

---

# Structured Transforms Across Spaces with Cost-Regularized Optimal Transport

---

Othmane Sebbouh  
ENS, CNRS, ENSAE

Marco Cuturi  
Apple

Gabriel Peyré  
ENS and CNRS

## Abstract

Matching a source to a target probability measure is often solved by instantiating a *linear* optimal transport (OT) problem, parameterized by a ground cost function that quantifies discrepancy between points. When these measures live in the same metric space, the ground cost often defaults to its distance. When instantiated across two different spaces, however, choosing that cost in the absence of aligned data is a conundrum. As a result, practitioners often resort to solving instead a quadratic Gromov-Wasserstein (GW) problem. We exploit in this work a parallel between GW and *cost-regularized* OT, the regularized minimization of a linear OT objective parameterized by a ground cost. We use this cost-regularized formulation to match measures across two different Euclidean spaces, where the cost is evaluated between transformed source points and target points. We show that several quadratic OT problems fall in this category, and consider enforcing *structure* in *linear* transform (e.g. sparsity), by introducing structure-inducing regularizers. We provide a proximal algorithm to extract such transforms from unaligned data, and demonstrate its applicability to single-cell spatial transcriptomics/multiomics matching tasks.

## 1 INTRODUCTION

Optimal Transport (OT) is by now an established tool in the machine learning playbook, one that plays a key role when matching probability distributions. OT has played a prominent role in generative modeling

(Montavon et al., 2016; Genevay et al., 2018; Tong et al., 2023; Neklyudov et al., 2023a; Lipman et al., 2022; Neklyudov et al., 2023b), adversarial training (Sinha et al., 2018; Wong et al., 2019), domain adaptation (Courty et al., 2017), neuroscience (Janati et al., 2020), or single-cell modeling (Schiebinger et al., 2019; Tong et al., 2020; Bunne et al., 2023). However, OT is not without drawbacks, and we target in this work an important pain point that hinders the use of OT in practical tasks: the challenges of choosing a ground cost function to compare two measures, particularly when they are supported on *heterogeneous* and *high-dimensional* Euclidean spaces.

**Choosing a cost or going quadratic?** OT is most often used to match probability distributions supported on the same space, in which case the ground cost is often picked to be an (exponentiated) distance on that space. Yet, an increasing number of applications require handling multimodality, i.e. comparing distributions supported on heterogeneous spaces, as in multi-omics single-cell data matching (Demetci et al., 2022; Klein et al., 2023a). For such problems, defining a ground cost *across* two spaces (e.g. between vectors of different dimensions), in the absence of other knowledge, is understandably difficult. As a result, practitioners often bypass this issue by resorting to the Gromov-Wasserstein (GW) framework (Mémoli, 2011). GW instead instantiates a *quadratic* problem in the space of couplings that involves *two* within ground cost functions, one for each space. However, solving the GW problem is fraught with computational challenges: it requires minimizing a concave quadratic function on the space of couplings, an NP hard problem. Moreover, some of the classic linear OT machinery, e.g. the existence of *Monge* maps, does not directly translate to quadratic settings, as shown by Vayer (2020); Dumont et al. (2022) who discussed their existence for a variety of costs. This is a very active field of research, as shown by recent contributions on the properties of entropic GW (Zhang et al., 2022; Rioux et al., 2023); closed-form formulas for Gaussians (Salmona et al., 2022; Le et al., 2022), or scalable computational schemes (Scetbon et al., 2022; Nekrashevich et al., 2023).

**High-dimensional issues.** Another issue with OT comes from the significant statistical challenges that arise when dealing with high-dimensional problems (Dudley, 1966; Weed and Bach, 2019; Genevay et al., 2019). One way around this is performing offline dimensionality reduction, typically via PCA or VAEs, or more elaborate schemes that carry out a joint estimation of a projection and transport (Bonneel et al., 2015; Niles-Weed and Rigollet, 2022; Paty and Cuturi, 2019; Deshpande et al., 2019; Le et al., 2019). A recent line of work proposes to leverage sparsity in displacements (Cuturi et al., 2023; Klein et al., 2023b). These challenges only add up when dealing with heterogeneous spaces and taking the quadratic route advocated in GW. For example, the standard iterative algorithm to solve the GW problem requires solving a linear OT problem *at each iteration* (Peyré et al., 2016), and is only able to handle dimensionality as a way to speed up these routines (Scetbon et al., 2022)

**Contributions.** In this work, we target both issues in the case where practitioners are faced with the task of transporting points across two modalities, where each lies in a high-dimensional space.

- We use a result by Paty and Cuturi (2020), who showed that minimizing a concave function over couplings can be expressed as the minimization of a cost-regularized OT problem, to show that most common quadratic OT formulations can be unified under that “cost-regularized OT” lens.
- We propose to leverage this fact to impose explicitly structure on costs using a suitable convex regularization. Focusing on dot-product costs parameterized by a *linear transform* across spaces, we propose a simple alternated minimization algorithm, PROX-ROT, to solve cost-regularized OT when the regularizer is evaluated on such transforms. PROX-ROT is particularly efficient when the regularization function has a closed-form proximity operator.
- We use this framework to introduce sparsifying norms on such transforms and show this is equivalent to variants of GW which only depend on an *adaptively selected* subset of dimensions: Using PROX-ROT with such norms consists in iteratively performing a feature selection step followed by an OT plan computation step until convergence. This makes cost-regularized OT particularly well-suited to high-dimensional problems, where linear OT and GW usually struggle.
- Building on Dumont et al. (2022), we show that there exist Monge maps for the linear cost-regularized OT problem. Extending the framework of entropic Monge maps (Pooladian and Niles-Weed, 2021), we derive entropic Monge maps *across spaces*, for linear cost-regularized OT, and demonstrate that they converge to the ground truth Monge maps under suitable assumptions.
- We apply our methods to toy and real-world data. We use sparsifying and low-rank regularizations to solve high-dimensional multiomics single-cell data integration tasks while improving over the best-known OT baselines, and we take advantage of the cost-regularized OT objective structure to use SGD to solve a large-scale spatial transcriptomics problem.

The closest setting to ours appeared in Alvarez-Melis et al. (2019), who proposed to use a constraint on the spectrum of the linear transforms and showed the equivalence between the Gromov-Wasserstein problem and solving a linear OT problem with a dot product cost using a Frobenius norm constraint on the linear operator. The PROX-ROT algorithm in Section 3 using the nuclear or the rank regularization results (up to reparametrization) in a similar algorithm to the one proposed in Alvarez-Melis et al. (2019). We take the alternative *regularization* route and consider general regularizations. Notably, we pay particular attention to sparsifying norms acting directly on the entries of the linear operator, and draw links with GW-like problems on a subset of the dimensions.

**Notations.** In what follows, we consider  $(\mathcal{X}, d_{\mathcal{X}})$  and  $(\mathcal{Y}, d_{\mathcal{Y}})$ , two metric spaces, and  $(\alpha, \beta) \in \mathcal{P}(\mathcal{X}) \times \mathcal{P}(\mathcal{Y})$ , two probability measures with compact supports. The space of couplings  $\Pi(\alpha, \beta)$  is defined as the probability distributions on  $(\mathcal{X} \times \mathcal{Y})$  whose marginals are  $\alpha$  and  $\beta$  (Santambrogio, 2015):

$$\Pi(\alpha, \beta) = \left\{ \gamma \in \mathcal{P}(\mathcal{X} \times \mathcal{Y}) : (\pi_x)_\# \gamma = \alpha, (\pi_y)_\# \gamma = \beta \right\},$$

where  $\pi_x$  and  $\pi_y$  are the projections of  $\mathcal{X} \times \mathcal{Y}$  onto  $\mathcal{X}$  and  $\mathcal{Y}$  respectively. We denote the linear OT cost between  $\alpha$  and  $\beta$  with cost  $c \in \mathcal{C}(\mathcal{X} \times \mathcal{Y})$  as

$$\mathcal{W}_c(\alpha, \beta) \triangleq \min_{\pi \in \Pi(\alpha, \beta)} \int_{\mathcal{X} \times \mathcal{Y}} c(x, y) d\pi(x, y),$$

which is a *linear* problem in  $\pi$ . When  $c = d_{\mathcal{X}}^p$  and  $\mathcal{X} = \mathcal{Y}$ ,  $\mathcal{W}_c^{1/p}$  defines a distance between probability measures for all  $p \geq 1$  (Villani, 2003).

## 2 OT ACROSS SPACES AS A CONCAVE MINIMIZATION PROBLEM

While the classical OT objective is a *linear* function of the coupling  $\pi$ , existing approaches to comparing probability distributions across spaces all require minimizing a *concave* function over  $\Pi(\alpha, \beta)$ .

**Problem 2.1.** *Let  $\mathcal{Q}$  be a concave function over  $\Pi(\alpha, \beta)$ . We define the concave minimization over*

couplings problem as

$$\min_{\pi \in \Pi(\alpha, \beta)} \mathcal{Q}(\pi). \quad (\text{QOT})$$

Remarkably, concave minimization problems over couplings are closely linked to linear OT via convex duality.

**Proposition 2.2** (Remark 2 in Paty and Cuturi (2020)). *Let  $\mathcal{Q}$  be a proper usc concave function over  $\Pi(\alpha, \beta)$ . Writing  $\mathcal{Q}^*$  for its convex conjugate, one has:*

$$\begin{aligned} \inf_{\pi \in \Pi(\alpha, \beta)} \mathcal{Q}(\pi) &= \inf_{\substack{\pi \in \Pi(\alpha, \beta) \\ c \in \mathcal{C}(\mathcal{X} \times \mathcal{Y})}} \int c(x, y) d\pi + (-\mathcal{Q})^*(-c) \\ &= \inf_{c \in \mathcal{C}(\mathcal{X} \times \mathcal{Y})} \mathcal{W}_c(\alpha, \beta) + (-\mathcal{Q})^*(-c). \end{aligned}$$

In other words, concave minimization problems over couplings can be reformulated as the minimization over cost functions of a linear OT problem plus a cost regularization. Consequently, we can define the following cost-regularized reformulation of Problem 2.1.

**Problem 2.3.** *Let  $\mathcal{R}$  be a convex function over  $\mathcal{C}(\mathcal{X} \times \mathcal{Y})$ . We define the cost-regularized optimal transport problem as*

$$\text{ROT}(\alpha, \beta) \triangleq \inf_{c \in \mathcal{C}(\mathcal{X} \times \mathcal{Y})} \mathcal{W}_c(\alpha, \beta) + \mathcal{R}(c). \quad (\text{ROT})$$

Note that, trivially,  $\text{ROT} = \text{QOT}$  when  $\mathcal{R}(c) = (-\mathcal{Q})^*(-c)$ . We show next how well-known OT formulations used to compare probability distributions across spaces can be reduced to either  $\text{QOT}$  or  $\text{ROT}$ . We say that two minimization problems are equivalent, or that one is an instance of the other, when they have the same minimizers in  $\Pi(\alpha, \beta)$ .

## 2.1 Examples

**Sturm’s distance.** Sturm (2006) defined a distance between metric measured spaces (MMS) that quantifies how “isometric“ they are. Given two MMS  $(\mathcal{X}, d_X, \alpha)$  and  $(\mathcal{Y}, d_Y, \beta)$ , Sturm defined the squared distance:

$$\min_{\pi \in \Pi(\alpha, \beta)} \min_{\substack{(Z, d_Z) \in \mathcal{C} \\ \phi: \mathcal{X} \rightarrow Z \\ \psi: \mathcal{Y} \rightarrow Z}} \int_{\mathcal{X} \times \mathcal{Y}} d_Z^2(\phi(x), \psi(y)) d\pi(x, y),$$

where  $\phi: \mathcal{X} \rightarrow Z$  and  $\psi: \mathcal{Y} \rightarrow Z$  are constrained to be isometries and  $(Z, d_Z)$  is a metric space. As such, Sturm’s distance is an instance of  $\text{ROT}$  with

$$\mathcal{R}(c) = \begin{cases} 0 & \text{if } \exists (Z, \phi, \psi) : c(x, y) = d_Z^2(\phi(x), \psi(y)) \\ \infty & \text{otherwise.} \end{cases}$$

**Wasserstein Procrustes.** A more tractable variant of Sturm’s distance when  $\mathcal{X}, \mathcal{Y} \subset \mathbb{R}^d$  is the so-called

Wasserstein Procrustes problem (Zhang et al., 2017; Grave et al., 2019). It is defined as

$$\min_{\pi \in \Pi(\alpha, \beta)} \min_{\mathbf{C} \in \mathcal{O}_d} \int_{\mathcal{X} \times \mathcal{Y}} \|\mathbf{C}x - y\|^2 d\pi(x, y). \quad (1)$$

where  $\mathcal{O}_d$  is the orthogonal manifold. Wasserstein Procrustes defines a distance between  $\mathcal{X}$  and  $\mathcal{Y}$  up to a rotation. It is an instance of  $\text{ROT}$  with

$$\mathcal{R}(c) = \begin{cases} 0 & \text{if } \exists \mathbf{C} \in \mathcal{O}_d : c(x, y) = \|\mathbf{C}x - y\|^2 \\ \infty & \text{otherwise.} \end{cases}$$

Remarkably, (1) is a weighted orthogonal Procrustes problem, and can be cast as an instance of  $\text{QOT}$ .

**Proposition 2.4.** *Let  $\mathcal{X}, \mathcal{Y} \subset \mathbb{R}^d$ . Then the Wasserstein Procrustes distance is an instance of the concave minimization problem (QOT) with*

$$\mathcal{Q}(\pi) = \int \|\mathbf{U}(\pi)\mathbf{V}(\pi)^\top x - y\|^2 d\pi,$$

where  $\mathbf{U}(\pi)\mathbf{\Sigma}(\pi)\mathbf{V}(\pi)^\top$  is an SVD of  $\int_{\mathcal{X} \times \mathcal{Y}} yx^\top d\pi$ .

**Arbitrary Transformations.** While Sturm considers isometries and the Wasserstein Procrustes approach uses rotations, any type of a transformation of  $\alpha$  and/or  $\beta$  could be used. For example, Cohen and Guibas (1999) studied a large range of transformations, including linear ones and translations. Finally, learning generative models with OT (Genevay et al., 2018; Salimans et al., 2018) with learned features to define costs is another instance of  $\text{ROT}$ .

**The Gromov-Wasserstein distance.** Given two costs  $c_X$  and  $c_Y$ , the Gromov-Wasserstein (GW) distance (Mémoli, 2011) is defined as

$$\min_{\pi \in \Pi(\alpha, \beta)} \int_{(\mathcal{X} \times \mathcal{Y})^2} (c_X(x, x') - c_Y(y, y'))^2 d\pi d\pi. \quad (2)$$

Mémoli (2011) showed that when  $c_X$  and  $c_Y$  are distances, (2) defines a distance between metric measured spaces up to isometry. This problem is not concave in general, but when  $c_X$  and  $c_Y$  define a conditionally negative kernel, the function  $\pi \mapsto \langle Q(\pi), \pi \rangle$ , where

$$Q((x, y), (x', y')) := (c_X(x, x') - c_Y(y, y'))^2,$$

is a concave quadratic function of  $\pi$  (Dumont et al., 2022, Proposition 1). In this case, (2) is an instance of  $\text{QOT}$  with  $\mathcal{Q}(\pi) = \langle Q(\pi), \pi \rangle$ .

Two common cases that fit this setting are the inner product and the squared Euclidean distance. In this case, it has been recently shown (Vayer, 2020) that the concave minimization problem has a simple cost-regularized equivalent.

**Proposition 2.5.** Let  $\mathcal{X} \subset \mathbb{R}^{d_x}$  and  $\mathcal{Y} \subset \mathbb{R}^{d_y}$ . For all  $\mathbf{M} \in \mathbb{R}^{d_y \times d_x}$ , define  $c_{\mathbf{M}}(x, y) = -\langle \mathbf{M}x, y \rangle$ .

- Let  $c_{\mathcal{X}} = \langle \cdot, \cdot \rangle_{\mathcal{X}}$  and  $c_{\mathcal{Y}} = \langle \cdot, \cdot \rangle_{\mathcal{Y}}$ . The problems

$$\min_{\pi \in \Pi(\alpha, \beta)} \int (\langle x, x' \rangle - \langle y, y' \rangle)^2 d\pi d\pi \quad (\text{GW-IP})$$

$$\text{and } \min_{\substack{\pi \in \Pi(\alpha, \beta) \\ \mathbf{M} \in \mathbb{R}^{d_y \times d_x}}} \int -\langle \mathbf{M}x, y \rangle d\pi + \frac{1}{2} \|\mathbf{M}\|_F^2 \quad (\text{GW-IP-ROT})$$

are equivalent: *GW-IP* is an instance of *ROT* with

$$\mathcal{R}(c) = \begin{cases} \frac{1}{2} \|\mathbf{M}\|_F^2 & \text{if } \exists \mathbf{M} \in \mathbb{R}^{d_y \times d_x} : c = c_{\mathbf{M}} \\ \infty & \text{otherwise.} \end{cases}$$

- Let  $c_{\mathcal{X}} = \|\cdot - \cdot\|_{\mathcal{X}}^2$  and  $c_{\mathcal{Y}} = \|\cdot - \cdot\|_{\mathcal{Y}}^2$ . The problems

$$\min_{\pi \in \Pi(\alpha, \beta)} \int (\|x - x'\|^2 - \|y - y'\|^2)^2 d\pi d\pi \quad (3)$$

$$\text{and } \min_{\substack{\pi \in \Pi(\alpha, \beta) \\ \mathbf{M} \in \mathbb{R}^{d_y \times d_x}}} \int -\langle \mathbf{M}x, y \rangle - \|x\|^2 \|y\|^2 d\pi + \frac{1}{2} \|\mathbf{M}\|_F^2$$

are equivalent. Hence, (3) is an instance of *ROT* with

$$\mathcal{R}(c) = \begin{cases} \frac{1}{2} \|\mathbf{M}\|_F^2 & \text{if } \exists \mathbf{M} : c = c_{\mathbf{M}} - \|x\|^2 \|y\|^2 \\ \infty & \text{otherwise.} \end{cases}$$

## 2.2 Leveraging cost-regularized OT

In this work, we aim to incorporate additional structure to OT across spaces by leveraging the *ROT* formulation. Inspired by GW, we will constrain our costs to be linear and consider new regularization functions. We will see that depending on the regularization function, we can draw interesting links to GW-like problems by writing the *ROT* problem in its equivalent *QOT* form.

**Entropic regularization.** A preferred way to compute OT and GW distances in practice is adding an entropic regularization to the linear OT term (Cuturi, 2013; Peyré et al., 2016). Given an  $\varepsilon > 0$  we can follow the same procedure and add a regularization to *ROT*. The entropic-regularized *ROT* problem writes

$$\min_{c \in \mathcal{C}(\mathcal{X} \times \mathcal{Y})} \mathcal{W}_c^\varepsilon(\alpha, \beta) + \mathcal{R}(c) \quad (4)$$

$$\text{where } \mathcal{W}_c^\varepsilon(\alpha, \beta) \triangleq \min_{\pi \in \Pi(\alpha, \beta)} \int c d\pi + \varepsilon \text{KL}(\pi || \alpha \otimes \beta),$$

with  $\text{KL}(\pi || \alpha \otimes \beta) = \int \log\left(\frac{d\pi}{d\alpha \otimes d\beta}\right) d\pi$ . In this case, the entropic version of *QOT* is simply

$$\min_{\pi \in \Pi(\alpha, \beta)} \mathcal{Q}(\pi) + \varepsilon \text{KL}(\pi || \alpha \otimes \beta). \quad (5)$$

**Fused ROT.** In some cases, we may even have an inter-space cost that can be used alongside the learned cost. This is the case in the context of the fused Gromov-Wasserstein cost (Vayer et al., 2020). Such a cost  $\tilde{c} : \mathcal{C}(\mathcal{X} \times \mathcal{Y}) \rightarrow \mathbb{R}$ , which isn't learned, can be naturally integrated as (with  $\eta \geq 0$ )

$$\text{ROT}_\varepsilon^{\tilde{c}}(\alpha, \beta) = \min_{c \in \mathcal{C}(\mathcal{X} \times \mathcal{Y})} \mathcal{W}_{c+\eta\tilde{c}}^\varepsilon(\alpha, \beta) + \mathcal{R}(c). \quad (6)$$

## 3 STRUCTURED OT ACROSS SPACES

### 3.1 Linear cost-regularized OT

Let  $\mathcal{X} \subset \mathbb{R}^{d_x}$  and  $\mathcal{Y} \subset \mathbb{R}^{d_y}$ . From this point on, we consider linear costs parameterized by a matrix  $\mathbf{M} \in \mathbb{R}^{d_y \times d_x}$ :  $c_{\mathbf{M}}(x, y) = -\langle \mathbf{M}x, y \rangle$ . In this case we can write (4) as

$$\min_{\substack{\pi \in \Pi(\alpha, \beta) \\ \mathbf{M} \in \mathbb{R}^{d_y \times d_x}}} \int -\langle \mathbf{M}x, y \rangle d\pi + \mathcal{R}(\mathbf{M}) + \varepsilon \text{KL}(\pi || \alpha \otimes \beta), \quad (\text{L-ROT}_\varepsilon)$$

where we directly define  $\mathcal{R}$  on  $\mathbb{R}^{d_y \times d_x}$  for brevity.

**Proposition 3.1.** The linear *ROT* problem (*L-ROT* $_\varepsilon$ ) is equivalent to the *QOT* problem (5) with

$$\mathcal{Q}(\pi) = -\mathcal{R}^* \left( \int_{\mathcal{X} \times \mathcal{Y}} yx^\top d\pi(x, y) \right),$$

where  $\mathcal{R}^*$  is the convex conjugate of  $\mathcal{R}$ .

Since for low values of  $\varepsilon$ , (*L-ROT* $_\varepsilon$ ) is a concave minimization problem, solving it exactly is generally out of reach. However, we can approximate it using a simple alternated minimization procedure on the objective. Initializing at a  $\pi_0 \in \Pi(\alpha, \beta)$ , let

$$\begin{aligned} \mathbf{M}_{k+1} &= \operatorname{argmin}_{\mathbf{M} \in \mathbb{R}^{d_y \times d_x}} \int -\langle \mathbf{M}x, y \rangle d\pi_k + \mathcal{R}(\mathbf{M}) \\ \pi_{k+1} &= \operatorname{argmin}_{\pi \in \Pi(\alpha, \beta)} -\langle \mathbf{M}_{k+1}x, y \rangle d\pi + \varepsilon \text{KL}(\pi || \alpha \otimes \beta) \end{aligned} \quad (7)$$

**Proposition 3.2.** Let  $\mathcal{X} = \{x_i\}_{i=1}^m \subset \mathbb{R}^{d_x}$ ,  $\mathcal{Y} = \{y_j\}_{j=1}^n \subset \mathbb{R}^{d_y}$ ,  $\alpha = \sum_{i=1}^m \alpha_i \delta_{x_i}$  and  $\beta = \sum_{j=1}^n \beta_j \delta_{y_j}$ . Let  $\varepsilon > 0$  and  $\mathcal{R}$  be strongly convex. Then, any limit point of  $(M_k, \pi_k)$  defined in (7) is a stationary point of the objective in (*L-ROT* $_\varepsilon$ ).

In practice, the  $\pi$ -step can be carried out efficiently using Sinkhorn's algorithm (Cuturi, 2013) with the cost  $c_{\mathbf{M}_k}$ . The  $\mathbf{M}$ -step is in general trickier since it might involve a difficult optimization problem depending on the regularization function  $\mathcal{R}$ .

### 3.2 Proximal ROT

We consider regularizations of the form  $\mathcal{R}(\mathbf{M}) = \frac{1}{2} \|\mathbf{M}\|_F^2 + \lambda g(\mathbf{M})$ . We aim to solve

$$\min_{\substack{\pi \in \Pi(\alpha, \beta) \\ \mathbf{M} \in \mathbb{R}^{d_y \times d_x}}} - \int \langle \mathbf{M}x, y \rangle d\pi + \frac{\|\mathbf{M}\|_F^2}{2} + \lambda g(\mathbf{M}) \quad (8) \\ + \varepsilon \text{KL}(\pi \| \alpha \otimes \beta),$$

where  $g : \mathbb{R}^{d_y \times d_x} \rightarrow \mathbb{R}$ . Note that when  $\lambda = 0$ , per Prop. 2.5, (8) is equivalent to (GW-IP). The minimization with respect to  $\mathbf{M}$  in (8) is a well-studied problem in the optimization and machine learning literature (Bauschke and Combettes, 2011).

**Lemma 3.3.** *For any  $\pi$ , the solution of (8) in  $\mathbf{M}$  is*

$$\mathbf{M}(\pi) \triangleq \text{prox}_{\lambda g} \left( \int_{\mathcal{X} \times \mathcal{Y}} yx^\top d\pi \right),$$

where  $\text{prox}_h(x) \triangleq \text{argmin}_{z \in \mathbb{R}^d} \frac{1}{2} \|x - z\|^2 + h(z)$ .

The alternated minimization scheme (7) becomes:

$$\begin{aligned} \mathbf{M}_{k+1} &= \text{prox}_{\lambda g} \left( \int_{\mathcal{X} \times \mathcal{Y}} yx^\top d\pi_k \right) \quad (\text{Prox-ROT}) \\ \pi_{k+1} &= \text{argmin}_{\pi \in \Pi(\alpha, \beta)} - \int_{\mathcal{X} \times \mathcal{Y}} \langle \mathbf{M}_{k+1}x, y \rangle d\pi_k \\ &\quad + \varepsilon \text{KL}(\pi_k \| \alpha \otimes \beta) \end{aligned}$$

As a result, the  $\mathbf{M}$ -step in Prox-ROT can be efficiently implemented when the proximal operator of  $g$  can be computed in closed form.

**Statistical benefits of sparsity.** Since OT famously suffers from the curse of dimensionality – the sample complexity for estimating OT maps is  $\mathcal{O}(n^{-1/d})$  (Niles-Weed and Rigollet, 2022; Pooladian and Niles-Weed, 2021) – the main factor we consider in choosing the regularization, beyond computational efficiency, is using inductive biases to reduce *effective dimension* of the OT problem.

**$\ell_1$  and  $\ell_{1,2}$  regularizations.** Among the choices of  $g$  where the  $\mathbf{M}$ -step can be implemented efficiently are the  $\ell_1$  and  $\ell_{1,2}$  regularizations, where  $\|\mathbf{M}\|_1 = \sum_{i=1}^{d_x} \sum_{j=1}^{d_y} |\mathbf{M}_{ij}|$  and  $\|\mathbf{M}\|_{1,2} = \sum_{i=1}^{d_x} \|\mathbf{M}_{:i}\|_2$ . Their proximal operators are given by (Chierchia et al., 2023)

$$\begin{aligned} \text{prox}_{\lambda \|\cdot\|_1}(\mathbf{M})_{ij} &= \text{sign}(\mathbf{M}_{ij}) (|\mathbf{M}_{ij}| - \lambda)_+ \\ \text{prox}_{\lambda \|\cdot\|_{1,2}}(\mathbf{M})_{:i} &= \left( 1 - \frac{\lambda}{\max\{\|\mathbf{M}_{:i}\|_2, \lambda\}} \right) \mathbf{M}_{:i}, \end{aligned}$$

where  $(a)_+ = \max(a, 0)$  and  $\mathbf{A}_{:i}$  denotes the  $i$ th column of  $\mathbf{A}$ . Intuitively, the reason behind using  $\ell_1$  regularization is to select the features in the spaces (or point

clouds)  $\mathcal{X}$  and  $\mathcal{Y}$  that are the most helpful in aligning those spaces. Similarly, using  $\ell_{1,2}$  regularization, we can discard features in one of the spaces alone.

This is best seen by examining the iterations of Prox-ROT in each of these cases: For  $\ell_1$ , the prox operator in the  $\mathbf{M}$  iteration results in a matrix  $\mathbf{M}_{k+1}$  where some entries  $\mathbf{M}_{k+1}(i, j)$  can be set to 0, and as a result the update of the OT plan  $\pi_{k+1}$  is oblivious to the corresponding features  $x_i$  and  $y_j$ ; similarly, for  $\ell_{1,2}$  iterations, the prox sets some columns of  $\mathbf{M}_{:i}$  to 0, and as a result the update of the OT plan doesn't use the corresponding dimensions  $x_i$ . In both cases, Prox-ROT alternates between a feature selection step and a linear OT step.

Rewriting (8) in its QOT form, we show that these regularizations result in problems that are closely related to the inner-product GW problem (GW-IP).

**Proposition 3.4.** *Consider Problem (8) with  $\varepsilon = 0$ .*

•  **$\ell_{1,2}$ -regularization:** *with  $g(M) = \|M\|_{1,2}$ , (8) with  $\varepsilon = 0$  is equivalent to QOT with*

$$\begin{aligned} \mathcal{Q}(\pi) &= \int_{(\mathcal{X} \times \mathcal{Y})^2} (\langle x_{I_\pi}, x'_{I_\pi} \rangle - \langle y, y' \rangle)^2 d\pi d\pi - \frac{\lambda^2}{2} |I_\pi| \\ &\quad - \int_{\mathcal{X}} \langle x_{I_\pi}, x'_{I_\pi} \rangle^2 d\alpha + \lambda \sum_{i \in I_\pi} \left\| \int_{\mathcal{X} \times \mathcal{Y}} x_i y d\pi(x, y) \right\|, \end{aligned}$$

where  $I_\pi = \{i : \left\| \int_{\mathcal{X} \times \mathcal{Y}} x_i y d\pi \right\| > \lambda\}$ . Moreover, with  $(\pi^*, M^*)$  a solution to (8), (8) is equivalent to

$$\min_{\pi \in \Pi(\alpha, \beta)} \int_{(\mathcal{X} \times \mathcal{Y})^2} (\langle x_{I_{\pi^*}}, x'_{I_{\pi^*}} \rangle - \langle y, y' \rangle)^2 d\pi^* d\pi.$$

•  **$\ell_1$ -regularization:** *With  $g(M) = \|M\|_1$ , (8) with  $\varepsilon = 0$  is equivalent to QOT with*

$$\begin{aligned} \mathcal{Q}(\pi) &= \int_{(\mathcal{X} \times \mathcal{Y})^2} \sum_{(i,j) \in I_\pi} (x_i x'_i - y_j y'_j)^2 d\pi d\pi - \frac{\lambda^2}{2} |I_\pi| \\ &\quad - \sum_{(i,j) \in I_\pi} \left( \int_{\mathcal{X} \times \mathcal{Y}} x_i^2 + y_j^2 d\alpha d\beta - \lambda \left| \int_{\mathcal{X} \times \mathcal{Y}} y_j x_i d\pi \right| \right), \end{aligned}$$

where  $I_\pi = \{(i, j) : \int_{\mathcal{X} \times \mathcal{Y}} x_i y_j d\pi > \lambda\}$ . Moreover, with  $(\pi^*, M^*)$  a solution to (8), (8) is equivalent to

$$\min_{\pi \in \Pi(\alpha, \beta)} \int_{(\mathcal{X} \times \mathcal{Y})^2} \sum_{(i,j) \in I_{\pi^*}} (x_i x'_i - y_j y'_j)^2 d\pi^* d\pi.$$

**Nuclear and Rank regularizations.** Two other possible choices of regularizations are the nuclear norm and rank regularizations. Let  $\mathbf{M} = \mathbf{U}^\top \mathbf{\Sigma} \mathbf{V}$  be an SVD of  $\mathbf{M}$ , where  $\mathbf{\Sigma}$  contains the vector of singular vectors in decreasing order  $\sigma = (\sigma_i)_{i \in [d_x]}$  of  $\mathbf{M}$ . The nuclear

norm of  $\mathbf{M}$  is defined as  $\|\mathbf{M}\|_* = \|\sigma\|_1$ . The proximal operator of  $\|\cdot\|_*$  and  $\text{rk}(\cdot)$  are respectively

$$\begin{aligned}\text{prox}_{\lambda\|\cdot\|_*}(\mathbf{M}) &= \mathbf{U}^\top \text{prox}_{\lambda\|\cdot\|_1}(\sigma) \mathbf{V} \\ \text{prox}_{\lambda\text{rk}}(\mathbf{M}) &= \mathbf{U}^\top \text{prox}_{\lambda\|\cdot\|_0}(\sigma) \mathbf{V},\end{aligned}$$

where  $\text{prox}_{\lambda\|\cdot\|_0}(\sigma)_i = \sigma_i \mathbb{1}_{\{\sigma_i^2 > 2\lambda\}}$ . Consider a solution  $(\mathbf{M}^*, \pi^*)$  of (8) with the nuclear norm. Then  $\mathbf{M}^* = \text{prox}_{\lambda\|\cdot\|_*}(\int yx^\top d\pi^*)$ . Let  $\lambda$  be large enough so that  $\text{rk}(\mathbf{M}^*) = r > 0$ . Denote by  $\tilde{\mathbf{U}} = \mathbf{U}_{:, :r}$  (resp.  $\tilde{\mathbf{V}} = \tilde{\sigma} \odot \mathbf{V}_{:, :r}$ ) the restriction of  $\mathbf{U}$  (resp.  $\mathbf{V}$ ) to its first  $r$  lines (resp. columns multiplied elementwise with  $\tilde{\sigma}$ , the restriction of  $\sigma$  to its first  $r$  elements). Then, we can write the cost of the linearized problem as  $c_{\mathbf{M}^*}(x, y) = -\langle \tilde{\mathbf{V}}x, \tilde{\mathbf{U}}y \rangle$ .  $\tilde{\mathbf{V}}x, \tilde{\mathbf{U}}y \in \mathbb{R}^r$ , thus the nuclear norm allows to similarly reduce the effective dimension of the problem from potential large  $d_x, d_y$  to a small  $r$  determined by the magnitude of the regularization  $\lambda$ .

Note that since computing the SVD at each iteration of **Prox-ROT** can be costly when  $d_x$  and  $d_y$  are large, we suggest in extremely high-dimensional problems to directly parametrize  $\mathbf{M}$  as a low-rank matrix  $\mathbf{M} = \mathbf{M}_2^\top \mathbf{M}_1$ , where  $r \ll \min(d_x, d_y)$ , and use the alternated minimization scheme (7). We write the corresponding iterations in the Appendix. Despite making the problem non-convex in  $\mathbf{M}$ , we see in our applications that the resulting method works well.

## 4 ENTROPIC MONGE MAPS FOR ROT

Optimal transport (Monge, 1781) seeks a map  $T: \mathcal{X} \rightarrow \mathcal{Y}$  that minimizes the average displacement cost  $c(x, T(x))$  between two measures  $\alpha$  and  $\beta$ , i.e. finding an OT coupling of the form  $\pi^* = (\text{id}, T)_\# \alpha$ . Although such a (so-called) Monge map does not exist for all costs, Brenier (1991) showed that in the case where  $c(x, y) = \frac{1}{2}\|x - y\|^2$ , if  $\alpha$  has a density, the optimal map exists, is unique, and can be written as the gradient of a convex function. A convenient approach to approximate that Monge map can be found in entropic regularization. Pooladian and Niles-Weed (2021) showed that one can build a map  $T_\varepsilon$  using entropy regularized transport (with regularization strength  $\varepsilon > 0$ ) such that  $T_\varepsilon \xrightarrow{\varepsilon \rightarrow 0} T$ . We show that there always exists a Monge map for linear ROT, and propose a formulation for an entropic map for this problem.

### 4.1 Monge maps for linear ROT

Dumont et al. (2022) recently showed that there exists a Monge map for the **GW-IP** problem, which, as seen in Prop. 2.5 is a special case of linear ROT. Here, we

extend their reasoning to show the existence of Monge maps for the general linear ROT problem.

**Proposition 4.1.** *Let  $\alpha \in \mathcal{P}(\mathbb{R}^{d_x})$  and  $\beta \in \mathcal{P}(\mathbb{R}^{d_y})$  with compact supports and  $d_x \geq d_y$ . Assume that  $\alpha \ll \mathcal{L}^{d_x}$ , the Lebesgue measure. Then there exists a map  $T: \mathcal{X} \rightarrow \mathcal{Y}$  such that  $\pi^* = (\text{id}, T)_\# \alpha$  and  $(\pi^*, \mathbf{M}(\pi^*))$  is optimal for **(L-ROT) $_\varepsilon$**  with  $\varepsilon = 0$ .*

*Proof.* Let  $(\mathbf{M}^*, \pi^*)$  be optimal for **(L-ROT) $_\varepsilon$**  with  $\varepsilon = 0$ . Then  $\pi^*$  is a solution to

$$\min_{\pi \in \Pi(\alpha, \beta)} \int_{\mathcal{X} \times \mathcal{Y}} -\langle \mathbf{M}^* x, y \rangle d\pi(x, y),$$

and (Dumont et al., 2022, Theorem 4) showed that for all  $\mathbf{M} \in \mathbb{R}^{d_y \times d_x}$ , under the assumptions of Prop. 4.1, there exists a Monge map for the cost  $c_{\mathbf{M}}(x, y) = -\langle \mathbf{M}x, y \rangle$  between  $\alpha$  and  $\beta$ .  $\square$

Now that the existence of a Monge map has been established, a natural question is whether, and how, we can approximate it using entropy-regularization. In the following sections, we define an entropic map for **L-ROT) $_\varepsilon$**  and show its convergence to a Monge map for this problem under suitable assumptions. Then, we discuss its convergence for the sparse and low-rank regularizations we considered in Section 3.2.

### 4.2 Entropic Monge maps for linear ROT

**Definition 4.2** (Entropic map for **L-ROT) $_\varepsilon$** ). *Let  $(\pi_\varepsilon^*, M_\varepsilon^*)$  be a solution of **L-ROT) $_\varepsilon$** . Let  $\varepsilon' > 0$ . Define*

$$T_{\varepsilon, \varepsilon'}(x) = \frac{\int y \exp((g_{\varepsilon, \varepsilon'}(y) + \langle M_\varepsilon^* x, y \rangle) / \varepsilon') d\beta(y)}{\int \exp((g_{\varepsilon, \varepsilon'}(y) + \langle M_\varepsilon^* x, y \rangle) / \varepsilon') d\beta(y)}. \quad (9)$$

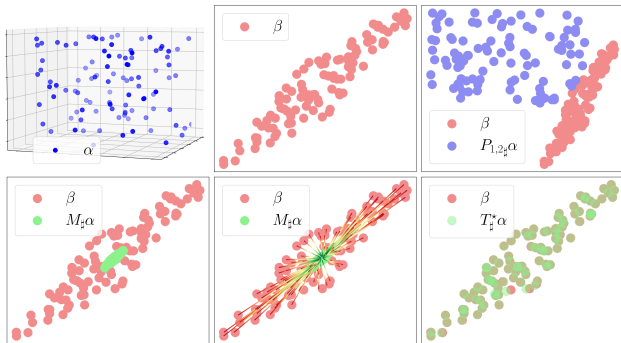
Here,  $(f_{\varepsilon, \varepsilon'}, g_{\varepsilon, \varepsilon'})$  are Sinkhorn potentials for the inner product cost between  $M_\varepsilon^* \# \alpha$  and  $\beta$  with an  $\varepsilon'$  entropic regularization: They are a solution of the problem

$$\begin{aligned}\max_{(f, g) \in \mathcal{C}(X) \times \mathcal{C}(Y)} & \int_X f(x) d\mathbf{M}_\varepsilon^* \alpha(x) + \int_Y g(y) d\beta(y) \\ -\varepsilon' \int_{X \times Y} & \exp\left(\frac{f(x) + g(y) + \langle x, y \rangle}{\varepsilon'}\right) d\mathbf{M}_\varepsilon^* \alpha \otimes \beta(x, y).\end{aligned}$$

#### 4.2.1 Computing the entropic map in practice

An important point in the theoretical definition of the entropic map (9) is to use two different regularization parameters  $\varepsilon, \varepsilon' > 0$ . In practice (as we do in our applications), we can simplify the definition and use a single  $\varepsilon$ . Given samples  $\{x_i\}_{i=1}^n \sim \alpha$  and  $\{y_i\}_{i=1}^n \sim \beta$ :

- Find a primal solution  $(\pi^*, \mathbf{M}_\varepsilon^*)$  and a dual solution  $(f, g)$  of the discrete **L-ROT) $_\varepsilon$**  problem with



**Figure 1:** Illustration of the entropic map (10) for GW-IP from a 3D to a 2D point cloud with  $\varepsilon \approx 0$ .

$\hat{\alpha} = \sum_{i=1}^n \alpha_i \delta_{x_i}$  and  $\hat{\beta} = \sum_{j=1}^m \beta_j \delta_{y_j}$  using (7).

- Define the entropic map as

$$T_\varepsilon(x) = \frac{\sum_{j=1}^m y_j \exp((\mathbf{g}_j + \langle \mathbf{M}_\varepsilon^* x, y_j \rangle) / \varepsilon)}{\sum_{k=1}^m \exp((\mathbf{g}_k + \langle \mathbf{M}_\varepsilon^* x, y_k \rangle) / \varepsilon)}. \quad (10)$$

**Illustration.** In Figure 1, we showcase the action of the entropic map (10) when  $\varepsilon \approx 0$ . We consider two point clouds  $\alpha$  and  $\beta$  on  $\mathbb{R}^3$  and  $\mathbb{R}^2$ . The entropic map implicitly acts in two steps: it first transforms the 3D point cloud  $\alpha$  into a 2D point cloud  $\mathbf{M}_\# \alpha$  which is aligned with  $\beta$ . It then maps the points from  $\mathbf{M}_\# \alpha$  to  $\beta$  using the inner-product entropic map.

#### 4.2.2 Convergence theory

**The case where  $\mathbf{M}^*$  is full rank.** Assume that  $d_x \geq d_y$  and consider the case where  $\text{rk}(\mathbf{M}^*) = d_y$ . Dumont et al. (2022) and Vayer (2020) showed that if  $\mathbf{M}^*$  is full rank, then a Monge map for GW-IP is given by  $T = -\nabla f \circ \mathbf{M}^*$ , where  $f$  is a Kantorovitch potential for  $\mathcal{W}_{\langle \cdot, \cdot \rangle}^\varepsilon(\mathbf{M}_\# \alpha, \beta)$ . The entropic map for  $L\text{-ROT}_\varepsilon$  verifies an analog result.

**Lemma 4.3.** *Let  $(f_{\varepsilon, \varepsilon'}, g_{\varepsilon, \varepsilon'})$  be Sinkhorn potentials for  $\mathcal{W}_{\langle \cdot, \cdot \rangle}^{\varepsilon'}(\mathbf{M}_\# \alpha, \beta)$ . Then, we can rewrite (9) as*

$$T_{\varepsilon, \varepsilon'} = -\nabla f_{\varepsilon, \varepsilon'} \circ \mathbf{M}_\varepsilon^*$$

Using the previous lemma and results from Pooladian and Niles-Weed (2021); Zhang et al. (2022), we have the following convergence result.

**Proposition 4.4.** *Let  $\alpha, \beta \in \mathcal{P}(\mathcal{X}) \times \mathcal{P}(\mathcal{Y})$  be two measures with compact supports. Assume that they are regular enough (namely that they verify assumptions (A1-3) from Pooladian and Niles-Weed (2021)). Let  $\varepsilon, \varepsilon' > 0$  and consider the map  $T_{\varepsilon, \varepsilon'}$  defined in (9). Then, we have that  $T_{\varepsilon, \varepsilon'} \xrightarrow{L^2(\alpha)} T_{\varepsilon, 0}$  as  $\varepsilon' \rightarrow 0$ , where  $T_{\varepsilon, 0}$  is a Monge map for  $\mathcal{W}_{c_{\mathbf{M}_\varepsilon^*}}^0(\alpha, \beta)$ , and  $c_{\mathbf{M}_\varepsilon^*}(x, y) = -\langle \mathbf{M}_\varepsilon^* x, y \rangle$ . Moreover, we have along a subsequence that  $T_{\varepsilon, 0} \xrightarrow{L^2(\alpha)} T_{0, 0}$  as  $\varepsilon \rightarrow 0$ , where  $T_{0, 0}$  is a Monge*

map for  $L\text{-ROT}_\varepsilon$ . Hence, along a subsequence,

$$\lim_{\varepsilon \rightarrow 0} \lim_{\varepsilon' \rightarrow 0} T_{\varepsilon, \varepsilon'} = T_{0, 0} \quad \text{in } L^2(\alpha).$$

**The case where  $\mathbf{M}^*$  is not full rank.** The case where  $\mathbf{M}^*$  is not full rank is more involved and with limited use in practice, as the theory in Dumont et al. (2022) requires solving an OT problem between conditional probabilities, which are not accessible in practice. In the appendix, we simulate problems where we explicitly constrain  $\mathbf{M}^*$  to being rank-deficient. We test on Gaussians (for which ground-truth OT maps are known (Salmona et al., 2022)) and show that rank-deficiency has no effect on the convergence of the entropic map.

**Sparsifying transforms and entropic Monge maps.** As we saw in the previous section, we can show the convergence of the entropic map for  $L\text{-ROT}_\varepsilon$  when  $\mathbf{M}^*$  is full rank. With sparsifying and low-rank norms, it is unlikely that such a matrix is full rank. However, when using  $\ell_{1,2}$  regularization, our goal is to operate feature selection on  $\mathcal{X}$ , so that we don't actually care about the dimensions on which the matrix  $\mathbf{M}_\varepsilon^*$  is 0 (which correspond to features we want to discard). Thus we can restrict our study to the existence of Monge maps between  $P_{I_\#} \alpha$  and  $\beta$  for the cost  $c_{\mathbf{M}_I^*}(x, y) = -\langle \mathbf{M}_I^* x, y \rangle$ , where  $I = \{i \in [d_x] : \|\mathbf{M}_{:,i}\| \neq 0\}$ ,  $P_I$  is the projection operator on the dimensions  $I$ , and  $\mathbf{M}_I$  the restriction of  $\mathbf{M}$  to the columns indexed by  $I$ . Note, though, that this matrix could still be rank-deficient.


Similarly, nuclear norm regularization inherently makes  $\mathbf{M}^*$  low-rank. As we saw earlier it results in a linear OT problem with an inner product cost between  $\tilde{\mathbf{U}}_\# \alpha$  and  $\tilde{\mathbf{V}}_\# \beta$  (i.e.  $\mathcal{W}_{\langle \cdot, \cdot \rangle}(\tilde{\mathbf{U}}_\# \alpha, \tilde{\mathbf{V}}_\# \beta)$ ). If we are interested in low-dimensional representations of  $\alpha$  and  $\beta$ , there always exists a Monge map between  $\tilde{\mathbf{U}}_\# \alpha$  and  $\tilde{\mathbf{V}}_\# \beta$ , and the entropic map converges to the Monge map since it is associated with the inner-product cost on  $\mathbb{R}^r$ , where  $r = \text{rk}(\mathbf{M}^*)$  (Pooladian and Niles-Weed, 2021).

## 5 APPLICATIONS

In all of the applications below, we use our proximal approach to compute transforms, jointly with the entropic map (10) to displace points from one space to the other in and/or out of sample. That map's  $\varepsilon$  regularization is selected using cross-validation on the training set using a grid in [5e-3, 1e-3, 5e-4, 1e-4, 5e-5, 1e-5].

### 5.1 Multi-omics data Integration: Sparse and Low-Rank Transforms

**Sparse Transforms.** We consider the scGM dataset (Cheow et al., 2016) containing the gene expression and DNA methylation modalities for human somatic



**Figure 2:** **Left:** Convergence of **Prox-ROT** with different sparsifying regularizations  $\mathcal{R}$  on the scGM dataset. X-axis: # updates of the OT plan. Y-axis: LTA (higher is better).  $\lambda = 0$  corresponds to standard **GW-IP**.  $\varepsilon = 0.05$ . **Right:**  $\mathbf{M}_\varepsilon^*$ . % zeros: *Top left*  $\lambda = 0$  : 6%, *right*  $\ell_{1,2}(\lambda = 7.5)$  : 67%; *Bottom left*  $\ell_1(\lambda = 1.5)$  : 71%, *right*  $\ell_1(\lambda = 2)$  : 84%.


cells (Welch et al., 2017). The goal is to match the 177 samples of the dataset across modalities using an entropic map from the gene expression ( $d_x = 34$ ) to the DNA methylation domain ( $d_y = 27$ ). Performance is measured in terms of Label Transfer Accuracy, a common way to evaluate single-cell data integration tasks (Demetci et al., 2022) (see Appendix). We compare solving **GW-IP** using nested Sinkhorn iterations (Peyré et al., 2016) to **Prox-ROT** with the  $\ell_1$  and  $\ell_{1,2}$  regularizations in (8). Adding a sparsity prior to the map  $\mathbf{M}$  results in a better performance (Figure 2). Using an  $\ell_1$  or  $\ell_{1,2}$  regularization also outperforms all the methods in (Demetci et al., 2022, Figure 2).

**Low-Rank Transforms.** We consider the Neurips 2021 multimodal single-cell integration dataset (Lance et al., 2022). We use the Site 1/Donor 1 stored in the MOSCOT package (Klein et al., 2023a). Like scGM, the task consists in mapping 6,224 cells from the chromatin accessibility ( $d_y = 8,000$ ) to the gene expression domain ( $d_x = 2,000$ ).

We consider random subsets of the data of sizes 25, 50, 100, 250, 500, 1000, 6224. On each subset, we run **Prox-ROT** with rank constraint between 5 and 12. We display for each subset the best-performing rank on average in Figure 3 (that rank is always 11 or less). Performance is measured in terms of FOSCTTM (fraction of samples closer than the true match) depending on the subset size. Compared to **GW**, low-rank **Prox-ROT** handles settings where  $n \ll d_y$  much better than **GW**. Note that **Prox-ROT-rk** still has benefits in the high  $n$  regime since the dimension of the OT problem solved at each iteration is at most 12 vs. 2,000 for **GW**.

## 5.2 Spatial Transcriptomics: Stochastic Fused **GW-IP**

An important benefit of the cost-regularized formulation of **GW-IP** is that it lends itself ideally to stochastic optimization. Indeed, let  $\alpha = \sum_{i=1}^m \alpha_i \delta_{x_i}$  and



**Figure 3:** Average performance (FOSCTTM - lower is better) across sampled subsets of **GW** vs. **Prox-ROT** depending on the size of training data. **Left:** Training performance on each subset size. **Right:** Evaluation on the full dataset of the map computed using a subset.

$\beta = \sum_{j=1}^n \beta_j \delta_{y_j}$  be two discrete distributions. Then, using entropic regularization, we can rewrite **GW-IP-ROT** using the dual formulation of linear OT as

$$\min_{\substack{\mathbf{M} \in \mathbb{R}^{d_y \times d_x} \\ (f,g) \in \mathbb{R}^m \times \mathbb{R}^n}} \sum_{i=1}^m f_i \alpha_i + \sum_{j=1}^n g_j \beta_j + \frac{1}{8} \|\mathbf{M}\|_F^2 \quad (11)$$

$$- \varepsilon \sum_{i,j} \exp\left(\frac{f_i + g_j + \langle \mathbf{M}x_i, y_j \rangle + \eta \tilde{c}(\tilde{x}_i, \tilde{y}_j)}{\varepsilon}\right) \alpha_i \beta_j.$$

where  $\eta > 0$ , and as in (6), we can use an additional known inter-space cost  $\tilde{c}$ , which corresponds to using Fused **GW-IP**. We can solve this problem using (epoch) stochastic gradient descent ascent.

Alg.	val $\rho$	test $\rho$	F1 macro	F1 micro	F1 weighted
LFGW	0.365	0.443	0.576	0.720	0.714
ULFGW	<b>0.379</b>	<b>0.463</b>	0.582	0.733	0.724
SFGW-IP	0.335	0.436	<b>0.609</b>	<b>0.768</b>	<b>0.765</b>

**Table 1:** Performance of different **GW** variants on a spatial transcriptomics task. Details for LFGW and ULFGW can be found in Scetbon et al. (2023). SFGW-IP refers to solving (11) via epoch SGDA.

We reproduce the experimental setting of Scetbon et al. (2023). The goal is to align cells from two coronal sections of a mouse brain (Shi et al., 2023). The linear OT term in the fused formulation comes from a 30d PCA computed in gene expression space. We use validation Pearson correlation for hyperparameter selection. We show in Table 1 that the stochastic approach, while having a lower per-iteration cost, is competitive with all **GW** variants, and significantly exceeds them in F1 scores (see appendix for more details).

**Conclusion.** We have leveraged a cost-regularized OT perspective on **GW** problems to propose new methods for OT across spaces that can induce structure (e.g. sparsity) on the transformation from one space to another. Our injection of sparsity into the transform  $\mathbf{M}$  is unrelated to the sparsity observed naturally for couplings when solving the original OT problem, from LP duality or from other regularizations, as in



e.g. (Blondel et al., 2018; Liu et al., 2022), nor is it related to that obtained for same-space Monge displacements by the MBO estimator (Cuturi et al., 2023). Our use of low-rank constraints for transform  $\mathbf{M}$  is also unrelated to the low rank constraint on couplings introduced by (Scetbon et al., 2021, 2022). The ability to add structural assumptions on across-space transforms opens up new perspectives to use transport *across* high-dimensional modalities. We also showed the existence of Monge maps for our formulation, and demonstrated their applicability with entropic maps. We used a proximal alternated minimization algorithm together with structure-inducing regularizations, demonstrating applicability to single-cell multiomics data matching and spatial transcriptomics tasks.

## References

- Alvarez-Melis, D., Jegelka, S., and Jaakkola, T. S. (2019). Towards optimal transport with global invariances. In *The 22nd International Conference on Artificial Intelligence and Statistics*, pages 1870–1879. PMLR. 2
- Bauschke, H. H. and Combettes, P. L. (2011). Convex analysis and monotone operator theory in hilbert spaces. *CMS Books in Mathematics*. 5
- Bertsekas, D. P. (1997). Nonlinear programming. *Journal of the Operational Research Society*, 48(3):334–334. 14
- Blondel, M., Seguy, V., and Rolet, A. (2018). Smooth and sparse optimal transport. In Storkey, A. and Perez-Cruz, F., editors, *Proceedings of the Twenty-First International Conference on Artificial Intelligence and Statistics*, volume 84 of *Proceedings of Machine Learning Research*, pages 880–889. PMLR. 9
- Bonneel, N., Rabin, J., Peyré, G., and Pfister, H. (2015). Sliced and radon wasserstein barycenters of measures. *Journal of Mathematical Imaging and Vision*, 51:22–45. 2
- Brenier, Y. (1991). Polar factorization and monotone rearrangement of vector-valued functions. *Communications on pure and applied mathematics*, 44(4):375–417. 6, 16
- Bunne, C., Stark, S. G., Gut, G., Del Castillo, J. S., Levesque, M., Lehmann, K.-V., Pelkmans, L., Krause, A., and Rätsch, G. (2023). Learning single-cell perturbation responses using neural optimal transport. *Nature Methods*, pages 1–10. 1
- Cheow, L. F., Courtois, E. T., Tan, Y., Viswanathan, R., Xing, Q., Tan, R. Z., Tan, D. S., Robson, P., Loh, Y.-H., Quake, S. R., et al. (2016). Single-cell multimodal profiling reveals cellular epigenetic heterogeneity. *Nature methods*, 13(10):833–836. 7
- Chierchia, G., Chouzenoux, E., Combettes, P. L., and Pesquet, J.-C. (2023). The proximity operator repository. user’s guide. 5
- Chizat, L., Roussillon, P., Léger, F., Vialard, F.-X., and Peyré, G. (2020). Faster wasserstein distance estimation with the sinkhorn divergence. *Advances in Neural Information Processing Systems*, 33:2257–2269. 17
- Cohen, S. and Guibas, L. (1999). The earth mover’s distance under transformation sets. In *Proceedings of the Seventh IEEE International Conference on Computer Vision*, volume 2, pages 1076–1083. IEEE. 3

- Courty, N., Flamary, R., Habrard, A., and Rakotomamonjy, A. (2017). Joint distribution optimal transportation for domain adaptation. *Advances in neural information processing systems*, 30. [1](#)
- Cuturi, M. (2013). Sinkhorn distances: Lightspeed computation of optimal transport. *Advances in neural information processing systems*, 26:2292–2300. [4](#)
- Cuturi, M., Klein, M., and Ablin, P. (2023). Monge, Bregman and occam: Interpretable optimal transport in high-dimensions with feature-sparse maps. In *Proceedings of the 40th International Conference on Machine Learning*, volume 202, pages 6671–6682. PMLR. [2](#), [9](#)
- De Philippis, G. (2013). *Regularity of optimal transport maps and applications*, volume 17. Springer Science & Business Media. [17](#)
- Demetci, P., Santorella, R., Sandstede, B., Noble, W. S., and Singh, R. (2022). Scot: single-cell multi-omics alignment with optimal transport. *Journal of Computational Biology*, 29(1):3–18. [1](#), [8](#), [20](#)
- Deshpande, I., Hu, Y.-T., Sun, R., Pyrros, A., Siddiqui, N., Koyejo, S., Zhao, Z., Forsyth, D., and Schwing, A. G. (2019). Max-sliced wasserstein distance and its use for gans. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10648–10656. [2](#)
- Dudley, R. M. (1966). Weak convergence of probabilities on nonseparable metric spaces and empirical measures on euclidean spaces. *Illinois Journal of Mathematics*, 10(1):109–126. [2](#)
- Dumont, T., Lacombe, T., and Vialard, F.-X. (2022). On the existence of monge maps for the gromov-wasserstein distance. *arXiv preprint arXiv:2210.11945*. [1](#), [2](#), [3](#), [6](#), [7](#), [17](#)
- Genevay, A., Chizat, L., Bach, F., Cuturi, M., and Peyré, G. (2019). Sample complexity of sinkhorn divergences. In *The 22nd International Conference on Artificial Intelligence and Statistics*, pages 1574–1583. [2](#)
- Genevay, A., Peyré, G., and Cuturi, M. (2018). Learning generative models with sinkhorn divergences. In *International Conference on Artificial Intelligence and Statistics, AISTATS 2018*, pages 1608–1617. [1](#), [3](#)
- Grave, E., Joulin, A., and Berthet, Q. (2019). Unsupervised alignment of embeddings with wasserstein procrustes. In *The 22nd International Conference on Artificial Intelligence and Statistics*, pages 1880–1890. PMLR. [3](#)
- Janati, H., Bazeille, T., Thirion, B., Cuturi, M., and Gramfort, A. (2020). Multi-subject meg/eeg source imaging with sparse multi-task regression. *NeuroImage*, 220:116847. [1](#)
- Klein, D., Palla, G., Lange, M., Klein, M., Piran, Z., Gander, M., Meng-Papaxanthos, L., Sterr, M., Bastidas-Ponce, A., Tarquis-Medina, M., et al. (2023a). Mapping cells through time and space with moscot. *bioRxiv*, pages 2023–05. [1](#), [8](#), [20](#)
- Klein, M., Pooladian, A.-A., Ablin, P., Ndiaye, E., Niles-Weed, J., and Cuturi, M. (2023b). Learning costs for structured monge displacements. *arXiv preprint arXiv:2306.11895*. [2](#)
- Lance, C., Luecken, M. D., Burkhardt, D. B., Cannoodt, R., Rautenstrauch, P., Laddach, A., Ubingazhibov, A., Cao, Z.-J., Deng, K., Khan, S., et al. (2022). Multimodal single cell data integration challenge: Results and lessons learned. In *NeurIPS 2021 Competitions and Demonstrations Track*, pages 162–176. PMLR. [8](#), [20](#)
- Le, K., Le, D. Q., Nguyen, H., Do, D., Pham, T., and Ho, N. (2022). Entropic gromov-wasserstein between Gaussian distributions. In *International Conference on Machine Learning*, pages 12164–12203. PMLR. [1](#)
- Le, T., Yamada, M., Fukumizu, K., and Cuturi, M. (2019). Tree-sliced variants of wasserstein distances. *Advances in neural information processing systems*, 32. [2](#)
- Lipman, Y., Chen, R. T., Ben-Hamu, H., Nickel, M., and Le, M. (2022). Flow matching for generative modeling. In *The Eleventh International Conference on Learning Representations*. [1](#)
- Liu, T., Puigcerver, J., and Blondel, M. (2022). Sparsity-constrained optimal transport. In *The Eleventh International Conference on Learning Representations*. [9](#)
- Mémoli, F. (2011). Gromov–wasserstein distances and the metric approach to object matching. *Foundations of computational mathematics*, 11(4):417–487. [1](#), [3](#)
- Monge, G. (1781). Mémoire sur la théorie des déblais et des remblais. *Mem. Math. Phys. Acad. Royale Sci.*, pages 666–704. [1](#), [6](#)
- Montavon, G., Müller, K.-R., and Cuturi, M. (2016). Wasserstein training of restricted boltzmann machines. *Advances in Neural Information Processing Systems*, 29. [1](#)
- Neklyudov, K., Brekelmans, R., Severo, D., and Makhzani, A. (2023a). Action matching: Learning stochastic dynamics from samples. In *International conference on machine learning*. [1](#)
- Neklyudov, K., Brekelmans, R., Severo, D., and Makhzani, A. (2023b). Action matching: Learning stochastic dynamics from samples. In Krause, A., Brunskill, E., Cho, K., Engelhardt, B., Sabato, S., and Scarlett, J., editors, *Proceedings of the 40th International Conference on Machine Learning*, volume

- 202 of *Proceedings of Machine Learning Research*, pages 25858–25889. PMLR. 1
- Nekrashevich, M., Korotin, A., and Burnaev, E. (2023). Neural gromov-wasserstein optimal transport. *arXiv preprint arXiv:2303.05978*. 1
- Niles-Weed, J. and Rigollet, P. (2022). Estimation of wasserstein distances in the spiked transport model. *Bernoulli*, 28(4):2663–2688. 2, 5
- Paty, F.-P. and Cuturi, M. (2019). Subspace robust wasserstein distances. In *International conference on machine learning*, pages 5072–5081. PMLR. 2
- Paty, F.-P. and Cuturi, M. (2020). Regularized optimal transport is ground cost adversarial. In *International Conference on Machine Learning*, pages 7532–7542. PMLR. 2, 3
- Peyré, G., Cuturi, M., and Solomon, J. (2016). Gromov-wasserstein averaging of kernel and distance matrices. In *International conference on machine learning*, pages 2664–2672. PMLR. 2, 4, 8
- Pooladian, A.-A. and Niles-Weed, J. (2021). Entropic estimation of optimal transport maps. *arXiv preprint arXiv:2109.12004*. 2, 5, 6, 7, 16, 17
- Rioux, G., Goldfeld, Z., and Kato, K. (2023). Entropic gromov-wasserstein distances: Stability, algorithms, and distributional limits. *arXiv preprint arXiv:2306.00182*. 1
- Salimans, T., Zhang, H., Radford, A., and Metaxas, D. (2018). Improving gans using optimal transport. In *International Conference on Learning Representations*. 3
- Salmona, A., Delon, J., and Desolneux, A. (2022). Gromov-wasserstein distances between Gaussian distributions. *Journal of Applied Probability*, 59(4). 1, 7, 19
- Santambrogio, F. (2015). Optimal transport for applied mathematicians. *Birkäuser, NY*, 55(58-63):94. 2
- Scetbon, M., Cuturi, M., and Peyré, G. (2021). Low-rank sinkhorn factorization. In *International Conference on Machine Learning*, pages 9344–9354. PMLR. 9
- Scetbon, M., Klein, M., Palla, G., and Cuturi, M. (2023). Unbalanced low-rank optimal transport solvers. *arXiv preprint arXiv:2305.19727*. 8, 20
- Scetbon, M., Peyré, G., and Cuturi, M. (2022). Linear-time gromov Wasserstein distances using low rank couplings and costs. In Chaudhuri, K., Jegelka, S., Song, L., Szepesvari, C., Niu, G., and Sabato, S., editors, *Proceedings of the 39th International Conference on Machine Learning*, volume 162 of *Proceedings of Machine Learning Research*, pages 19347–19365. PMLR. 1, 2, 9
- Schiebinger, G., Shu, J., Tabaka, M., Cleary, B., Subramanian, V., Solomon, A., Gould, J., Liu, S., Lin, S., Berube, P., et al. (2019). Optimal-transport analysis of single-cell gene expression identifies developmental trajectories in reprogramming. *Cell*, 176(4):928–943. 1
- Shi, H., He, Y., Zhou, Y., Huang, J., Maher, K., Wang, B., Tang, Z., Luo, S., Tan, P., Wu, M., et al. (2023). Spatial atlas of the mouse central nervous system at molecular resolution. *Nature*, pages 1–10. 8
- Sinha, A., Namkoong, H., and Duchi, J. (2018). Certifying some distributional robustness with principled adversarial training. In *International Conference on Learning Representations*. 1
- Sturm, K.-T. (2006). On the geometry of metric measure spaces. *Acta Mathematica*, 196(1):65–131. 3
- Tong, A., Huang, J., Wolf, G., Van Dijk, D., and Krishnaswamy, S. (2020). TrajectoryNet: A dynamic optimal transport network for modeling cellular dynamics. In *International conference on machine learning*, pages 9526–9536. PMLR. 1
- Tong, A., Malkin, N., Hugué, G., Zhang, Y., Rector-Brooks, J., FATRAS, K., Wolf, G., and Bengio, Y. (2023). Improving and generalizing flow-based generative models with minibatch optimal transport. In *ICML Workshop on New Frontiers in Learning, Control, and Dynamical Systems*. 1
- Vaswani, S., Mishkin, A., Laradji, I., Schmidt, M., Gidel, G., and Lacoste-Julien, S. (2019). Painless stochastic gradient: Interpolation, line-search, and convergence rates. *Advances in neural information processing systems*, 32:3732–3745. 20
- Vayer, T. (2020). A contribution to optimal transport on incomparable spaces. *arXiv preprint arXiv:2011.04447*. 1, 3, 7, 17
- Vayer, T., Chapel, L., Flamary, R., Tavenard, R., and Courty, N. (2020). Fused gromov-wasserstein distance for structured objects. *Algorithms*, 13(9):212. 4
- Villani, C. (2003). Topics in optimal transportation. *Graduate Studies in Mathematics*. 2
- Weed, J. and Bach, F. (2019). Sharp asymptotic and finite-sample rates of convergence of empirical measures in wasserstein distance. *Bernoulli*, 25(4A). 2
- Welch, J. D., Hartemink, A. J., and Prins, J. F. (2017). Matcher: manifold alignment reveals correspondence between single cell transcriptome and epigenome dynamics. *Genome biology*, 18(1):1–19. 8
- Wong, E., Schmidt, F., and Kolter, Z. (2019). Wasserstein adversarial examples via projected sinkhorn iterations. In *International Conference on Machine Learning*, pages 6808–6817. PMLR. 1

Zhang, M., Liu, Y., Luan, H., and Sun, M. (2017). Earth mover’s distance minimization for unsupervised bilingual lexicon induction. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 1934–1945. [3](#), [13](#)

Zhang, Z., Goldfeld, Z., Mroueh, Y., and Sriperumbudur, B. K. (2022). Gromov-wasserstein distances: Entropic regularization, duality, and sample complexity. *arXiv preprint arXiv:2212.12848*. [1](#), [7](#), [17](#)

## Checklist

1. For all models and algorithms presented, check if you include:
  - (a) A clear description of the mathematical setting, assumptions, algorithm, and/or model. Yes
  - (b) An analysis of the properties and complexity (time, space, sample size) of any algorithm. Yes
  - (c) (Optional) Anonymized source code, with specification of all dependencies, including external libraries. Not Applicable. We will release it by the supplementary materials deadline.
2. For any theoretical claim, check if you include:
  - (a) Statements of the full set of assumptions of all theoretical results. Yes
  - (b) Complete proofs of all theoretical results. Yes
  - (c) Clear explanations of any assumptions. Yes
3. For all figures and tables that present empirical results, check if you include:
  - (a) The code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL). Yes
  - (b) All the training details (e.g., data splits, hyperparameters, how they were chosen). Yes
  - (c) A clear definition of the specific measure or statistics and error bars (e.g., with respect to the random seed after running experiments multiple times). Yes
  - (d) A description of the computing infrastructure used. (e.g., type of GPUs, internal cluster, or cloud provider). Not Applicable
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets, check if you include:
  - (a) Citations of the creator If your work uses existing assets. Yes

- (b) The license information of the assets, if applicable. Not Applicable
- (c) New assets either in the supplemental material or as a URL, if applicable. Not Applicable
- (d) Information about consent from data providers/curators. Not Applicable
- (e) Discussion of sensible content if applicable, e.g., personally identifiable information or offensive content. Not Applicable

5. If you used crowdsourcing or conducted research with human subjects, check if you include:
  - (a) The full text of instructions given to participants and screenshots. Not Applicable
  - (b) Descriptions of potential participant risks, with links to Institutional Review Board (IRB) approvals if applicable. Not Applicable
  - (c) The estimated hourly wage paid to participants and the total amount spent on participant compensation. Not Applicable

## SUPPLEMENTARY MATERIAL

### Structured Transforms Across Spaces with Cost-Regularized Optimal Transport

---

#### A APPENDIX FOR SECTION 2

##### A.1 Proof of Proposition 2.4

Recall that the Wasserstein Procrustes problem (1) is

$$\min_{\pi \in \Pi(\alpha, \beta)} \min_{\mathbf{C} \in \mathcal{O}_d} \int_{\mathcal{X} \times \mathcal{Y}} \|\mathbf{C}x - y\|^2 d\pi(x, y). \quad (12)$$

As noticed in e.g. Eq. 11 in Zhang et al. (2017), this formulation is an extension of the orthogonal Procrustes problem. Given two matrices  $\mathbf{A}, \mathbf{B} \in \mathbb{R}^{d \times d}$  The orthogonal Procrustes problem aims to solve

$$\min_{\mathbf{C} \in \mathcal{O}_d} \|\mathbf{C}\mathbf{A} - \mathbf{B}\|_F.$$

The solution of this problem is  $\mathbf{C}^* = \mathbf{U}\mathbf{V}^\top$ , where  $\mathbf{B}\mathbf{A}^\top = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^\top$ . One can follow a similar procedure to show that a solution to the inner problem in (12) is given by  $\mathbf{C}^\top = \mathbf{U}(\pi)\mathbf{V}(\pi)^\top$ , where  $\int yx^\top d\pi(x, y) = \mathbf{U}(\pi)\mathbf{\Sigma}(\pi)\mathbf{V}(\pi)^\top$  is an SVD. Replacing with this solution in the inner problem gives the desired QOT formulation:

$$\min_{\pi \in \Pi(\alpha, \beta)} \int_{\mathcal{X} \times \mathcal{Y}} \|\mathbf{U}(\pi)\mathbf{V}(\pi)^\top x - y\|^2 d\pi(x, y).$$

#### B APPENDIX FOR SECTION 3

##### B.1 Proof of Proposition 3.1

Starting from  $\text{L-ROT}_\varepsilon$ , we have

$$\begin{aligned} & \min_{\substack{\pi \in \Pi(\alpha, \beta) \\ \mathbf{M} \in \mathbb{R}^{d_y \times d_x}}} \int -\langle \mathbf{M}x, y \rangle d\pi + \mathcal{R}(\mathbf{M}) + \varepsilon \text{KL}(\pi || \alpha \otimes \beta), \\ &= \min_{\pi \in \Pi(\alpha, \beta)} \varepsilon \text{KL}(\pi || \alpha \otimes \beta) + \min_{\mathbf{M} \in \mathbb{R}^{d_y \times d_x}} \int -\langle \mathbf{M}x, y \rangle d\pi + \mathcal{R}(\mathbf{M}) \\ &= \min_{\pi \in \Pi(\alpha, \beta)} \varepsilon \text{KL}(\pi || \alpha \otimes \beta) + \min_{\mathbf{M} \in \mathbb{R}^{d_y \times d_x}} -\langle \mathbf{M}, \int yx^\top d\pi \rangle_F + \mathcal{R}(\mathbf{M}) \\ &= \min_{\pi \in \Pi(\alpha, \beta)} \varepsilon \text{KL}(\pi || \alpha \otimes \beta) - \max_{\mathbf{M} \in \mathbb{R}^{d_y \times d_x}} -\langle \mathbf{M}, \int yx^\top d\pi \rangle_F + \mathcal{R}(\mathbf{M}) \\ &= \min_{\pi \in \Pi(\alpha, \beta)} \varepsilon \text{KL}(\pi || \alpha \otimes \beta) - \mathcal{R}^* \left( \int_{\mathcal{X} \times \mathcal{Y}} yx^\top d\pi(x, y) \right), \end{aligned}$$

which is the desired result.

## B.2 Proof of Proposition 3.2

Consider  $L\text{-ROT}_\varepsilon$ :

$$\min_{\substack{\pi \in \Pi(\alpha, \beta) \\ \mathbf{M} \in \mathbb{R}^{d_y \times d_x}}} \int -\langle \mathbf{M}x, y \rangle d\pi + \mathcal{R}(\mathbf{M}) + \varepsilon \text{KL}(\pi | | \alpha \otimes \beta).$$

Under the assumption of Proposition 3.2, the minimizer over  $\pi$  is unique for any fixed  $\mathbf{M}$ , and the minimizer over  $\mathbf{M}$  is unique for any fixed  $\pi$ . Thus, we can directly apply (Bertsekas, 1997, Prop. 2.7.1 and Ex. 2.7.1) to conclude.

## B.3 Proof of Lemma 3.3

For any  $\pi \in \Pi(\alpha, \beta)$ ,

$$\begin{aligned} \mathbf{M}(\pi) &\triangleq \underset{\mathbf{M}}{\text{argmin}} - \int_{\mathcal{X} \times \mathcal{Y}} \langle \mathbf{M}x, y \rangle d\pi + \frac{\|\mathbf{M}\|_F^2}{2} + \lambda g(\mathbf{M}) \\ &= \underset{\mathbf{M}}{\text{argmin}} - \langle \mathbf{M}, \int_{\mathcal{X} \times \mathcal{Y}} yx^\top d\pi \rangle_F + \frac{\|\mathbf{M}\|_F^2}{2} + \lambda g(\mathbf{M}) \\ &= \underset{\mathbf{M}}{\text{argmin}} \frac{1}{2} \left\| \mathbf{M} - \int_{\mathcal{X} \times \mathcal{Y}} yx^\top d\pi \right\|_F^2 + \lambda g(\mathbf{M}) \\ &= \text{prox}_{\lambda g} \left( \int_{\mathcal{X} \times \mathcal{Y}} yx^\top d\pi \right), \end{aligned}$$

where in the third line we introduced an additional term  $\frac{1}{2} \left\| \int_{\mathcal{X} \times \mathcal{Y}} yx^\top d\pi \right\|_F^2$  which doesn't depend on  $M$ .

## B.4 Proof of Proposition 3.4

Recall the cost-regularized OT problem with  $\varepsilon = 0$

$$\underset{\substack{\pi \in \Pi(\alpha, \beta) \\ \mathbf{M} \in \mathbb{R}^{d_y \times d_x}}} {\text{argmin}} - \int \langle \mathbf{M}x, y \rangle d\pi + \frac{\|\mathbf{M}\|_F^2}{2} + \lambda g(\mathbf{M}). \quad (13)$$

Denote by  $h : \Pi(\alpha, \beta) \times \mathbb{R}^{d_y \times d_x} \rightarrow \mathbb{R}$  and  $H : \Pi(\alpha, \beta) \rightarrow \mathbb{R}$ :

$$h(\pi, \mathbf{M}) = - \int_{\mathcal{X} \times \mathcal{Y}} \langle \mathbf{M}x, y \rangle d\pi + \frac{\|\mathbf{M}\|_F^2}{2} + \lambda g(\mathbf{M}) \quad \text{and} \quad H(\pi) = \min_{\mathbf{M} \in \mathbb{R}^{d_y \times d_x}} h(\pi, \mathbf{M}).$$

As shown in Lemma 3.3, given  $\pi$ , the solution of 13 is

$$\mathbf{M}(\pi) \triangleq \text{prox}_{\lambda g} \left( \int_{\mathcal{X} \times \mathcal{Y}} yx^\top d\pi \right),$$

Proposition 3.4 follows by evaluating the objective (13) in the solution  $\mathbf{M}(\pi)$  using the corresponding proximal operator for each regularization.

**$\ell_1$ -regularization.** Consider the case where  $g(\mathbf{M}) = \|\mathbf{M}\|_1$ . Recall that for  $\mathbf{M} \in \mathbb{R}^{d_y \times d_x}$  and all  $i, j \in [d_x] \times [d_y]$ ,

$$\text{prox}_{\lambda \|\cdot\|_1}(\mathbf{M})_{ij} = \text{sign}(\mathbf{M}_{ij}) (|\mathbf{M}_{ij}| - \lambda)_+.$$

Thus, for  $\mathbf{M}(\pi)$ ,

$$\text{prox}_{\lambda \|\cdot\|_1}(\mathbf{M}(\pi))_{ij} = \text{sign} \left( \int_{\mathcal{X} \times \mathcal{Y}} x_i y_j d\pi \right) \left( \int_{\mathcal{X} \times \mathcal{Y}} x_i y_j d\pi - \lambda \right)_+. \quad (14)$$

Define  $I_\pi \triangleq \{(i, j) \in [d_x] \times [d_y] : \int_{\mathcal{X} \times \mathcal{Y}} x_i y_j d\pi(x, y) > \lambda\}$ . Plugging (16) in (13) gives, after reducing,

$$\begin{aligned} & \operatorname{argmin}_{\pi \in \Pi(\alpha, \beta)} \sum_{(i, j) \in I_\pi} \int_{\mathcal{X} \times \mathcal{Y}} \int_{\mathcal{X} \times \mathcal{Y}} -x_i y_j x'_i y'_j d\pi(x, y) d\pi(x', y') + \lambda \sum_{(i, j) \in I_\pi} \left| \int_{\mathcal{X} \times \mathcal{Y}} x_i y_j d\pi \right| - \frac{\lambda^2}{2} |I_\pi| \\ &= \int_{(\mathcal{X} \times \mathcal{Y})^2} \sum_{(i, j) \in I_\pi} (x_i x'_i - y_j y'_j)^2 d\pi d\pi - \sum_{(i, j) \in I_\pi} \left( \int_{\mathcal{X}} x_i^2 d\alpha + \int_{\mathcal{Y}} y_j^2 d\beta - \lambda \left| \int_{\mathcal{X} \times \mathcal{Y}} y_j x_i d\pi(x, y) \right| \right) - \frac{\lambda^2}{2} |I_\pi|. \end{aligned} \quad (15)$$

Let  $(\pi^*, \mathbf{M}(\pi^*))$  be a solution to (8). Then

$$\pi \in \operatorname{argmin}_{\pi \in \Pi(\alpha, \beta)} H(\pi) \iff \pi \in \operatorname{argmin}_{\pi \in \Pi(\alpha, \beta)} h(\pi, \mathbf{M}(\pi^*)).$$

Hence, using the same calculations as in (15), (13) is equivalent to

$$\min_{\pi \in \Pi(\alpha, \beta)} \int_{(\mathcal{X} \times \mathcal{Y})^2} \sum_{(i, j) \in I_{\pi^*}} (x_i x'_i - y_j y'_j)^2 d\pi^* d\pi - \sum_{(i, j) \in I_{\pi^*}} \left( \int_{\mathcal{X}} x_i^2 d\alpha + \int_{\mathcal{Y}} y_j^2 d\beta - \lambda \left| \int_{\mathcal{X} \times \mathcal{Y}} y_j x_i d\pi^*(x, y) \right| \right).$$

Thus, ignoring the terms that don't depend on  $\pi$ , (13) is equivalent to

$$\min_{\pi \in \Pi(\alpha, \beta)} \int_{(\mathcal{X} \times \mathcal{Y})^2} \sum_{(i, j) \in I_{\pi^*}} (x_i x'_i - y_j y'_j)^2 d\pi^* d\pi,$$

which is the desired result.

**$\ell_{1,2}$ -regularization.** We proceed similarly. Recall that for  $\mathbf{M} \in \mathbb{R}^{d_y \times d_x}$  and all  $i \in [d_x]$ ,

$$\operatorname{prox}_{\lambda \|\cdot\|_{1,2}}(\mathbf{M})_{:i} = \left( 1 - \frac{\lambda}{\max\{\|\mathbf{M}_{:i}\|, \lambda\}} \right) \mathbf{M}_{:i}, \quad (16)$$

i.e.

$$\mathbf{M}_{:i} = \begin{cases} \left( 1 - \frac{\lambda}{\|\mathbf{M}_{:i}\|} \right) \mathbf{M}_{:i} & \text{if } \|\mathbf{M}_{:i}\| > \lambda \\ 0 & \text{otherwise.} \end{cases}$$

Let  $\mathbf{A} \in \mathbb{R}^{d_y \times d_x}$  and  $\tilde{\mathbf{A}} = \operatorname{prox}_{\lambda \|\cdot\|_{1,2}}(\mathbf{A})$ . Let  $I \triangleq \{\|\mathbf{A}_{:i}\| > \lambda\}$ . Then,

$$\begin{aligned} & -\langle \mathbf{A}, \tilde{\mathbf{A}} \rangle_F + \frac{\|\tilde{\mathbf{A}}\|}{2} + \lambda \|\tilde{\mathbf{A}}\|_{1,2} \\ &= \lambda \sum_{i \in I} \|\mathbf{A}_{:i}\| - \lambda^2 + \sum_{i \in I} \sum_j \left( -\mathbf{A}_{ij}^2 \left( 1 - \frac{\lambda}{\|\mathbf{A}_{:i}\|} \right) + \frac{1}{2} \left( 1 - \frac{\lambda}{\|\mathbf{A}_{:i}\|} \right)^2 \mathbf{A}_{ij}^2 \right) \\ &= \lambda \sum_{i \in I} \|\mathbf{A}_{:i}\| - \lambda^2 |I| - \frac{1}{2} \sum_{i \in I} \sum_j (\|\mathbf{A}_{:i}\| - \lambda^2) \frac{\mathbf{A}_{ij}^2}{\|\mathbf{A}_{:i}\|} \\ &= \lambda \sum_{i \in I} \|\mathbf{A}_{:i}\| - \lambda^2 |I| - \frac{1}{2} \sum_{i \in I} (\|\mathbf{A}_{:i}\|^2 - \lambda^2) \\ &= -\frac{1}{2} \sum_{i \in I} \|\mathbf{A}_{:i}\|^2 + \lambda \sum_{i \in I} \|\mathbf{A}_{:i}\| - \frac{\lambda^2}{2} |I| \end{aligned}$$

Thus, using  $\mathbf{A} = \mathbf{M}(\pi)$  and  $I_\pi \triangleq \{i \in [d_x] : \|\int_{\mathcal{Y} \times \mathcal{Y}} x_i y d\pi\| > \lambda\}$ , (13) is equivalent to

$$\min_{\pi \in \Pi(\alpha, \beta)} \int_{(\mathcal{X} \times \mathcal{Y})^2} (\langle x_{I_\pi}, x'_{I_\pi} \rangle - \langle y, y' \rangle)^2 d\pi d\pi - \int_{\mathcal{X}} \langle x_{I_\pi}, x'_{I_\pi} \rangle^2 d\alpha + \sum_{i \in I_\pi} \left\| \int_{\mathcal{X} \times \mathcal{Y}} x_i y d\pi(x, y) \right\| - \frac{\lambda^2}{2} |I_\pi|.$$

Proceeding as we did for  $\ell_1$  regularization, let  $(\pi^*, \mathbf{M}(\pi^*))$  be a solution to (8). Then (13) is equivalent to

$$\min_{\pi \in \Pi(\alpha, \beta)} \int_{(\mathcal{X} \times \mathcal{Y})^2} \int_{\mathcal{X} \times \mathcal{Y}} (\langle x_{I_{\pi^*}}, x'_{I_{\pi^*}} \rangle - \langle y, y' \rangle)^2 d\pi^* d\pi - \int_{\mathcal{X}} \langle x_{I_{\pi^*}}, x'_{I_{\pi^*}} \rangle^2 d\alpha + \lambda \sum_{i \in I_{\pi^*}} \left\| \int_{\mathcal{X} \times \mathcal{Y}} x_i y d\pi^* \right\| - \frac{\lambda^2}{2} |I_{\pi^*}|$$

Thus, ignoring the terms that don't depend on  $\pi$ , (13) is equivalent to

$$\min_{\pi \in \Pi(\alpha, \beta)} \int_{\mathcal{X} \times \mathcal{Y}} (\langle x_{I_{\pi^*}}, x'_{I_{\pi^*}} \rangle - \langle y, y' \rangle)^2 d\pi^* d\pi,$$

which is the desired result.

## C APPENDIX FOR SECTION 4

### C.1 Proof for 4.1

Before moving to the proof of Proposition 4.1, we first clarify a point concerning the existence and convergence of entropic maps for the inner-product linear OT.

The results of Pooladian and Niles-Weed (2021), which define an entropic map and prove its convergence to the Monge map, are dedicated to the linear OT problem with the squared euclidean cost with  $\mathcal{X} = \mathcal{Y} \subset \mathbb{R}^d$ :

$$\operatorname{argmin}_{\pi \in \Pi(\alpha, \beta)} \int_{\mathcal{X} \times \mathcal{Y}} \frac{1}{2} \|x - y\|^2 d\pi(x, y).$$

In our case, an intermediary result that we will need is the convergence of the entropic map for the inner-product cost  $c(x, y) = -\langle x, y \rangle$ . However, by developing the square, one can see that the two problems are equivalent:

$$\begin{aligned} & \operatorname{argmin}_{\pi \in \Pi(\alpha, \beta)} \int_{\mathcal{X} \times \mathcal{Y}} \frac{1}{2} \|x - y\|^2 d\pi(x, y) \\ &= \operatorname{argmin}_{\pi \in \Pi(\alpha, \beta)} \int_{\mathcal{X} \times \mathcal{Y}} \frac{1}{2} \|x\|^2 d\pi(x, y) - \int_{\mathcal{X} \times \mathcal{Y}} \langle x, y \rangle d\pi(x, y) + \frac{1}{2} \|y\|^2 d\pi(x, y) \\ &= \operatorname{argmin}_{\pi \in \Pi(\alpha, \beta)} \int_{\mathcal{X}} \frac{1}{2} \|x\|^2 d\alpha(x) - \int_{\mathcal{X} \times \mathcal{Y}} \langle x, y \rangle d\pi(x, y) + \frac{1}{2} \int_{\mathcal{Y}} \|y\|^2 d\beta(y) \\ &= \operatorname{argmin}_{\pi \in \Pi(\alpha, \beta)} - \int_{\mathcal{X} \times \mathcal{Y}} \langle x, y \rangle d\pi(x, y). \end{aligned}$$

Thus, squared euclidean linear OT and inner-product OT have the same OT plan, i.e. (since the OT plans for the squared euclidean cost are induced by Monge maps) the same Monge maps. Hence, the results of Pooladian and Niles-Weed (2021) directly extend to the inner-product cost. An additional, easy to verify fact, is the form of the Monge maps for the inner-product cost.

**Lemma C.1.** *Let  $f$  be a Kantorovitch potential for the inner product cost between two measures  $\mu$  and  $\nu$  with compact supports in  $\mathbb{R}^d$ . Then the Monge map for the inner-product cost (or equivalently for the squared euclidean cost) can be written as*

$$T(x) = -\nabla f(x).$$

*Proof.* One can show that if  $f$  is a Kantorovitch potential for the inner product cost, then  $\tilde{f} : x \mapsto f(x) + \frac{\|x\|^2}{2}$  is a Kantorovitch potential for the squared Euclidean cost. And we know from Brenier (1991) that given any Kantorovitch potential  $\tilde{f}$  for the squared euclidean cost, the unique (since Kantorovitch potentials are equal up to a constant) Monge map can be written as  $T(x) = x - \nabla \tilde{f}(x)$ . Thus, given a Kantorovitch potential  $f$  for the IP cost, we can write the Monge map for the inner product cost (or for squared euclidean costs since they have the same Monge map) as  $T(x) = x - \nabla(f(x) + \frac{\|x\|^2}{2}) = -\nabla f(x)$ .  $\square$

For completeness, we rewrite the entropic map of Definition 4.2.

$$T_{\varepsilon, \varepsilon'}(x) = \frac{\int y \exp((g_{\varepsilon, \varepsilon'}(y) + \langle \mathbf{M}_{\varepsilon}^* x, y \rangle) / \varepsilon) d\beta(y)}{\int \exp((g_{\varepsilon, \varepsilon'}(y) + \langle \mathbf{M}_{\varepsilon}^* x, y \rangle) / \varepsilon) d\beta(y)}.$$



**Proposition C.2.** Let  $(f_{\varepsilon, \varepsilon'}, g_{\varepsilon, \varepsilon'})$  be Sinkhorn potentials for the inner product cost between  $\mathbf{M}_{\varepsilon\sharp}^* \alpha$  and  $\beta$  with an  $\varepsilon'$  entropic regularization. Then,

$$T_{\varepsilon, \varepsilon'} = -\nabla f_{\varepsilon, \varepsilon'} \circ \mathbf{M}_{\varepsilon}^*$$

*Proof.* Using optimality conditions for the  $\varepsilon'$ -entropy regularized OT problem, we have that for all  $\tilde{y} \in \mathcal{Y}$ ,

$$f_{\varepsilon, \varepsilon'}(\tilde{y}) = -\varepsilon \log \left( \int \exp((g_{\varepsilon, \varepsilon'}(y) + \langle \tilde{y}, y \rangle) / \varepsilon) d\beta(y) \right)$$

Differentiating w.r.t.  $\tilde{y}$  gives

$$\nabla f_{\varepsilon, \varepsilon'}(\tilde{y}) = -\frac{\int y \exp((g_{\varepsilon, \varepsilon'}(y) + \langle \tilde{y}, y \rangle) / \varepsilon) d\beta(y)}{\int \exp((g_{\varepsilon, \varepsilon'}(y) + \langle \tilde{y}, y \rangle) / \varepsilon) d\beta(y)}.$$

Hence for all  $x \in \mathcal{X}$ , evaluating at  $\tilde{y} = \mathbf{M}_{\varepsilon}^* x$

$$\nabla f_{\varepsilon, \varepsilon'}(\mathbf{M}_{\varepsilon}^* x) = -\frac{\int y \exp((g_{\varepsilon, \varepsilon'}(y) + \langle \mathbf{M}_{\varepsilon}^* x, y \rangle) / \varepsilon) d\beta(y)}{\int \exp((g_{\varepsilon, \varepsilon'}(y) + \langle \mathbf{M}_{\varepsilon}^* x, y \rangle) / \varepsilon) d\beta(y)} = -T_{\varepsilon, \varepsilon'}(x).$$

That is,  $T_{\varepsilon, \varepsilon'} = -\nabla f_{\varepsilon, \varepsilon'} \circ \mathbf{M}_{\varepsilon}^*$ .  $\square$

Finally, before moving on to the proof, recall the following result from [Vayer \(2020\)](#), also appearing in [\(Dumont et al., 2022, Proposition 6\)](#).

**Proposition C.3** (Theorem 4.2.3 in [Vayer \(2020\)](#), Informal.). *Finding a Monge map for the inner-product GW problem reduces to finding a Monge map between  $\alpha$  and  $\beta$  for the cost  $c_{\mathbf{M}}(x, y) = -\langle \mathbf{M}^* x, y \rangle$ , where  $\mathbf{M}^* = \int y x^\top d\pi^*$ , and  $\pi^*$  is a solution to the inner-product GW problem. Such a map is given by*

$$T = -\nabla f \circ \mathbf{M}^*,$$

where  $f$  is a Kantorovitch potential for  $\mathcal{W}_{0, \langle \cdot, \cdot \rangle}(\mathbf{M}_{\sharp} \alpha, \beta)$ .

*Proof of Proposition 4.4.* Assume that  $d_x \geq d_y$  and consider the case where  $rk(\mathbf{M}^*) = d_y$ .

Using Lemma C.1, we now that the Monge map for  $\mathcal{W}_{\langle \cdot, \cdot \rangle}^0(\mathbf{M}_{\varepsilon\sharp}^* \alpha, \beta)$  is  $T_{\varepsilon, 0}(x) = -\nabla f_{\varepsilon, 0}(x)$ , where  $f_{\varepsilon, 0}$  is a Kantorovitch potential for  $\mathcal{W}_{\langle \cdot, \cdot \rangle}^0(\mathbf{M}_{\varepsilon\sharp}^* \alpha, \beta)$ . Since  $f_{\varepsilon, \varepsilon'}$  is a Sinkhorn potential for  $\mathcal{W}_{\langle \cdot, \cdot \rangle}^{\varepsilon'}(\mathbf{M}_{\varepsilon\sharp}^* \alpha, \beta)$ , we can apply [\(Pooladian and Niles-Weed, 2021, Corollary 1\)](#) (here with the inner product cost rather than the squared euclidean cost), with  $P = \mathbf{M}_{\varepsilon\sharp}^* \alpha$  and  $Q = \beta$ . We have that

$$\int_{\mathcal{Y}} \|\nabla f_{\varepsilon, \varepsilon'}(y) - \nabla f_{\varepsilon, 0}(y)\|^2 d\mathbf{M}_{\varepsilon\sharp}^* \alpha(y) \leq \varepsilon'^2 I_0(\mathbf{M}_{\varepsilon\sharp}^* \alpha, \beta) + \varepsilon'^{(\bar{\alpha}+1)/2},$$

Hence,

$$\int_{\mathcal{X}} \|\nabla f_{\varepsilon, \varepsilon'}(\mathbf{M}_{\varepsilon} x) - \nabla f_{\varepsilon, 0}(\mathbf{M}_{\varepsilon} x)\|^2 d\alpha(x) \leq \varepsilon'^2 I_0(\mathbf{M}_{\varepsilon\sharp}^* \alpha, \beta) + \varepsilon'^{(\bar{\alpha}+1)/2},$$

that is, using Lemma C.1 and Proposition C.2,

$$\int_{\mathcal{X}} \|T_{\varepsilon, \varepsilon'}(x) - T_{\varepsilon, 0}(x)\|^2 d\alpha(x) \leq \varepsilon'^2 I_0(\mathbf{M}_{\varepsilon\sharp}^* \alpha, \beta) + \varepsilon'^{(\bar{\alpha}+1)/2}.$$

And since  $I_0(\alpha'_\varepsilon, \beta'_\varepsilon) < \infty$  and doesn't depend on  $\varepsilon'$  [Chizat et al. \(2020\)](#), we have that for any  $\varepsilon > 0$ ,  $T_{\varepsilon, \varepsilon'} \xrightarrow{L^2(\alpha)} T_{\varepsilon, 0}$  as  $\varepsilon' \rightarrow 0$ .

Now for the second statement. [Zhang et al. \(2022\)](#) showed that  $\pi_\varepsilon^* \rightarrow_\varepsilon \pi_0^*$  along a subsequence. Hence,  $\mathbf{M}_\varepsilon := \int y x^\top d\pi_\varepsilon^* \rightarrow_\varepsilon \mathbf{M} := \int y x^\top d\pi^*(x, y)$ , which implies that  $\mathbf{M}_{\varepsilon\sharp}^* \alpha \rightarrow \mathbf{M}_{\sharp}^* \alpha$  along a subsequence.

Since  $\nabla f_{\varepsilon, 0}$  is a Monge map between  $\mathbf{M}_{\varepsilon\sharp}^* \alpha$  and  $\beta$ , and  $\nabla f_{0, 0}$  is a Monge map between  $\mathbf{M}_{\sharp}^* \alpha$  and  $\beta$ , it follows from the proof of [\(De Philippis, 2013, Theorem 4.2\)](#) that  $T_{\varepsilon, 0} \xrightarrow{\frac{L_2(\alpha)}{\varepsilon}} T_{0, 0}$ . Indeed, it is shown that this implies that  $\nabla f_{\varepsilon, 0}$  converges locally uniformly to  $\nabla f_{0, 0}$ , which in turn implies that  $\nabla f_{\varepsilon, 0}$  converges uniformly to  $\nabla f_{0, 0}$  since  $\mathcal{X}$  is compact. Thus, since  $\mathbf{M}_\varepsilon \rightarrow \mathbf{M}$  and  $\mathcal{X}$  is compact, we have  $T_{\varepsilon, 0} \xrightarrow{u} T_{0, 0}$ , and consequently  $T_{\varepsilon, 0} \xrightarrow{\frac{L_2(\alpha)}{\varepsilon}} T_{0, 0}$  along a subsequence, which gives the desired result.  $\square$

## C.2 Using rank constraints in extremely high-dimensional problems

As we suggested at the end of Section 3, using the proximal operators of the nuclear norm or of the rank can be infeasible in extremely high dimension because they involve computing the SVD of a high-dimensional matrix in the  $\mathbf{M}$ -step of  $\text{Prox-ROT}$ . Thus, in such cases, we suggest to choose a rank  $r \ll d_y$  and use the following regularization:

$$\mathcal{R}(c) = \begin{cases} \frac{1}{2} \|\mathbf{M}_2^\top \mathbf{M}_1\|_F^2 & \text{if } \exists (\mathbf{M}_1, \mathbf{M}_2) \in \mathbb{R}^{d_x \times r} \times \mathbb{R}^{d_y \times r} : \forall (x, y), c(x, y) = \langle \mathbf{M}_1 x, \mathbf{M}_2 y \rangle \\ \infty & \text{otherwise.} \end{cases}$$

In other words, we directly parametrize the matrix  $\mathbf{M}$  in  $\text{L-ROT}_\varepsilon$  as a low-rank matrix:  $\mathbf{M} = \mathbf{M}_2^\top \mathbf{M}_1$ . In this case,  $\text{L-ROT}_\varepsilon$  rewrites:

$$\min_{\substack{\pi \in \Pi(\alpha, \beta) \\ (\mathbf{M}_1, \mathbf{M}_2) \in \mathbb{R}^{r \times d_x} \times \mathbb{R}^{r \times d_y}}} \int -\langle \mathbf{M}_1 x, \mathbf{M}_2 y \rangle d\pi + \frac{1}{2} \|\mathbf{M}_2^\top \mathbf{M}_1\|_F^2 + \varepsilon \text{KL}(\pi \| \alpha \otimes \beta). \quad (\text{L-ROT}_\varepsilon)$$

To solve this problem, we use alternated minimization (a.k.a. block coordinate descent) on each of the variables. Notably, we never have to compute the product  $\mathbf{M}_2^\top \mathbf{M}_1$ . Fixing  $\pi$  and considering the first-order conditions on  $(\mathbf{M}_1, \mathbf{M}_2)$  gives

$$\begin{aligned} (\mathbf{M}_2 \mathbf{M}_2^\top) \mathbf{M}_1 &= \mathbf{M}_2 \int y x^\top d\pi(x, y) \\ (\mathbf{M}_1 \mathbf{M}_1^\top) \mathbf{M}_2 &= \mathbf{M}_1 \int x y^\top d\pi(x, y). \end{aligned}$$

Thus, assuming that  $(\mathbf{M}_2 \mathbf{M}_2^\top)$  and  $(\mathbf{M}_1 \mathbf{M}_1^\top)$  are invertible,

$$\begin{aligned} \mathbf{M}_1 &= (\mathbf{M}_2 \mathbf{M}_2^\top)^{-1} \int (\mathbf{M}_2 y) x^\top d\pi(x, y) \\ \mathbf{M}_2 &= (\mathbf{M}_1 \mathbf{M}_1^\top)^{-1} \int (\mathbf{M}_1 x) y^\top d\pi(x, y). \end{aligned}$$

So in practice, given  $\pi_k$ , we successively update

$$\begin{aligned} \mathbf{M}_1^{k+1} &= (\mathbf{M}_2^k \mathbf{M}_2^{k\top})^\dagger \int (\mathbf{M}_2^k y) x^\top d\pi_k(x, y) \\ \mathbf{M}_2^{k+1} &= (\mathbf{M}_1^{k+1} \mathbf{M}_1^{k+1\top})^\dagger \int (\mathbf{M}_1^{k+1} x) y^\top d\pi_k(x, y), \end{aligned} \quad (17)$$

where  $(\mathbf{A})^\dagger$  denotes the Moore pseudoinverse of  $\mathbf{A}$ . We can see that here all the operations are linear in  $d_x$  and  $d_y$ . Moreover, as we can choose a small rank  $r \ll d_y$ , the  $\mathcal{O}(r^3)$  dependency of computing the pseudoinverse can be much smaller than  $\mathcal{O}(d_y)$ .

**Application of the updates (17).** We test the updates of (17) in place of the proximal rank operator we used in Section 5. These updates result in very similar performance in practice (Figure 5). We display Figure 3 here again for the convenience of the reader.




Figure 4: Copy of Figure 3.

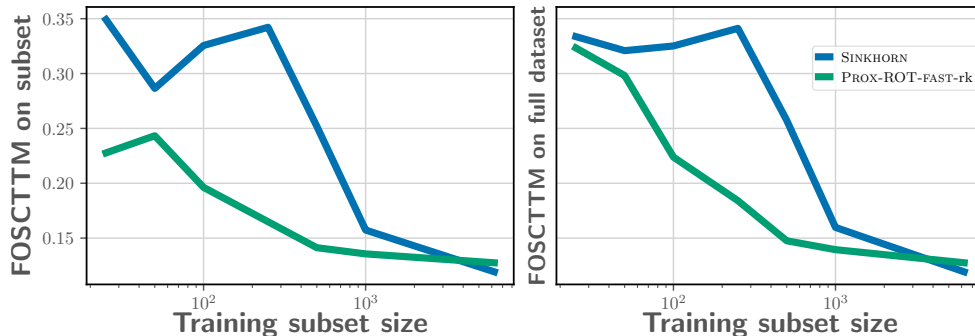


Figure 5: Reproduction of Figure 4 but with the updates (17) instead of Prox-ROT.

## D APPENDIX FOR SECTION 5

### D.1 Illustration of the action of the entropic map in the case where $M^*$ is not full-rank

One condition that we need for the convergence of the entropic map in Proposition 4.4 is that the matrix  $M^*$  (where  $(\pi^*, M^*)$  is a solution of  $L\text{-ROT}_\varepsilon$ ) is full-rank. In Section 5, we showed through an experiment (see Figure 3) that even when the matrix  $M^*$  is explicitly constrained to being low-rank, the entropic map still results in good performance.

Here, we show through a more controlled experiment that indeed the fact that  $M^*$  is not full rank doesn't result in a divergent entropic map (as measured by the performance of the map when  $\varepsilon \approx 0$ ).

To do so, we simulate two point clouds coming from two Gaussian distributions, one in 20D and the other in 10D. We then use rank regularizations and take  $\lambda = 12.5$ , which is equivalent to constraining the matrix  $M$  to being of at most of rank 5. The reason we use Gaussians is that we know ground truth OT map between them (Salmona et al., 2022).

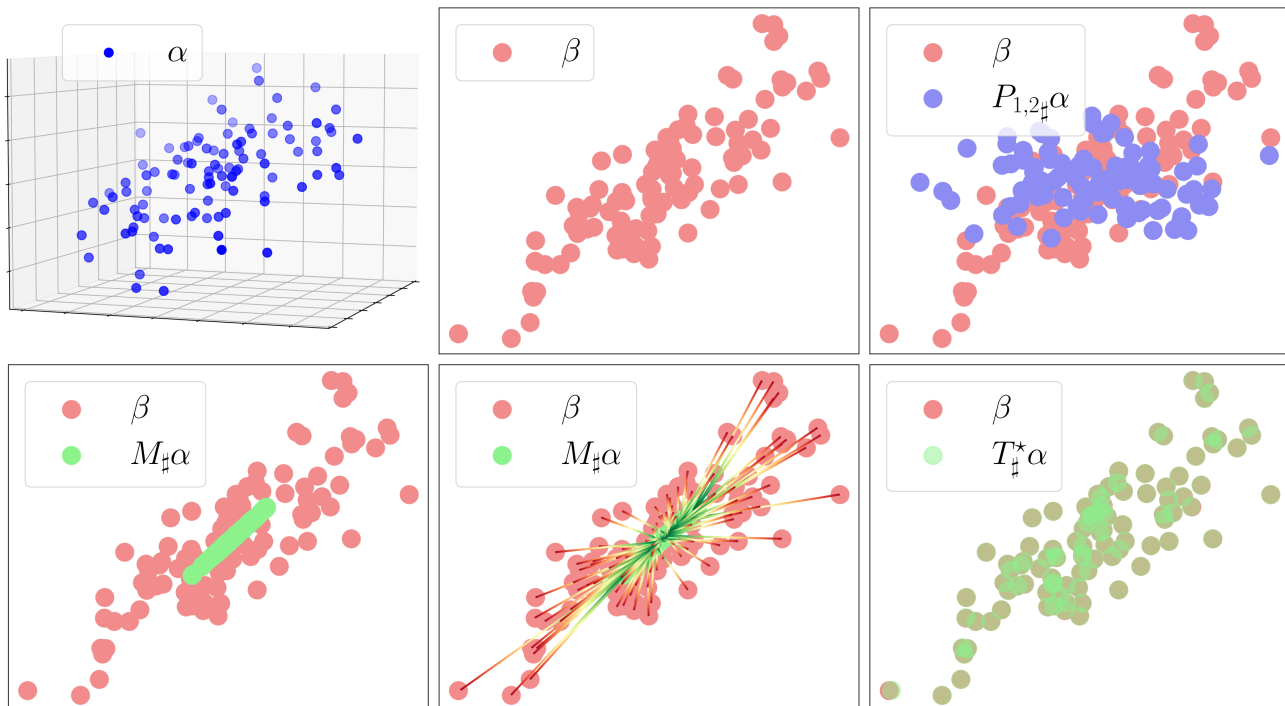


Figure 6: Illustration of the action of the entropic map when  $M_\varepsilon^*$  is low-rank.

## D.2 Details on the evaluation metrics for sparse and low-rank transforms applications

For our sparse transforms experiment on the scGM dataset (Section 5.1), we measured the performance of the algorithms using Label Transfer Accuracy, as done in Demetci et al. (2022). Label Transfer Accuracy uses a k-NN classifier fitted on the data in the anchor domain (the one where we map using the entropic map, for us the data in the lower-dimensional data methylation domain  $\{y_i\}_{i=1}^N$ ). Then the score is given as the test accuracy of the classifier on the mapped  $\{T_\varepsilon(x_i)\}_{i=1}^N$ .

For our low-rank transforms experiment on the Neurips 2021 multimodal single-cell integration dataset (Lance et al., 2022), we measured the performance of the algorithms using the Fraction of Samples Closer than the True Match (FOSCTTM) as done in the MOSCOT package (Klein et al., 2023a). The score computes the distances between each mapped sample  $\{T_\varepsilon(x_i)\}_{i=1}^N$  and the data  $\{y_i\}_{i=1}^N$  and, given ground truth correspondences between data in  $\mathcal{X}$  and data in  $\mathcal{Y}$ , computes the FOSCTTM score as the average proportion of  $y_i$ 's closer to  $T_\varepsilon(x_i)$  than its true corresponding datapoint in  $\{y_i\}_{i=1}^N$ .

## D.3 Details on the spatial transcriptomics experiment

As done in Scetbon et al. (2023), to determine the best hyperparameters for the experiment, we ran a grid search and picked the best hyperparameters combination using performance on 10 validation genes and use pearson correlation as a validation metric. The hyperparameters we considered in our grid search are the entropic regularization  $\varepsilon$  and the fused cost parameter  $\eta$  (6). We used a single minibatch size of 1000 and a single number of ascent steps in  $(f, g)$  of 1000. For the SGD step in  $\mathbf{M}$ , we use stochastic line search (Vaswani et al., 2019).