\mathcal{G}}^*$ is simply multiplication by a matrix, we can further say

$$\nabla_{\mathcal{G}}^* : \mathcal{M}(\Omega, \mathbb{R}^k) \to \mathcal{M}(\Omega, \mathbb{R}^k) \subset (C^1(\Omega, \mathbb{R}^k))^*.$$

We write $\operatorname{div}_{\mathcal{G}} = -\nabla_{\mathcal{G}}^*$.

5.4 Zero-flux boundary condition

Let $\vec{\mathbf{m}} : \Omega \to \mathbb{R}^{k \times d}$ a smooth function. Then integration by parts tells us that

$$\int_\Omega \langle \nabla_{\mathbf{x}} \vec{\varphi}(\mathbf{x}), \vec{\mathbf{m}}(\mathbf{x}) \rangle \, d\mathbf{x} = -\int_\Omega \langle \vec{\varphi}(\mathbf{x}), \operatorname{div}_{\mathbf{x}} \vec{\mathbf{m}}(\mathbf{x}) \rangle \, d\mathbf{x}$$

holds for all smooth $\vec{\varphi} : \Omega \to \mathbb{R}^k$ if and only if $\vec{\mathbf{m}}(\mathbf{x})$ satisfies the zero-flux boundary condition, i.e., $\vec{\mathbf{m}}(\mathbf{x}) \mathbf{n}(\mathbf{x}) = 0$ for all $\mathbf{x} \in \partial\Omega$, where $\mathbf{n}(\mathbf{x})$ denotes the normal vector at $\mathbf{x}$. Here $\operatorname{div}_{\mathbf{x}}$ denotes the usual (spatial) divergence. To put it differently, $\nabla_{\mathbf{x}}^* = -\operatorname{div}_{\mathbf{x}}$ holds when the zero-flux boundary condition holds.

We generalize this notion to measures. We say $\vec{\mathbf{u}} \in \mathcal{M}(\Omega, \mathbb{R}^{k \times d})$ satisfies the zero-flux boundary condition in the weak sense if there is a $\vec{g} \in \mathcal{M}(\Omega, \mathbb{R}^k) \subset (C^1(\Omega, \mathbb{R}^k))^*$ such that

$$\int_\Omega \langle \nabla_{\mathbf{x}} \vec{\phi}(\mathbf{x}), \vec{\mathbf{u}}(d\mathbf{x}) \rangle = -\int_\Omega \langle \vec{\phi}(\mathbf{x}), \vec{g}(d\mathbf{x}) \rangle$$

holds for all $\vec{\phi} \in C^1(\Omega, \mathbb{R}^k)$. In other words, $\vec{\mathbf{u}}$ satisfies the zero-flux boundary condition if $\nabla_{\mathbf{x}}^* \vec{\mathbf{u}} \in \mathcal{M}(\Omega, \mathbb{R}^k) \subset (C^1(\Omega, \mathbb{R}^k))^*$. In this case, we write $\operatorname{div}_{\mathbf{x}}(\vec{\mathbf{u}}) = \vec{g}$ and $\operatorname{div}_{\mathbf{x}}(\vec{\mathbf{u}}) = -\nabla_{\mathbf{x}}^*(\vec{\mathbf{u}})$. This definition is often used in elasticity theory.

5.5 Duality

To establish duality, we view the dual problem (6) as the primal problem and obtain the primal problem (5) as the dual of the dual. We do this because the dual of $C(\Omega, \mathbb{R}^{k \times d})$ is known, while the dual of $\mathcal{M}(\Omega, \mathbb{R}^{k \times d})$ is difficult to characterize.

Consider the problem

$$
\begin{aligned}
\underset{\vec{\phi} \in C^1(\Omega, \mathbb{R}^k)}{\text{maximize}} \quad & \int_\Omega \langle \vec{\phi}(\mathbf{x}), \vec{\lambda}^1(\mathbf{x}) - \vec{\lambda}^0(\mathbf{x}) \rangle \, d\mathbf{x} \\
\text{subject to} \quad & \|\nabla_{\mathbf{x}} \vec{\phi}(\mathbf{x})\|_{u*} \leq 1 \\
& \|\nabla_{\mathcal{G}} \vec{\phi}(\mathbf{x})\|_{w*} \leq \alpha \quad \text{for all } \mathbf{x} \in \Omega,
\end{aligned}
\tag{10}
$$

which is equivalent to (9). Define

$$L : C^1(\Omega, \mathbb{R}^k) \to C(\Omega, \mathbb{R}^{k \times d}) \times C(\Omega, \mathbb{R}^k)$$
$$\vec{\phi} \mapsto (\nabla_{\mathbf{x}} \vec{\phi}, \nabla_{\mathcal{G}} \vec{\phi})$$

and

$$g(a, b) = \begin{cases} 0 & \text{if } \|a(\mathbf{x})\|_{u*} \leq 1, \|b(\mathbf{x})\|_{w*} \leq \alpha \text{ for all } \mathbf{x} \in \Omega \\ \infty & \text{otherwise.} \end{cases}$$

Rewrite (10) as

$$\underset{\vec{\phi}\in C^1(\Omega,\mathbb{R}^k)}{\text{maximize}} \quad \int_\Omega \langle \vec{\phi}(\mathbf{x}), \vec{\lambda}^1(\mathbf{x}) - \vec{\lambda}^0(\mathbf{x})\rangle \, d\mathbf{x} - g(L\vec{\phi}),$$

and consider its Fenchel-Rockafellar dual

$$\underset{\substack{\vec{\mathbf{u}}\in\mathcal{M}(\Omega,\mathbb{R}^{k\times d}) \\ \vec{w}\in\mathcal{M}(\Omega,\mathbb{R}^k)}}{\text{minimize}} \quad \int_\Omega \|\vec{\mathbf{u}}(\mathbf{x})\|_u + \alpha\|\vec{w}(\mathbf{x})\|_w \, d\mathbf{x}$$

$$\text{subject to} \quad -\nabla_{\mathbf{x}}^*\vec{\mathbf{u}} - \nabla_{\mathcal{G}}^*\vec{w} = \vec{\lambda}^0 - \vec{\lambda}^1 \text{ as members of } (C^1(\Omega,\mathbb{R}^k))^*.$$

The constraint

$$-\nabla_{\mathbf{x}}^*\vec{\mathbf{u}} = \vec{\lambda}^0 - \vec{\lambda}^1 - \operatorname{div}_{\mathcal{G}}\vec{w} \in \mathcal{M}(\Omega,\mathbb{R}^k) \subset (C^1(\Omega,\mathbb{R}^k))^*$$

implies

$$-\nabla_{\mathbf{x}}^*\vec{\mathbf{u}} \in \mathcal{M}(\Omega,\mathbb{R}^k),$$

i.e., $\vec{\mathbf{u}}$ satisfies the zero-flux boundary condition.

We now state the rigorous form of the primal problem (5)

$$\underset{\substack{\vec{\mathbf{u}}\in\mathcal{M}(\Omega,\mathbb{R}^{k\times d}) \\ \vec{w}\in\mathcal{M}(\Omega,\mathbb{R}^k)}}{\text{minimize}} \quad \int_\Omega \|\vec{\mathbf{u}}(\mathbf{x})\|_u + \alpha\|\vec{w}(\mathbf{x})\|_w \, d\mathbf{x}$$

$$\text{subject to} \quad \operatorname{div}_{\mathbf{x}}\vec{\mathbf{u}} + \operatorname{div}_{\mathcal{G}}\vec{w} = \vec{\lambda}^0 - \vec{\lambda}^1 \text{ as members of } \mathcal{M}(\Omega,\mathbb{R}^k) \tag{11}$$

$$\vec{\mathbf{u}} \text{ satisfies zero-flux b.c in the weak sense.}$$

The point $\vec{\phi} = 0$ satisfies the assumption of Theorem 5. Furthermore, it is easy to verify that the optimal value of the dual problem (6) is bounded. This implies strong duality, (11) is feasible, and (11) has a solution.

Given that $V(\vec{\lambda}^0,\vec{\lambda}^1) < \infty$ for all $\vec{\lambda}^0, \vec{\lambda}^1 \in \mathcal{P}(\Omega,\mathbb{R}^k)$, it is not hard to prove that $V : \mathcal{P}(\Omega,\mathbb{R}^k)\times\mathcal{P}(\Omega,\mathbb{R}^k) \to \mathbb{R}_+$ defines a metric. Interested readers can find the argument in [14].

6 Algorithmic preliminaries

Consider the Lagrangian for the vector optimal transport problems (5) and its dual (6)

$$L(\vec{\mathbf{u}},\vec{w},\vec{\phi}) = \int_\Omega \|\vec{\mathbf{u}}(\mathbf{x})\|_u + \alpha\|\vec{w}(\mathbf{x})\|_w \, d\mathbf{x}$$

$$+ \int_\Omega \langle \vec{\phi}(\mathbf{x}), \operatorname{div}_{\mathbf{x}}(\vec{\mathbf{u}})(\mathbf{x}) + \operatorname{div}_{\mathcal{G}}(\vec{w}(\mathbf{x})) - \vec{\lambda}^0(\mathbf{x}) + \vec{\lambda}^1(\mathbf{x})\rangle \, d\mathbf{x}, \tag{12}$$

which is convex with respect to $\vec{\mathbf{u}}$ and $\vec{w}$ and concave with respect to ϕ.

Finding a saddle point of (12) is equivalent to solving (5) and (6), when the primal problem (5) has a solution, the dual problem (6) has a solution, and the optimal values of (5) and (6) are equal. See [4, Theorem 7.1], [35, Theorem 2], or any reference on standard convex analysis such as [48] for further discussion on this point.

To solve the optimal transport problems, we discretize the continuous problems and apply PDHG method to solve the discretized convex-concave saddle point problem.

6.1 PDHG method

Consider the convex-concave saddle function

$$L(x,y,z) = f(x) + g(y) + \langle Ax + By, z\rangle - h(z),$$

where f, g, and h are (closed and proper) convex functions and $x \in \mathbb{R}^n$, $y \in \mathbb{R}^m$, $z \in \mathbb{R}^l$, $A \in \mathbb{R}^{l\times n}$, and $B \in \mathbb{R}^{l\times m}$. Note L is convex in x and y and concave in z. Assume L has a saddle point and step sizes $\mu,\nu,\tau > 0$ satisfy

$$1 > \tau\mu\lambda_{\max}(A^T A) + \tau\nu\lambda_{\max}(B^T B).$$

9

Write $\|\cdot\|_2$ for the standard Euclidean norm. Then the method

$$x^{k+1} = \underset{x \in \mathbb{R}^n}{\operatorname{argmin}} \left\{ L(x, y^k, z^k) + \frac{1}{2\mu} \|x - x^k\|_2^2 \right\}$$

$$y^{k+1} = \underset{y \in \mathbb{R}^m}{\operatorname{argmin}} \left\{ L(x^k, y, z^k) + \frac{1}{2\nu} \|y - y^k\|_2^2 \right\} \tag{13}$$

$$z^{k+1} = \underset{z \in \mathbb{R}^l}{\operatorname{argmax}} \left\{ L(2x^{k+1} - x^k, 2y^{k+1} - y^k, z) - \frac{1}{2\tau} \|z - z^k\|_2^2 \right\}$$

converges to a saddle point. This method is called the the Primal-Dual Hybrid Gradient (PDHG) method or the (preconditioned) Chambolle-Pock method [23, 10, 46].

PDHG can be interpreted as a proximal point method under a certain metric [30]. The quantity

$$R^k = \frac{1}{\mu} \|x^{k+1} - x^k\|_2^2 + \frac{1}{\mu} \|y^{k+1} - y^k\|_2^2 + \frac{1}{\tau} \|z^{k+1} - z^k\|_2^2$$
$$- 2\langle z^{k+1} - z^k, A(x^{k+1} - x^k) + B(y^{k+1} - y^k)\rangle.$$

is the fixed-point residual of the non-expansive mapping defined by the proximal point method. Therefore $R^k = 0$ if and only if (x^k, y^k, z^k) is a saddle point of L, and R^k decreases monotonically to 0, cf., review paper [49]. We can use R^k as a measure of progress and as a termination criterion.

6.2 Shrink operators

As the subproblems of PDHG (13) are optimization problems themselves, PDHG is most effective when these subproblems have closed-form solutions.

The problem definitions of scalar, vector, and matrix-OMT involve norms. For some, but not all, choices of norms, the "shrink" operators

$$\operatorname{shrink}(x^0; \mu) = \underset{x \in \mathbb{R}^n}{\operatorname{argmin}} \left\{ \mu\|x\| + (1/2)\|x - x^0\|_2^2 \right\}$$

have closed-form solutions. Therefore, when possible, it is useful to choose such norms for computational efficiency. Readers familiar with the compressed sensing or proximal methods literature may be familiar with this notion.

For the vector-OMT, we focus on norms

$$\|\vec{\mathbf{u}}\|_2^2 = \sum_{s=1}^d \|\mathbf{u}_s\|_2^2 \qquad \|\vec{\mathbf{u}}\|_{1,2} = \sum_{s=1}^d \|\mathbf{u}_s\|_2 \qquad \|\vec{\mathbf{u}}\|_1 = \sum_{s=1}^k \|\mathbf{u}_s\|_1$$

for $\vec{\mathbf{u}} \in \mathbb{R}^{k \times d}$ and

$$\|\vec{w}\|_2^2 = \sum_{s=1}^\ell (w_s)^2 \qquad \|\vec{w}\|_1 = \sum_{s=1}^\ell |w_s|$$

for $\vec{w} \in \mathbb{R}^\ell$. The shrink operators of these norms have closed-form solutions.

For the matrix-OMT, we focus on norms

$$\|\mathbf{U}\|_2^2 = \sum_{s=1}^d \sum_{i,j=1}^k |(U_s)_{ij}|^2 \quad \|\mathbf{U}\|_1 = \sum_{s=1}^d \sum_{i,j=1}^k |(U_s)_{ij}| \quad \|\mathbf{U}\|_{1,\mathrm{nuc}} = \sum_{s=1}^d \|U_s\|_{\mathrm{nuc}}$$

for $\mathbf{U} \in \mathcal{H}^d$ and $\|\cdot\|_2$, $\|\cdot\|_1$, and $\|\cdot\|_{1,\mathrm{nuc}}$ for $\mathbf{W} \in \mathcal{S}^\ell$, which are defined likewise. The nuclear norm $\|\cdot\|_{\mathrm{nuc}}$ is the sum of the singular values. The shrink operators of these norms have closed-form solutions.

We provide further information and details on shrink operators in the appendix.

7 Algorithms

We now present simple and parallelizable algorithms for the OMT problems. These algorithms are, in particular, very well-suited for GPU computing.

In Section 7 and 8 we deal with discretized optimization variables that approximate solutions to the continuous problems. For simplicity of notation, we use the same symbol to denote the discretizations and their continuous counterparts. Whether we are referring to the continuous variable or its discretization should be clear from context.

As mentioned in Section 6, these methods are the PDHG method applied to discretizations of the continuous problems. In the implementation, it is important to get the discretization at the boundary correct in order to respect the zero-flux boundary conditions. For interested readers, the details are provided in the appendix.

Instead of detailing the somewhat repetitive derivations of the algorithms in full, we simply show the key steps and arguments for the $\vec{\mathbf{u}}$ update of vector-OMT. The other steps follow from similar logic.

When we discretize the primal and dual vector-OMT problems and apply PDHG to the discretized Lagrangian form of (12), we get

$$\vec{\mathbf{u}}^{k+1} = \operatorname*{argmin}_{\vec{\mathbf{u}} \in \mathbb{R}^{n \times n \times k \times d}} \left\{ \sum_{ij} \|\vec{\mathbf{u}}_{ij}\|_u + \langle \vec{\phi}_{ij}, (\operatorname{div}_{\mathbf{x}} \mathbf{u})_{ij} \rangle + \frac{1}{2\mu} \|\vec{\mathbf{u}}_{ij} - \vec{\mathbf{u}}_{ij}^k\|_2^2 \right\}$$

$$= \operatorname*{argmin}_{\vec{\mathbf{u}} \in \mathbb{R}^{n \times n \times k \times d}} \left\{ \sum_{ij} \mu \|\vec{\mathbf{u}}_{ij}\|_u - \mu \langle (\nabla_{\mathbf{x}} \vec{\phi})_{ij}, \mathbf{u}_{ij} \rangle + (1/2) \|\vec{\mathbf{u}}_{ij} - \vec{\mathbf{u}}_{ij}^k\|_2^2 \right\}.$$

Since the minimization splits over the i, j indices, we write

$$\vec{\mathbf{u}}_{ij}^{k+1} = \operatorname*{argmin}_{\vec{\mathbf{u}}_{ij} \in \mathbb{R}^{k \times d}} \left\{ \mu \|\vec{\mathbf{u}}_{ij}\|_u - \mu \langle (\nabla_{\mathbf{x}} \vec{\phi})_{ij}, \mathbf{u}_{ij} \rangle + (1/2) \|\vec{\mathbf{u}}_{ij} - \vec{\mathbf{u}}_{ij}^k\|_2^2 \right\}$$

$$= \operatorname*{argmin}_{\vec{\mathbf{u}}_{ij} \in \mathbb{R}^{k \times d}} \left\{ \mu \|\vec{\mathbf{u}}_{ij}\|_u + (1/2) \|\vec{\mathbf{u}}_{ij} - (\vec{\mathbf{u}}_{ij}^k + \mu (\nabla_{\mathbf{x}} \vec{\phi})_{ij})\|_2^2 \right\}$$

$$= \operatorname{shrink}(\vec{\mathbf{u}}_{ij}^k + \mu (\nabla_{\mathbf{x}} \vec{\phi})_{ij}; \mu).$$

At the boundary, these manipulations need special care. When we incorporate ghost cells in our discretization, these seemingly cavalier manipulations are also correct on the boundary. We further explain the ghost cells and discretization in the appendix.

7.1 Scalar-OMT algorithm

The scalar-OMT algorithm can be viewed as a special case of vector-OMT or matrix-OMT algorithms. This scalar-OMT algorithm was presented in [34], but we restate it here for completeness.

First-order Method for S-OMT
    Input: Problem data λ^0, λ^1
           Initial guesses $\mathbf{u}^0$, ϕ^0 and step sizes μ, τ
    Output: Optimal $\mathbf{u}^\star$ and $\phi^\star$

---

    for $k = 1, 2, \cdots$ (Iterate until convergence)
        $\mathbf{u}_{ij}^{k+1} = \operatorname{shrink}(\mathbf{u}_{ij}^k + \mu (\nabla \Phi^k)_{ij}, \mu)$ for $i, j = 1, \ldots, n$
        $\phi_{ij}^{k+1} = \phi_{ij}^k + \tau (\operatorname{div}_{\mathbf{x}}(2\mathbf{u}^{k+1} - \mathbf{u}^k)_{ij} + \lambda_{ij}^1 - \lambda_{ij}^0)$ for $i, j = 1, \ldots, n$
    end

This method converges for step sizes $\mu, \tau > 0$ that satisfy

$$1 > \tau \mu \lambda_{\max}(-\Delta_{\mathbf{x}}).$$

11

For the particular setup of $\Omega = [0,1] \times [0,1]$ and $\Delta x = 1/(n-1)$, the bound $\lambda_{\max}(-\Delta_{\mathbf{x}}) \leq 8/(\Delta x)^2 = 8(n-1)^2$ is known. In our experiments, we use $\mu = 1/(16\tau(n-1)^2)$, a choice that ensures convergence for any $\tau > 0$. We tune τ for the fastest convergence.

7.2   Vector-OMT algorithm

Write shrink_u and shrink_w for the shrink operators with respect to $\|\cdot\|_u$ and $\|\cdot\|_w$. The vector-OMT algorithm is as follows:

First-order Method for V-OMT

    Input: Problem data $\mathcal{G}$, $\vec{\lambda}^0$, $\vec{\lambda}^1$, α

            Initial guesses $\vec{\mathbf{u}}^0$, $\vec{w}^0$, $\vec{\phi}^0$ and step sizes μ, ν, τ

    Output: Optimal $\vec{\mathbf{u}}^\star$, $\vec{w}^\star$, and $\vec{\phi}^\star$

---

    for $k = 1, 2, \cdots$ (Iterate until convergence)

        $\vec{\mathbf{u}}_{ij}^{k+1} = \text{shrink}_u(\vec{\mathbf{u}}_{ij}^k + \mu(\nabla\vec{\phi}^k)_{ij}, \mu)$ for $i,j = 1, \ldots, n$

        $\vec{w}_{ij}^{k+1} = \text{shrink}_w(\vec{w}_{ij}^k + \nu(\nabla_{\mathcal{G}}\vec{\phi}_{ij}^k), \alpha\nu)$ for $i,j = 1, \ldots, n$

        $\vec{\phi}_{ij}^{k+1} = \vec{\phi}_{ij}^k + \tau(\text{div}_{\mathbf{x}}(2\vec{\mathbf{u}}^{k+1} - \vec{\mathbf{u}}^k)_{ij} + \text{div}_{\mathcal{G}}(2\vec{w}^{k+1} - \vec{w}^k)_{ij} + \vec{\lambda}_{ij}^1 - \vec{\lambda}_{ij}^0)$

                                    for $i,j = 1, \ldots, n$

    end

This method converges for step sizes $\mu, \nu, \tau > 0$ that satisfy

$$1 > \tau\mu\lambda_{\max}(-\Delta_{\mathbf{x}}) + \tau\nu\lambda_{\max}(-\Delta_{\mathcal{G}}).$$

For the particular setup of $\Omega = [0,1] \times [0,1]$ and $\Delta x = 1/(n-1)$, the bound $\lambda_{\max}(-\Delta_{\mathbf{x}}) \leq 8(n-1)^2$ is known. Given a graph $\mathcal{G}$, we can compute $\lambda_{\max}(-\Delta_{\mathcal{G}})$ with a standard eigenvalue routine. In our experiments, we use $\mu = 1/(32\tau(n-1)^2)$ and $\nu = 1/(4\tau\lambda_{\max}(-\Delta_{\mathcal{G}}))$, a choice that ensures convergence for any $\tau > 0$. We tune τ for the fastest convergence.

7.3   Matrix-OMT algorithm

Write shrink_u and shrink_w for the shrink operators with respect to $\|\cdot\|_u$ and $\|\cdot\|_w$. The matrix-OMT algorithm is as follows:

First-order Method for M-OMT

    Input: Problem data $\mathbf{L}$, Λ^0, Λ^1, α,

            Initial guesses $\mathbf{U}^0$, $\mathbf{W}^0$, Φ^0 and step sizes μ, ν, τ

    Output: Optimal $\mathbf{U}^\star$, $\mathbf{W}^\star$, and $\Phi^\star$

---

    for $k = 1, 2, \cdots$ (Iterate until convergence)

        $\mathbf{U}_{ij}^{k+1} = \text{shrink}_u(\mathbf{U}_{ij}^k + \mu(\nabla\Phi^k)_{ij}, \mu)$ for $i,j = 1, \ldots, n$

        $\mathbf{W}_{ij}^{k+1} = \text{shrink}_w(\mathbf{W}_{ij}^k + \nu(\nabla_{\mathbf{L}}\Phi_{ij}^k), \alpha\nu)$ for $i,j = 1, \ldots, n$

        $\Phi_{ij}^{k+1} = \Phi_{ij}^k + \tau(\text{div}_{\mathbf{x}}(2\mathbf{U}^{k+1} - \mathbf{U}^k)_{ij} + \text{div}_{\mathbf{L}}(2\mathbf{W}^{k+1} - \mathbf{W}^k)_{ij} + \Lambda_{ij}^1 - \Lambda_{ij}^0)$

                                    for $i,j = 1, \ldots, n$

    end

This method converges for step sizes $\mu, \nu, \tau > 0$ that satisfy

$$1 > \tau\mu\lambda_{\max}(-\Delta_{\mathbf{x}}) + \tau\nu\lambda_{\max}(-\Delta_{\mathbf{L}}).$$

For the particular setup of $\Omega = [0,1] \times [0,1]$ and $\Delta x = 1/(n-1)$, the bound $\lambda_{\max}(-\Delta_{\mathbf{x}}) \leq 8(n-1)^2$ is known. Given $\mathbf{L}$, we can compute the value of $\lambda_{\max}(-\Delta_{\mathbf{L}})$ by explicitly forming a $k^2 \times k^2$ matrix representing the linear operator $-\Delta_{\mathbf{L}}$ and applying a standard eigenvalue routine. In our experiments, we use $\mu = 1/(32\tau(n-1)^2)$ and $\nu = 1/(4\tau\lambda_{\max}(-\Delta_{\mathbf{L}}))$, a choice that ensures convergence for any $\tau > 0$. We tune τ for the fastest convergence.
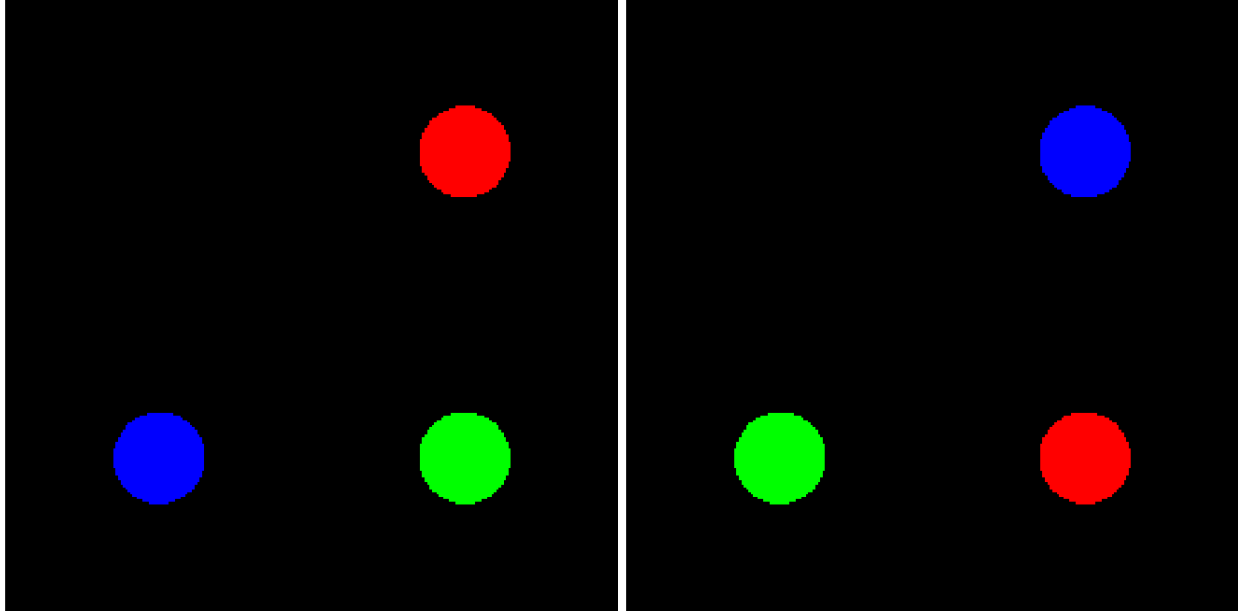
(a) $\vec{\lambda}^0$

(b) $\vec{\lambda}^1$

(c) Velocity field $\mathbf{u}$ with $c_1 = 1$, $c_2 = 1$, $c_3 = 1$, $\alpha = 1$ and norms $\|\vec{\mathbf{u}}\|_{1,2}, \|\vec{w}\|_1$. $V(\vec{\lambda}^0, \vec{\lambda}^1) = 0.57$.

(d) Velocity field $\mathbf{u}$ with $c_1 = 1$, $c_2 = 1$, $c_3 = 1$, $\alpha = 1$ and norms $\|\vec{\mathbf{u}}\|_1, \|\vec{w}\|_1$. $V(\vec{\lambda}^0, \vec{\lambda}^1) = 0.67$.

Figure 2: Color image vector-OMT example.

Grid Size	Iteration count	Time per-iter	Iteration time	τ	$V(\vec{\lambda}^0, \vec{\lambda}^1)$
32×32	0.5×10^4	$27\mu s$	$0.14s$	1	0.57
64×64	0.5×10^4	$27\mu s$	$0.14s$	1	0.57
128×128	2×10^4	$23\mu s$	$0.47s$	3	0.57
256×256	2×10^4	$68\mu s$	$1.36s$	3	0.57

Table 1: Computation cost for vector-OMT as a function of grid size.

7.4 Parallelization

For the vector-OMT, the computation for the $\vec{\mathbf{u}}$, $\vec{w}$, and $\vec{\phi}$ updates splits over the indices (i, j), i.e., the computation splits pixel-by-pixel. Furthermore, the $\vec{\mathbf{u}}$ and $\vec{w}$ updates can be done in parallel. Parallel processors handing jobs split over (i, j) must be synchronized before and after the $\vec{\phi}$ update. The same is true for the scalar-OMT and matrix-OMT.

This highly parallel and regular algorithmic structure makes the proposed algorithms very well suited for CUDA GPUs. We demonstrate this through our experiments.

8 Examples

In this section, we provide example applications of vector and matrix-OMT with numerical tests to demonstrate the effectiveness of our algorithms. As mentioned in the introduction, potential applications of vector and matrix-OMT are broad. Here, we discuss two of the simplest applications.

We implemented the algorithm on C++ CUDA and ran it on a Nvidia GPU. For convenience, we MEXed this code, i.e., the code is made available as a Matlab function. For scientific reproducibility, we release this code.

8.1 Color images

Color images in RGB format is one of the more immediate examples of vector-valued densities. At each spatial position of a 2D color image, the color is represented as a combination of the three basic colors red (R), green (G), and blue (B). We allow any basic color to change to another basic color with cost c_1 for R to G, c_2 for R to B, and c_3 for G to B. So the graph $\mathcal{G}$ as described in Section 3.1 has 3 nodes and 3 edges.

Consider the two color images on the domain $\Omega = [0, 1] \times [0, 1]$ with 256×256 discretization shown in Figures 2a and 2b. The initial and target densities $\vec{\lambda}^0$ and $\vec{\lambda}^1$ both have three disks at the same location, but with different colors. The optimal flux depends on the choice of norms. Figures 2c and 2d, show fluxes optimal with respect to different norms.

Whether it is optimal to spatially transport the colors or to change the colors depends on the parameters c_1, c_2, c_3, α as well as the norms $\| \cdot \|_u$ and $\| \cdot \|_w$. With the parameters of Figure 3a it is optimal to spatially transport the colors, while with the parameters of Figure 3b it is optimal to change the colors.

Finally, we test our algorithm on the setup of Figure 2c with grid sizes $32 \times 32, 64 \times 64, 128 \times 128$, and 256×256. Table 1 shows the number of iterations and runtime tested on a Nvidia Titan Xp GPU required to achieve a 10^{-3} precision, measured as the ratio between the duality gap and primal objective value.

8.2 Diffusion tensor imaging

Diffusion tensor imaging (DTI) is a technique used in magnetic resonance imaging. DTI captures orientations and intensities of brain fibers at each spatial position as ellipsoids and gives us a matrix-valued density. Therefore, the metric defined by the matrix-OMT provides a natural way to compare the differences between brain diffusion images. In Figure 4 we visualize diffusion tensor images by using colors to indicate different orientations of the tensors at each voxel. In this paper, we present simple 2D examples as a proof of concept and leave actual 3D imaging for a topic of future work.
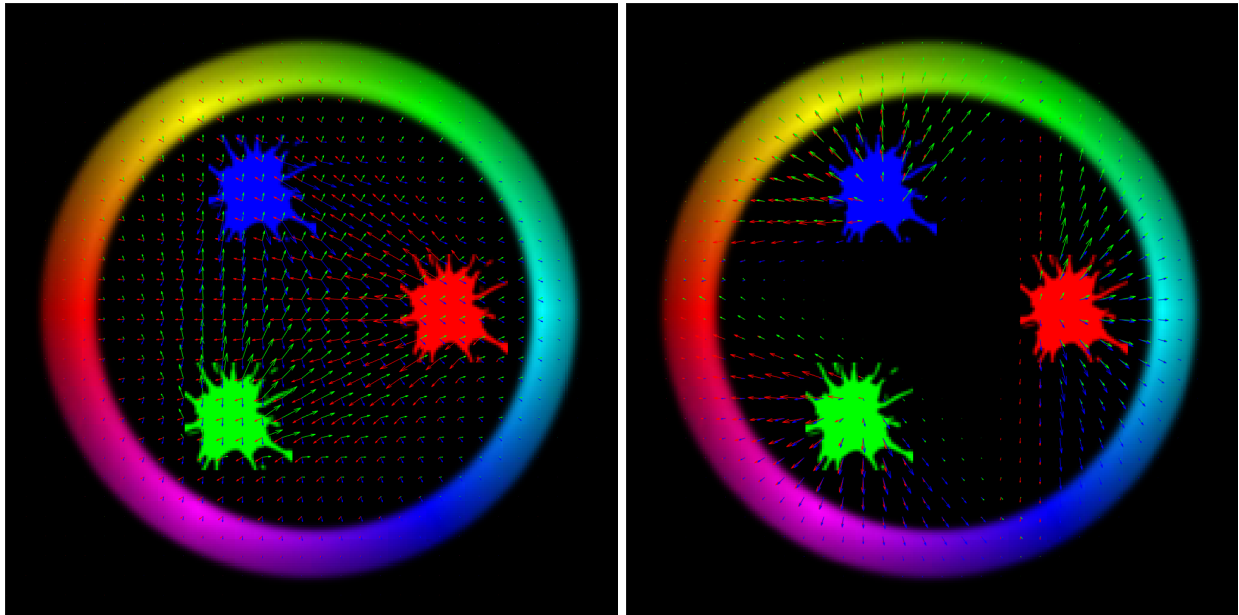
(a) $c_1 = 1, c_2 = 1, c_3 = 1, \alpha = 10$.　　　　　　　　　(b) $c_1 = 1, c_2 = 1, c_3 = 1, \alpha = 0.1$.

Figure 3: Color image vector-OMT example, with a more complicated shape and norms $\|\vec{\mathbf{u}}\|_{1,2}, \|\vec{w}\|_2$.

Parameters	$M(\Lambda^0, \Lambda^1)$	$M(\Lambda^0, \Lambda^2)$	$M(\Lambda^1, \Lambda^2)$
$\alpha = 10$	2.71	0.27	3.37
$\alpha = 3$	0.81	0.27	1.08
$\alpha = 1$	0.27	0.27	0.54
$\alpha = 0.3$	0.081	0.27	0.35
$\alpha = 0.1$	0.027	0.27	0.29

Table 2: Distances between the three images Λ^0, Λ^1 and Λ^2.

Consider three synthetic matrix-valued densities Λ^0, Λ^1, and Λ^2 in Figure 5. The densities Λ^0 and Λ^1 have mass at the same spatial location, but the ellipsoids have different shapes. The densities Λ^0 and Λ^2 have the same ellipsoids at different spatial locations.

We compute the distance between Λ^0, Λ^1, and Λ^2 for different parameters α and fixed $\mathbf{L} = [L_1, L_2]^*$ with

$$L_1 = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 2 & 0 \\ 0 & 0 & 0 \end{bmatrix}, \quad L_2 = \begin{bmatrix} 1 & 1 & 1 \\ 1 & 0 & 0 \\ 1 & 0 & 0 \end{bmatrix}.$$

Table 2, shows the results with $\|\mathbf{U}\|_2$ and $\|\mathbf{W}\|_1$ and grid size 128×128. As we can see, whether Λ^0 is more "similar" to Λ^1 or Λ^2, i.e., whether $M(\lambda^0, \lambda^1) < M(\lambda^0, \lambda^2)$ or $M(\lambda^0, \lambda^1) > M(\lambda^0, \lambda^2)$, depends on whether the cost on $\mathbf{U}$, spatial transport, is higher than the cost on $\mathbf{W}$, changing the ellipsoids.

Again, we test our algorithm on the setup of Figure 5 with $\alpha = 1$ on grid sizes $32 \times 32, 64 \times 64, 128 \times 128$, and 256×256. Table 3 shows the number of iterations and runtime tested on a Nvidia Titan Xp GPU required to achieve a precision of 10^{-3}, measured as the ratio between the duality gap and primal objective value.

9　Conclusions

In this paper, we studied the extensions of Wasserstein-1 optimal transport to vector and matrix-valued densities. This extension, as a tool of applied mathematics, is interesting if the mathematics is sound, if the

Figure 4: Example of 2D diffusion tensor images.

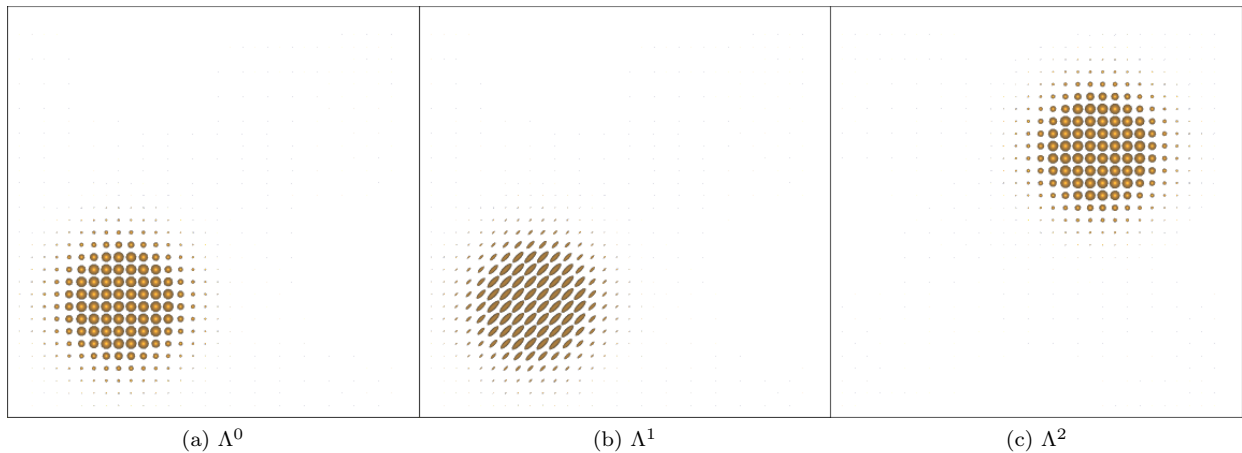
(a) Λ^0 (b) Λ^1 (c) Λ^2

Figure 5: Synthetic matrix-valued distributions.

numerical methods are good, and if the applications are interesting. In this paper, we investigated all three concerns.

From a practical viewpoint, that we can solve vector and matrix-OMT problems of realistic sizes in modest time with GPU computing is the most valuable observation. Applying our algorithms to tackle real world problems in signal/imaging processing, medical imaging, and machine learning would be interesting directions of future research.

Another interesting direction of study is quadratic regularization. In general, the solutions to vector and matrix-OMT problems are not unique. However, the regularized version of (5)

$$\begin{aligned}
\underset{\vec{\mathbf{u}},\vec{w}}{\text{minimize}} \quad & \int_\Omega \|\vec{\mathbf{u}}(\mathbf{x})\|_u + \alpha\|\vec{w}(\mathbf{x})\|_w + \epsilon(\|\vec{\mathbf{u}}(\mathbf{x})\|_u^2 + \|\vec{w}(\mathbf{x})\|_w^2)\, d\mathbf{x} \\
\text{subject to} \quad & \text{div}_\mathbf{x}(\vec{\mathbf{u}})(\mathbf{x}) + \text{div}_\mathcal{G}(\vec{w}(\mathbf{x})) = \vec{\lambda}^0(\mathbf{x}) - \vec{\lambda}^1(\mathbf{x}) \\
& \vec{\mathbf{u}} \text{ satisfies zero-flux b.c}
\end{aligned}$$

is strictly convex and therefore has a unique solution. A similar regularization can be done for matrix-OMT. As discussed in [34, 50], this form of regularization is particularly useful as a slight modification to the proposed method solves the regularized problem.

16

Grid Size	Iteration count	Time per-iter	Iteration time	τ	$M(\Lambda^0, \Lambda^1)$
32×32	1×10^4	$39 \mu s$	$0.39 s$	10	0.27
64×64	1.5×10^4	$39 \mu s$	$0.59 s$	10	0.27
128×128	2×10^4	$85 \mu s$	$1.70 s$	30	0.27
256×256	4×10^4	$330 \mu s$	$13.2 s$	60	0.27

Table 3: Computation cost for matrix-OMT as a function of grid size.

Acknowledgments

A   Discretization

Here, we discribe the discretization for the vect-OMT problem (5) and (6). The discretization for the matrix-OMT problem is similar. We consider the 2D square domain, but (with more complicated notation) our approach immediately generalizes to more general domains. Again, we use the same symbol for the discretization and its continuous counterpart.

Consider a $n \times n$ discretization of Ω with finite difference Δx in both x and y directions. Write the x and y coordinates of the points as $x_1, \ldots, x_n$ and $y_1, \ldots, y_n$. So we are approximating the domain Ω with $\{x_1, \ldots, x_n\} \times \{y_1, \ldots, y_n\}$. Write $C(x, y)$ for the $\Delta x \times \Delta x$ cube centered at (x, y), i.e.,

$$C(x, y) = \{(x', y') \in \mathbb{R}^2 \,|\, |x' - x| \le \Delta x/2, \, |y' - y| \le \Delta x/2\} \,.$$

We use a finite volume approximation for $\vec{\lambda}^0$, $\vec{\lambda}^1$, $\vec{w}$ and $\vec{\phi}$. Specifically, we write $\vec{\lambda}^0 \in \mathbb{R}^{n \times n \times k}$ with

$$\vec{\lambda}^0_{ij} \approx \int_{C(x_i, y_j)} \vec{\lambda}^0(x, y) \, dxdy,$$

for $i, j = 1, \ldots, n$. The discretizations $\vec{\lambda}^1, \vec{\phi} \in \mathbb{R}^{n \times n \times k}$ and $\vec{w} \in \mathbb{R}^{n \times n \times \ell}$ are defined the same way.

Write $\vec{\mathbf{u}} = (\vec{u}_x, \vec{u}_y)$ for both the continuous variables and their discretizations. To be clear, the subscripts of $\vec{u}_x$ and $\vec{u}_y$ do not denote differentiation. We use the discretization $\vec{u}_x \in \mathbb{R}^{(n-1) \times n \times k}$ and $\vec{u}_y \in \mathbb{R}^{n \times (n-1) \times k}$. For $i = 1, \ldots, n-1$ and $j = 1, \ldots, n$

$$\vec{u}_{x, ij} \approx \int_{C(x_i + \Delta x/2, y_j)} \vec{u}_x(x, y) \, dxdy,$$

and for $i = 1, \ldots, n$ and $j = 1, \ldots, n-1$

$$\vec{u}_{y, ij} \approx \int_{C(x_i, y_j + \Delta x/2)} \vec{u}_y(x, y) \, dxdy.$$

In defining $\vec{u}_x$ and $\vec{u}_y$, the center points are placed between the $n \times n$ grid points to make the finite difference operator symmetric.

Define the discrete spacial divergence operator $\text{div}_{\mathbf{x}}(\vec{\mathbf{u}}) \in \mathbb{R}^{n \times n \times k}$ as

$$\text{div}_{\mathbf{x}}(\vec{\mathbf{u}})_{ij} = \frac{1}{\Delta x}(\vec{u}_{x, ij} - \vec{u}_{x, (i-1)j} + \vec{u}_{y, ij} - \vec{u}_{y, i(j-1)}),$$

for $i, j = 1, \ldots, n$, where we mean $\vec{u}_{x, 0j} = \vec{u}_{x, nj} = 0$ for $j = 1, \ldots, n$ and $\vec{u}_{y, i0} = \vec{u}_{y, in} = 0$ for $i = 1, \ldots, n$. This definition of $\text{div}_{\mathbf{x}}(\vec{\mathbf{u}})$ makes the discrete approximation consistent with the zero-flux boundary condition.

For $\vec{\phi} \in \mathbb{R}^{n \times n}$, define the discrete gradient operator $\nabla_{\mathbf{x}}\vec{\phi} = ((\nabla\vec{\phi})_x, (\nabla\vec{\phi})_y)$ as

$$(\nabla\vec{\phi})_{x,ij} = (1/\Delta x)\left(\vec{\phi}_{i+1,j} - \vec{\phi}_{i,j}\right) \quad \text{for } i = 1, \ldots, n-1, \; j = 1, \ldots, n$$

$$(\nabla\vec{\phi})_{y,ij} = (1/\Delta x)\left(\vec{\phi}_{i,j+1} - \vec{\phi}_{i,j}\right) \quad \text{for } i = 1, \ldots, n, \; j = 1, \ldots, n-1.$$

So $(\nabla\vec{\phi})_x \in \mathbb{R}^{(n-1)\times n \times k}$ and $(\nabla\vec{\phi})_y \in \mathbb{R}^{n \times (n-1) \times k}$, and the $\nabla_{\mathbf{x}}$ is the transpose (as a matrix) of $-\text{div}_{\mathbf{x}}$.

Ghost cells are convenient for both describing and implementing the method. This approach is similar to that of [34]. We redefine the variable $\vec{\mathbf{u}} = (\vec{u}_x, \vec{u}_y)$ so that

$$\vec{u}_{x,ij} = \begin{cases} \vec{u}_{x,ij} & \text{for } i < n \\ 0 & \text{for } i = n \end{cases} \qquad \vec{u}_{y,ij} = \begin{cases} \vec{u}_{y,ij} & \text{for } j < n \\ 0 & \text{for } j = n \end{cases},$$

for $i, j = 1, \ldots, n$, and $\vec{u}_x, \vec{u}_y \in \mathbb{R}^{n \times n \times k}$. We also redefine $\nabla\vec{\phi} = ((\nabla\vec{\phi})_x, (\nabla\vec{\phi})_y)$ so that

$$(\nabla\vec{\phi})_{x,ij} = \begin{cases} (\nabla\vec{\phi})_{x,ij} & \text{for } i < n \\ 0 & \text{for } i = n \end{cases} \qquad (\nabla\vec{\phi})_{y,ij} = \begin{cases} (\nabla\vec{\phi})_{y,ij} & \text{for } j < n \\ 0 & \text{for } j = n \end{cases},$$

for $i, j = 1, \ldots, n$, and $(\nabla\vec{\phi})_x, (\nabla\vec{\phi})_y \in \mathbb{R}^{n \times n \times k}$.

With some abuse of notation, we write

$$\|\vec{\mathbf{u}}\|_u = \sum_{i=1}^{n}\sum_{j=1}^{n} \|(\vec{u}_{x,ij}, \vec{u}_{y,ij})\|_u \qquad \|\vec{w}\|_w = \sum_{i=1}^{n}\sum_{j=1}^{n} \|\vec{w}_{ij}\|_w.$$

Using this notation, we write the discretization of (5) as

$$\underset{\vec{\mathbf{u}}, \vec{w}}{\text{minimize}} \quad \|\vec{\mathbf{u}}\|_u + \alpha\|\vec{w}\|_w$$

$$\text{subject to} \quad \text{div}_{\mathbf{x}}(\vec{\mathbf{u}}) + \text{div}_{\mathcal{G}}(\vec{w}) = \vec{\lambda}^0 - \vec{\lambda}^1,$$

where the boundary conditions are implicitly handled by the discretization.

B Shrink operators with closed-form solutions

Define $\text{shrink}_1 : \mathbb{C} \to \mathbb{C}$ as

$$\text{shrink}_1(x; \mu) = \underset{z \in \mathbb{C}}{\text{argmin}}\left\{\mu|z| + (1/2)|z - x|^2\right\}$$

$$= \begin{cases} (1 - \mu/|x|)x & \text{for } |x| \geq \mu \\ 0 & \text{otherwise} \end{cases}$$

for $\mu > 0$. Define $\text{shrink}_2 : \mathbb{C}^k \to \mathbb{C}^k$ as

$$\text{shrink}_2(x; \mu) = \underset{z \in \mathbb{C}^k}{\text{argmin}}\left\{\mu\|z\|_2 + (1/2)\|z - x\|_2^2\right\}$$

$$= \begin{cases} (1 - \mu/\|x\|_2)x & \text{for } \|x\|_2 \geq \mu \\ 0 & \text{otherwise} \end{cases}$$

for $\mu > 0$, where $\|\cdot\|_2$ is the standard Euclidean norm. Define $\text{shrink}_{\text{nuc}} : \mathcal{C} \to \mathcal{C}$ as

$$\text{shrink}_{\text{nuc}}(X; \mu) = \underset{Z \in \mathcal{C}}{\text{argmin}}\left\{\mu\|Z\|_* + (1/2)\|Z - X\|_2^2\right\}$$

$$= U \, \text{diag}\left([(\sigma_1 - \mu)_+, \ldots, (\sigma_n - \mu)_+]\right) V^T$$

where $\|\cdot\|_*$ is the nuclear norm and $X = U\,\text{diag}\left([\sigma_1, \ldots, \sigma_n]\right)V^T$ is the singular-value decomposition of X [9]. If $x = (x_1, x_2, \ldots, x_k)$ and $\|x\| = \|x_1\|_1 + \|x_2\|_2 + \cdots + \|x_k\|_k$ for some norms $\|\cdot\|_1, \ldots, \|\cdot\|_k$, then

the shrink operators can be applied component-by-component. So if $\text{shrink}_1, \text{shrink}_2, \ldots, \text{shrink}_k$ are the individual shrink operators, then

$$\text{shrink}(x; \mu) = \begin{bmatrix} \text{shrink}_1(x_1; \mu) \\ \text{shrink}_2(x_2; \mu) \\ \vdots \\ \text{shrink}_k(x_k; \mu) \end{bmatrix}.$$

All shrink operators we consider in this paper can be built from these shrink operators. These ideas are well-known to the compressed sensing and proximal methods community [44].

However, there is a subtlety we must address when applying the shrink operators to matrix-OMT: the **U** update is defined as the minimization over $\mathcal{H}$, not all of $\mathcal{C}$, and the **W** update is defined as the minimization over $\mathcal{S}$, not all of $\mathcal{C}$. Fortunately, this is not a problem when we use unitarily invariant norms (such as the nuclear norm) thanks to the following lemma.

Lemma 1. *Let $M \in \mathcal{H}$ and $N \in \mathcal{S}$. Let $\| \cdot \|$ be a unitarily invariant norm, i.e., $\|UAV\| = \|A\|$ for any $A \in \mathcal{C}$, and $U, V \in \mathcal{C}$ unitary. Then*

$$M^+ = \operatorname*{argmin}_{A \in \mathcal{H}} \left\{ \|A\| + \frac{1}{2} \|A - M\|_2^2 \right\} = \operatorname*{argmin}_{A \in \mathcal{C}} \left\{ \|A\| + \frac{1}{2} \|A - M\|_2^2 \right\}$$

is Hermitian

$$N^+ = \operatorname*{argmin}_{A \in \mathcal{S}} \left\{ \|A\| + \frac{1}{2} \|A - N\|_2^2 \right\} = \operatorname*{argmin}_{A \in \mathcal{C}} \left\{ \|A\| + \frac{1}{2} \|A - N\|_2^2 \right\}$$

is Skew-Hermitian.

Proof. The minimum over $\mathcal{H} \subset \mathcal{C}$ is the same as the minimum over all of $\mathcal{C}$, if the minimum over all of $\mathcal{C}$ is in $\mathcal{H}$. Likewise, the minimum over $\mathcal{S} \subset \mathcal{C}$ is the same as the minimum over all of $\mathcal{C}$, if the minimum over all of $\mathcal{C}$ is in $\mathcal{S}$.

Write $\sigma : \mathcal{C} \to \mathbb{R}^n$ for the function that outputs singular values in decreasing order. Because of unitary invariance, we can write $\|X\| = f(\sigma(X))$, where $f : \mathbb{R}^n \to \mathbb{R}$ is a norm on $\mathbb{R}^n$ [56]. Furthermore, it's subdifferential can be written as [33]

$$\partial \| \cdot \|(X) = \{U \operatorname{diag}(\mu) V^T \mid \mu \in (\partial f)(\sigma(X)), \ U \operatorname{diag} \sigma(X) V^T = X \text{ is SVD}\}.$$

Write $M = U \Sigma U^T$ for M's eigenvalue decomposition, which is also its singular value decomposition. Define

$$\sigma^+ = \operatorname*{argmin}_{s} \{f(s) + (1/2)\|s - \sigma(M)\|_2^2\},$$

which implies $0 \in \partial f(\sigma^+) + \sigma^+ - \sigma(M)$. With this, we can verify that

$$0 \in U \operatorname{diag}(\partial f(\sigma^+)) U^T + U \operatorname{diag}(\sigma^+) U^T - U \operatorname{diag}(\sigma(M)) U^T,$$

i.e., $U \operatorname{diag}(\sigma^+) U^T$ satisfies the optimality conditions of the optimization problem that defines M^+. So $M^+ = U \operatorname{diag}(\sigma^+) U^T$, which is Hermitian.

Likewise, write $N = U \Lambda U^T$ for N's eigenvalue decomposition. N has orthonormal eigenvectors and its eigenvalues are purely imaginary. We separate out the magnitude and phase of Λ and write $N = U \operatorname{diag}(\sigma(N)) P U^T$. More precisely, $\sigma(N) = |\operatorname{diag}(\Lambda)|$ and P is a diagonal matrix with diagonal components $\pm i$. Then the SVD of N is $N = U \operatorname{diag}(\sigma(N)) V^T$ where $P U^T = V^T$. With the same argument as before, we conclude that $N^+ = U \operatorname{diag}(\sigma^+) V^T$ for some $\sigma^+ \in \mathbb{R}^n$. So $N^+ = U \Lambda^+ U^T$ where $\Lambda^+ = \operatorname{diag}(\sigma^+) P$ is diagonal and purely imaginary, and we conclude N^+ is skew-Hermitian. $\square$

The norm $\| \cdot \|_1$ as described in Section 6.2, is not unitarily invariant. However, its shrink operator acts element-by-element, so it is easy to arrive at a conclusion similar to that of Lemma 1.

References

[1] R. K. Ahuja, T. L. Magnanti, and J. B. Orlin, *Network Flows: Theory, Algorithms, and Applications*, Prentice Hall, 1993.

[2] L. Ambrosio, N. Gigli, and G. Savaré, *Gradient Flows: In Metric Spaces and in the Space of Probability Measures*, Springer, 2006.

[3] S. Angenent, S. Haker, and A. Tannenbaum, *Minimizing flows for the Monge–Kantorovich problem*, SIAM J. Math. Anal., 35 (2003), pp. 61–97.

[4] H. H. Bauschke, R. I. Boţ, W. L. Hare, and W. M. Moursi, *Attouch-Théra duality revisited: Paramonotonicity and operator splitting*, J. Approx. Theory, 164 (2012), pp. 1065–1084.

[5] J.-D. Benamou and Y. Brenier, *A computational fluid mechanics solution to the Monge-Kantorovich mass transfer problem*, Numer. Math., 84 (2000), pp. 375–393.

[6] J.-D. Benamou, G. Carlier, M. Cuturi, L. Nenna, and G. Peyré, *Iterative Bregman projections for regularized transportation problems*, SIAM J. Sci. Comput., 37 (2015), pp. A1111–A1138.

[7] J.-D. Benamou, B. D. Froese, and A. M. Oberman, *Numerical solution of the optimal transportation problem using the Monge-Ampere equation*, J. Comput. Phys., 260 (2014), pp. 107–126.

[8] Y. Brenier, *Polar factorization and monotone rearrangement of vector-valued functions*, Comm. Pure Appl. Math., 44 (1991), pp. 375–417.

[9] J.-F. Cai, E. J. Candès, and Z. Shen, *A singular value thresholding algorithm for matrix completion*, SIAM J. Optim., 20 (2010), pp. 1956–1982.

[10] A. Chambolle and T. Pock, *A first-order primal-dual algorithm for convex problems with applications to imaging*, J. Math. Imaging Vis., 40 (2011), pp. 120–145.

[11] Y. Chen, *Modeling and Control of Collective Dynamics: From Schrödinger bridges to Optimal Mass Transport*, PhD thesis, University of Minnesota, 2016.

[12] Y. Chen, W. Gangbo, T. T. Georgiou, and A. Tannenbaum, *On the matrix Monge-Kantorovich problem*, arXiv, (2017).

[13] Y. Chen, T. Georgiou, and M. Pavon, *Entropic and displacement interpolation: a computational approach using the Hilbert metric*, SIAM J. Appl. Math., 76 (2016), pp. 2375–2396.

[14] Y. Chen, T. T. Georgiou, L. Ning, and A. Tannenbaum, *Matricial Wasserstein-1 distance*, IEEE Control Syst. Lett., (2017).

[15] Y. Chen, T. T. Georgiou, and M. Pavon, *On the relation between optimal transport and Schrödinger bridges: A stochastic control viewpoint*, J. Optim. Theory Appl., 169 (2016), pp. 671–691.

[16] Y. Chen, T. T. Georgiou, and M. Pavon, *Optimal transport over a linear dynamical system*, IEEE Trans. Automat. Control, 62 (2017), pp. 2137–2152.

[17] Y. Chen, T. T. Georgiou, and A. Tannenbaum, *Interpolation of density matrices and matrix-valued measures: The unbalanced case*, arXiv, (2016).

[18] Y. Chen, T. T. Georgiou, and A. Tannenbaum, *Matrix optimal mass transport: a quantum mechanical approach*, arXiv, (2016).

[19] Y. Chen, T. T. Georgiou, and A. Tannenbaum, *Vector-valued optimal mass transport*, arXiv, (2016).

[20] S.-N. Chow, W. Huang, Y. Li, and H. Zhou, *Fokker-Planck equations for a free energy functional or Markov process on a graph*, Arch. Ration. Mech. Anal., 203 (2012), pp. 969–1008.

[21] S.-N. Chow, W. Li, and H. Zhou, *Entropy dissipation of Fokker-Planck equations on graphs*, arXiv, (2017).

[22] M. Cuturi, *Sinkhorn distances: Lightspeed computation of optimal transport*, in Neural Inf. Process. Syst., 2013, pp. 2292–2300.

[23] E. Esser, X. Zhang, and T. F. Chan, *A general framework for a class of first order primal-dual algorithms for convex optimization in imaging science*, SIAM J. Imaging Sci., 3 (2010), pp. 1015–1046.

[24] L. C. Evans and W. Gangbo, *Differential Equations Methods for the Monge-Kantorovich Mass Transfer Problem*, American Mathematical Society, 1999.

[25] J. H. Fitschen, F. Laus, and G. Steidl, *Transport between RGB images motivated by dynamic optimal transport*, J. Math. Imaging Vis., 56 (2016), pp. 409–429.

[26] W. Gangbo and R. J. McCann, *The geometry of optimal transportation*, Acta Math., 177 (1996), pp. 113–161.

[27] A. Genevay, M. Cuturi, G. Peyré, and F. Bach, *Stochastic optimization for large-scale optimal transport*, in Neural Inf. Process. Syst., 2016, pp. 3440–3448.

[28] E. Haber and R. Horesh, *A multilevel method for the solution of time dependent optimal transport*, Numer. Math. Theory Methods Appl., 8 (2015), pp. 97–111.

[29] S. Haker, L. Zhu, A. Tannenbaum, and S. Angenent, *Optimal mass transport for registration and warping*, Int. J. Comput. Vis., 60 (2004), pp. 225–240.

[30] B. He and X. Yuan, *Convergence analysis of primal-dual algorithms for a saddle-point problem: From contraction perspective*, SIAM J. Imaging Sci., 5 (2012), pp. 119–149.

[31] R. Jordan, D. Kinderlehrer, and F. Otto, *The variational formulation of the Fokker–Planck equation*, SIAM J. Math. Anal., 29 (1998), pp. 1–17.

[32] L. V. Kantorovich, *On the transfer of masses*, Dokl. Akad. Nauk, 37 (1942), pp. 227–229.

[33] A. Lewis, *The convex analysis of unitarily invariant matrix functions.*, J. Convex Anal., 2 (1995), pp. 173–183.

[34] W. Li, E. K. Ryu, S. Osher, W. Yin, and W. Gangbo, *A parallel method for earth mover's distance*, J. Sci. Comput., (2017).

[35] Y. Liu, E. K. Ryu, and W. Yin, *A new use of Douglas-Rachford splitting and ADMM for identifying infeasible, unbounded, and pathological conic programs*, arXiv, (2017).

[36] J. Maas, *Gradient flows of the entropy for finite Markov chains*, J. Funct. Anal., 261 (2011), pp. 2250–2292.

[37] R. J. McCann, *A convexity principle for interacting gases*, Adv. Math., 128 (1997), pp. 153–179.

[38] G. Monge, *Mémoire sur la théorie des déblais et des remblais*, De l'Imprimerie Royale, 1781.

[39] M. Mueller, P. Karasev, I. Kolesov, and A. Tannenbaum, *Optical flow estimation for flame detection in videos*, IEEE Trans. Image Process., 22 (2013), pp. 2786–2797.

[40] L. Ning and T. T. Georgiou, *Metrics for matrix-valued measures via test functions*, in IEEE Conf. Decis. Control, IEEE, 2014, pp. 2642–2647.

[41] L. Ning, T. T. Georgiou, and A. Tannenbaum, *Matrix-valued Monge-Kantorovich optimal mass transport*, in IEEE Conf. Decis. Control, IEEE, 2013, pp. 3906–3911.

[42] L. Ning, T. T. Georgiou, and A. Tannenbaum, *On matrix-valued Monge-Kantorovich optimal mass transport*, IEEE Trans. Automat. Control, 60 (2015), pp. 373–382.

[43] F. Otto and C. Villani, *Generalization of an inequality by Talagrand and links with the logarithmic Sobolev inequality*, J. Funct. Anal., 173 (2000), pp. 361–400.

[44] N. Parikh and S. Boyd, *Proximal algorithms*, Found. Trends Optim., 1 (2014), pp. 127–239.

[45] G. Peyre, L. Chizat, F.-X. Vialard, and J. Solomon, *Quantum optimal transport for tensor field processing*, European J. Appl. Math., (2017).

[46] T. Pock and A. Chambolle, *Diagonal preconditioning for first order primal-dual algorithms in convex optimization*, IEEE Intern. Conf. Comput. Vis., (2011), pp. 1762–1769.

[47] S. T. Rachev and L. Rüschendorf, *Mass Transportation Problems: Volume I: Theory*, vol. 1, Springer, 1998.

[48] R. Rockafellar, *Conjugate Duality and Optimization*, Society for Industrial and Applied Mathematics, 1974.

[49] E. K. Ryu and S. Boyd, *Primer on monotone operator methods*, Appl. Comput. Math., 15 (2016), pp. 3–43.

[50] E. K. Ryu, W. Li, P. Yin, and S. Osher, *Unbalanced and partial L_1 Monge-Kantorovich problem: A scalable parallel first-order method*, J. Sci. Comput., (2017).

[51] P. Stoica and R. L. Moses, *Spectral Analysis of Signals*, Prentice Hall, 2005.

[52] E. Tannenbaum, T. Georgiou, and A. Tannenbaum, *Signals and control aspects of optimal mass transport and the Boltzmann entropy*, in IEEE Conf. Decis. Control, IEEE, 2010, pp. 1885–1890.

[53] C. Villani, *Topics in Optimal Transportation*, American Mathematical Society, 2003.

[54] C. Villani, *Optimal Transport: Old and New*, Springer, 2008.

[55] T. Vogt and J. Lellmann, *Measure-valued variational models with applications to diffusion-weighted imaging*, arXiv, (2017).

[56] J. von Neumann, *Some matrix inequalities and metrization of matric-space*, Tomsk Univ. Rev., 1 (1937), pp. 286–300.