


## Efficient Approximation of Gromov-Wasserstein Distance Using Importance Sparsification

Mengyu Li, Jun Yu, Hongteng Xu & Cheng Meng


To cite this article: Mengyu Li, Jun Yu, Hongteng Xu & Cheng Meng (2023) Efficient Approximation of Gromov-Wasserstein Distance Using Importance Sparsification, Journal of Computational and Graphical Statistics, 32:4, 1512-1523, DOI: 10.1080/10618600.2023.2165500

To link to this article: <https://doi.org/10.1080/10618600.2023.2165500>

 View supplementary material [↗](#)

 Published online: 15 Feb 2023.

 Submit your article to this journal [↗](#)

 Article views: 536

 View related articles [↗](#)

 View Crossmark data [↗](#)

 Citing articles: 3 View citing articles [↗](#)



# Efficient Approximation of Gromov-Wasserstein Distance Using Importance Sparsification

Mengyu Li <sup>a</sup>, Jun Yu <sup>\*b</sup>, Hongteng Xu <sup>c</sup>, and Cheng Meng <sup>d</sup>

<sup>a</sup>Institute of Statistics and Big Data, Renmin University of China, Beijing, China; <sup>b</sup>School of Mathematics and Statistics, Beijing Institute of Technology, Beijing, China; <sup>c</sup>Gaoling School of Artificial Intelligence, Renmin University of China, Beijing, China; <sup>d</sup>Center for Applied Statistics, Institute of Statistics and Big Data, Renmin University of China, Beijing, China

## ABSTRACT

As a valid metric of metric-measure spaces, Gromov-Wasserstein (GW) distance has shown the potential for matching problems of structured data like point clouds and graphs. However, its application in practice is limited due to the high computational complexity. To overcome this challenge, we propose a novel importance sparsification method, called Spar-GW, to approximate GW distance efficiently. In particular, instead of considering a dense coupling matrix, our method leverages a simple but effective sampling strategy to construct a sparse coupling matrix and update it with few computations. The proposed Spar-GW method is applicable to the GW distance with arbitrary ground cost, and it reduces the complexity from  $O(n^4)$  to  $O(n^{2+\delta})$  for an arbitrary small  $\delta > 0$ . Theoretically, the convergence and consistency of the proposed estimation for GW distance are established under mild regularity conditions. In addition, this method can be extended to approximate the variants of GW distance, including the entropic GW distance, the fused GW distance, and the unbalanced GW distance. Experiments show the superiority of our Spar-GW to state-of-the-art methods in both synthetic and real-world tasks. Supplementary materials for this article are available online.

## ARTICLE HISTORY

Received September 2022  
Accepted December 2022

## KEYWORDS



Element-wise sampling;  
Importance sampling;  
Sinkhorn-scaling algorithm;  
Unbalanced  
Gromov-Wasserstein  
distance

## 1. Introduction


Gromov-Wasserstein (GW) distance, as an extension of classical optimal transport distance, is originally proposed to measure the distance between different metric-measure spaces (Sturm 2006; Mémoli 2011). Recently, it attracts wide attention due to its potential for tackling challenging machine learning tasks, including but not limited to shape matching (Mémoli 2011; Ezuz et al. 2017; Titouan et al. 2019b), graph analysis (Chowdhury and Mémoli 2019; Xu et al. 2019; Titouan et al. 2019a; Chowdhury and Needham 2021; Brogat-Motte et al. 2022; Vincent-Cuaz et al. 2022; Xu et al. 2023), point cloud alignment (Peyré, Cuturi, and Solomon 2016; Alvarez-Melis and Jaakkola 2018; Alaux et al. 2019; Blumberg et al. 2020), and distribution comparison across different spaces (Yan et al. 2018; Bunne et al. 2019; Chapel, Alaya, and Gasso 2020; Gong, Nie, and Xu 2022).

Despite the wide application, calculating the GW distance is NP-hard, which corresponds to solving a non-convex non-smooth optimization problem. To bypass this obstacle, many efforts have been made to approximate the GW distance with low complexity. One major strategy is applying the conditional gradient algorithm (or its variants) to solve the GW distance iteratively in an alternating optimization framework (Peyré, Cuturi, and Solomon 2016; Titouan et al. 2019a). By introducing an entropic regularizer (Solomon et al. 2016) or a proximal

term based on the Bergman divergence (Xu et al. 2019), the subproblem in each iteration will be strictly convex and can be solved via the Sinkhorn-scaling algorithm (Sinkhorn and Knopp 1967; Cuturi 2013). Another strategy is computing sliced Gromov-Wasserstein distance (Titouan et al. 2019b), which projects the samples to different 1D spaces and calculates the expectation of the GW distances defined among the projected 1D samples. More recently, to further reduce the computational complexity, more variants of the GW distance have been proposed, which achieve acceleration via imposing structural information (e.g., tree (Le, Ho, and Yamada 2021), low-rank (Xu, Luo, and Carin 2019; Sato et al. 2020; Chowdhury, Miller, and Needham 2021), and sparse structure (Xu, Luo, and Carin 2019)) on the ground cost  $\mathcal{L}$ , the coupling matrix  $\mathbf{T}$ , or both (Scetbon, Peyré, and Cuturi 2022). However, most of these methods mainly focus on the GW distance using decomposable ground cost functions (Peyré, Cuturi, and Solomon 2016). Moreover, some of them are only applicable for specific data types (e.g., point clouds in Euclidean space (Titouan et al. 2019b) and sparse graphs with clustering structures (Xu, Luo, and Carin 2019; Blumberg et al. 2020)), and they are not able to approximate the original GW distance. See Table 1 for an overall comparison. Therefore, it is urgent to develop a new approximation of the GW distance that has better efficiency and applicability.

**CONTACT** Cheng Meng  [chengmeng@ruc.edu.cn](mailto:chengmeng@ruc.edu.cn)  Center for Applied Statistics, Institute of Statistics and Big Data, Renmin University of China, Beijing, China.

\*Joint first author

 Supplementary materials for this article are available online. Please go to [www.tandfonline.com/r/JCGS](http://www.tandfonline.com/r/JCGS).

© 2023 American Statistical Association and Institute of Mathematical Statistics

**Table 1.** Comparison for various GW distance approximation methods on their time complexity, assumptions imposed on their ground cost functions and coupling matrices, and required data types.

Method	Time $O(\cdot)$	Ground cost $\mathcal{L}$	Coupling $\mathbf{T}$	Data type
Entropic GW (Peyré, Cuturi, and Solomon 2016)	$n^3$	Decomposable	-	-
Sliced GW (Titouan et al. 2019b)	$n^2$	$\ell_2$ loss	-	Points
S-GWL (Xu, Luo, and Carin 2019)	$n^2 \log(n)$	Decomposable	Sparse & low-rank	-
AE (Sato et al. 2020)	$n^2 \log(n^2)$	$p$ -power ( $p \in \mathbb{Z}_+$ )	-	-
FlowAlign (Le, Ho, and Yamada 2021)	$n^2$	$\ell_2$ loss	Tree-structure	Points
Linear-time GW (Scetbon, Peyré, and Cuturi 2022)	$r_2(r_1 + r_2)n$	$\ell_2$ loss	Low-rank	Points
SaGroW (Kerdoncuff, Emonet, and Sebban 2021)	$n^2(s' + \log(n))$	-	-	-
Spar-GW (Proposed)	$n^2 + s^2$	-	-	-

<sup>1</sup> “-” means no constraints. For the column of data type, “-” means the data can be sample points and/or their relation matrices.

<sup>2</sup> For the column of time complexity, the time of calculating relation matrices of points is also included.

<sup>3</sup>  $n$  represents the sample size. For Linear-time GW,  $r_1$  and  $r_2$  are the assumed ranks of relation matrices and coupling matrix, respectively; for SaGroW,  $s'$  denotes the number of sampled matrix; for Spar-GW,  $s$  denotes the number of sampled elements.

*Major contribution.* In this article, we propose a randomized sparsification method, called SPAR-GW, to approximate the GW distance and its variants. In particular, during the iterative optimization of the GW distance, the proposed SPAR-GW method leverages an importance sparsification mechanism to derive a sparse coupling matrix and the corresponding kernel matrix. Replacing dense multiplications with sparse ones leads to an efficient approximation of GW distance with  $O(n^2 + s^2)$  time complexity, where  $s$  is the number of selected elements from an  $n \times n$  kernel matrix. The SPAR-GW is compatible with various computational methods, including the proximal gradient algorithm for original GW distance and the Sinkhorn-scaling algorithm for entropic GW distance, and it is capable of arbitrary ground cost. In theory, we show the proposed estimator is asymptotically unbiased when  $s = O(n^{1+\delta})$  for an arbitrary small  $\delta > 0$ , under some regularity conditions. Table 1 highlights the advantage of our method. Moreover, this method can be extended to compute the variants of GW distance, for example, a straightforward application to approximate the fused GW (FGW) distance and a nontrivial extension called SPAR-UGW to approximate the unbalanced GW (UGW) distance. Experiments show the superiority of the proposed methods to state-of-the-art competitors in both synthetic and real-world tasks.

The remainder of this article is organized as follows. We start in Section 2 by introducing computational optimal transport and GW distance. In Section 3, we develop the sampling probabilities and provide details of our main algorithm. The theoretical properties of the proposed estimator are presented in Section 4. Section 5 extends the proposed method to unbalanced problems. We examine the performance of the proposed algorithms through extensive synthetic and real-world datasets in Section 6. Technical proofs and more experimental results are provided in the supplementary material.

## 2. Background

In the following, we adopt the common convention of using uppercase boldface letters for matrices, lowercase boldface letters for vectors, and regular font for scalars. We denote nonnegative real numbers by  $\mathbb{R}_+$  and the set of integers  $\{1, \dots, n\}$  by  $[n]$ . We use  $\mathbf{1}_n$  and  $\mathbf{0}_n$  to denote the all-ones and all-zeroes vectors in  $\mathbb{R}^n$ , respectively. For a matrix  $\mathbf{A} = (A_{ij})$ , its spectral norm

(i.e., the largest singular value) and Frobenius norm are denoted as  $\|\mathbf{A}\|_2$  and  $\|\mathbf{A}\|_F$ , respectively. The condition number of  $\mathbf{A}$  is defined as  $\|\mathbf{A}\|_2/\sigma_{\min}(\mathbf{A})$ , where  $\sigma_{\min}(\cdot)$  stands for the smallest singular value. We denote by  $\exp(\mathbf{A})$  the matrix with entries  $\exp(A_{ij})$ . For two matrices  $\mathbf{A}$  and  $\mathbf{B}$  of the same dimension, we denote their Frobenius inner product by  $\langle \mathbf{A}, \mathbf{B} \rangle = \sum_{i,j} A_{ij}B_{ij}$ .

### 2.1. Computational Optimal Transport

Consider two samples  $X = \{\mathbf{x}_i\}_{i=1}^m$  and  $Y = \{\mathbf{y}_j\}_{j=1}^n$  that are generated from the distributions  $\mathbf{a} \in \Delta^{m-1}$  and  $\mathbf{b} \in \Delta^{n-1}$ , respectively, where  $\Delta^{n-1}$  represents the  $(n-1)$ -Simplex. When  $\mathbf{a}$  and  $\mathbf{b}$  lie in the same space, optimal transport (OT) and its associated Wasserstein distance (Villani 2009) are used extensively to quantify the discrepancy between these two probability distributions. The modern Kantorovich formulation (Kantorovich 1942) of OT writes

$$W(\mathbf{a}, \mathbf{b}) := \min_{\mathbf{T} \in \Pi(\mathbf{a}, \mathbf{b})} \langle \mathbf{M}, \mathbf{T} \rangle, \tag{1}$$

where  $\mathbf{M} \in \mathbb{R}^{m \times n}$  is a given distance matrix,  $\Pi(\mathbf{a}, \mathbf{b}) = \{\mathbf{T} \in \mathbb{R}_+^{m \times n} : \mathbf{T}\mathbf{1}_n = \mathbf{a}, \mathbf{T}^\top \mathbf{1}_m = \mathbf{b}\}$  is the set of admissible coupling matrices, that is, all joint probability distributions with marginals  $\mathbf{a}, \mathbf{b}$ , and the  $(i, j)$ th entry of  $\mathbf{T}$  represents the amount of probability mass shifted from  $i$  to  $j$ . The solution to (1) is called the optimal transport plan. If  $\mathbf{M}$  is a distance matrix of order  $p$ ,  $W_p(\cdot, \cdot) = W(\cdot, \cdot)^{1/p}$  is called the  $p$ -Wasserstein distance.

Despite the wide applications, the computational complexity of directly solving (1) using a linear program grows cubically as  $m$  or  $n$  increases (Brenier 1997; Benamou, Brenier, Y., and Guittet 2002). To approximate the optimal transport plan efficiently, Cuturi (2013) added an entropic regularization term on (1), which leads to a strongly convex and smooth problem

$$\min_{\mathbf{T} \in \Pi(\mathbf{a}, \mathbf{b})} \langle \mathbf{M}, \mathbf{T} \rangle + \varepsilon H(\mathbf{T}), \tag{2}$$

where  $\varepsilon > 0$  is a regularization parameter and  $H(\mathbf{T}) = \langle \mathbf{T}, \log \mathbf{T} \rangle$  is the negative Shannon entropy of  $\mathbf{T}$ . By introducing a kernel matrix  $\mathbf{K} := \exp(-\mathbf{M}/\varepsilon)$ , it is known that the solution to (2) is a projection onto  $\Pi(\mathbf{a}, \mathbf{b})$  of  $\mathbf{K}$  (Peyré and Cuturi 2019). Therefore, the problem (2) can be solved by using iterative matrix scaling (Sinkhorn and Knopp 1967), called the Sinkhorn-scaling algorithm (Cuturi 2013); see Step 5 in Algorithm 1 for details. The Sinkhorn-scaling algorithm enables researchers

to approximate the OT solution efficiently, and thus has been extensively studied in recent years (Genevay et al. 2019; Lin, Ho, and Jordan 2019; Scetbon and Cuturi 2020; Liao et al. 2022a, 2022b). There also exist projection-based methods to approximate the Wasserstein distance (Bonneel et al. 2015; Meng et al. 2019; Deshpande et al. 2019; Li et al. 2022). We refer to Nadjahi (2021) and Zhang et al. (2022) for recent reviews.

## 2.2. Computational GW Distance

When  $\mathbf{a}$  and  $\mathbf{b}$  lie in different spaces, the distance matrix  $\mathbf{M}$  is unavailable and thus optimal transport can not be used. As a replacement, the Gromov-Wasserstein distance is applicable to measuring the discrepancy between two samples located in different sample spaces by comparing their structural similarity. Similar to Wasserstein distance, the intuition of GW distance is still to minimize the transportation effort, but GW only relies on the structure in each space separately.

In particular, consider two relation matrices  $\mathbf{C}^X = (C_{ij}^X) \in \mathbb{R}^{m \times m}$  and  $\mathbf{C}^Y = (C_{ij}^Y) \in \mathbb{R}^{n \times n}$ , each of which can be the distance/kernel matrix defined on a sample (Mémoli 2011; Peyré, Cuturi, and Solomon 2016), or the adjacency matrix of a graph constructed by the sample (Xu et al. 2019; Titouan et al. 2019a). Let  $\mathcal{L} : \mathbb{R} \times \mathbb{R} \mapsto \mathbb{R}$  be the ground cost function, for example, the  $\ell_1$  loss (i.e.,  $\mathcal{L}(x_1, x_2) = |x_1 - x_2|$ ), the  $\ell_2$  loss (i.e.,  $\mathcal{L}(x_1, x_2) = (x_1 - x_2)^2$ ), and the Kullback-Leibler (KL) divergence (i.e.,  $\mathcal{L}(x_1, x_2) = x_1 \log(x_1/x_2) - x_1 + x_2$ ). The GW distance is defined as the following non-convex non-smooth optimization problem:

$$\begin{aligned} \text{GW}((\mathbf{C}^X, \mathbf{a}), (\mathbf{C}^Y, \mathbf{b})) & \quad (3) \\ & := \min_{\mathbf{T} \in \Pi(\mathbf{a}, \mathbf{b})} \sum_{i, i', j, j'} \mathcal{L}(C_{ii'}^X, C_{jj'}^Y) T_{ij} T_{i'j'} \\ & = \min_{\mathbf{T} \in \Pi(\mathbf{a}, \mathbf{b})} \langle \mathcal{L}(\mathbf{C}^X, \mathbf{C}^Y) \otimes \mathbf{T}, \mathbf{T} \rangle \\ & = \langle \mathcal{L}(\mathbf{C}^X, \mathbf{C}^Y) \otimes \mathbf{T}^*, \mathbf{T}^* \rangle, \end{aligned}$$

where  $\mathcal{L}(C_{ii'}^X, C_{jj'}^Y)$  measures similarity between pairs of points or graphs,  $T_{ij}$  is the  $(i, j)$ th entry of the coupling matrix  $\mathbf{T}$ , and then the term  $\mathcal{L}(C_{ii'}^X, C_{jj'}^Y) T_{ij} T_{i'j'}$  represents the transport cost between two pairs  $(i, i')$  and  $(j, j')$ . As shown in (3), this optimization problem can be written in a matrix format (Peyré, Cuturi, and Solomon 2016), where  $\mathcal{L}(\mathbf{C}^X, \mathbf{C}^Y)$  is a tensor, and  $\mathcal{L}(\mathbf{C}^X, \mathbf{C}^Y) \otimes \mathbf{T} := [(\sum_{i', j'} \mathcal{L}(C_{ii'}^X, C_{jj'}^Y) T_{i'j'})_{i,j}] \in \mathbb{R}^{m \times n}$  is a tensor-matrix multiplication. As before,  $\Pi(\mathbf{a}, \mathbf{b})$  is the set of admissible coupling matrices, and the solution to problem (3), denoted as  $\mathbf{T}^*$ , is called the optimal transport plan.

In general, this optimization problem can be solved in an iterative framework: at the  $r$ th iteration, the coupling matrix  $\mathbf{T}$  is updated via solving the following subproblem:

$$\mathbf{T}^{(r+1)} := \arg \min_{\mathbf{T} \in \Pi(\mathbf{a}, \mathbf{b})} \underbrace{\langle \mathcal{L}(\mathbf{C}^X, \mathbf{C}^Y) \otimes \mathbf{T}^{(r)}, \mathbf{T} \rangle}_{\mathcal{C}(\mathbf{T}^{(r)})} + \underbrace{\varepsilon \mathcal{R}(\mathbf{T})}_{\text{Optional Reg.}}. \quad (4)$$

Here,  $\mathcal{C}(\mathbf{T}^{(r)})$  is a cost matrix determined by the previous coupling matrix  $\mathbf{T}^{(r)}$ ,  $\mathcal{R}(\mathbf{T})$  is an optional regularizer of  $\mathbf{T}$ , whose significance is controlled by the weight  $\varepsilon \geq 0$ . The subproblem is essentially the (regularized) optimal transport problem (1)

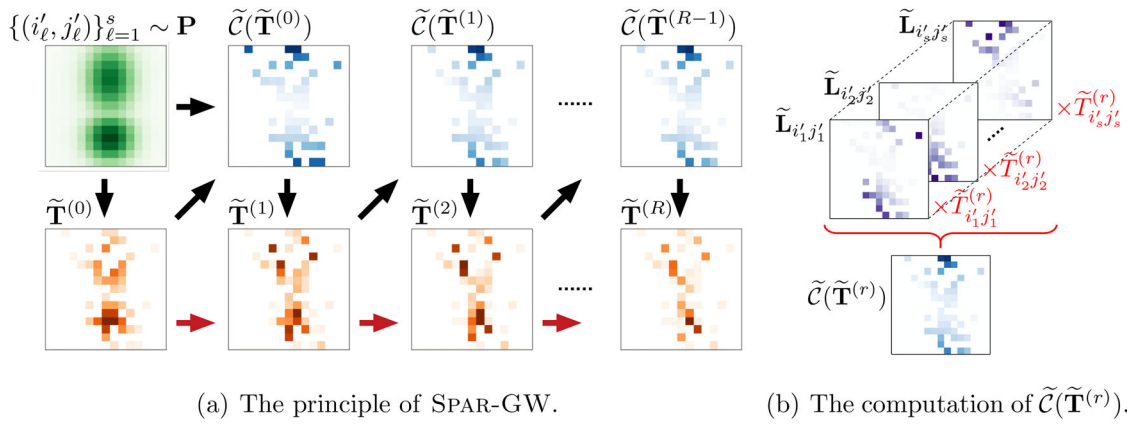
### Algorithm 1 Computation of GW distance

- 1: **Input:** Sample distributions  $\mathbf{a}, \mathbf{b}$ , relation matrices  $\mathbf{C}^X, \mathbf{C}^Y$ , ground cost function  $\mathcal{L}$ , regularization parameter  $\varepsilon$ , number of outer/inner iterations  $R, H$
- 2: Initialize  $\mathbf{T}^{(0)} = \mathbf{a}\mathbf{b}^\top$
- 3: **For**  $r = 0$  **to**  $R - 1$ :
- 4: **Construct a kernel matrix:**  $O(m^2 n^2)$ 
  - a) Compute the cost matrix  $\mathcal{C}(\mathbf{T}^{(r)}) = \mathcal{L}(\mathbf{C}^X, \mathbf{C}^Y) \otimes \mathbf{T}^{(r)}$
  - b)  $\mathbf{K}^{(r)} = \begin{cases} \exp(-\frac{\mathcal{C}(\mathbf{T}^{(r)})}{\varepsilon}) \odot \mathbf{T}^{(r)} & \text{if } \mathcal{R}(\mathbf{T}) = \text{KL}(\mathbf{T} \parallel \mathbf{T}^{(r)}) \\ \exp(-\frac{\mathcal{C}(\mathbf{T}^{(r)})}{\varepsilon}) & \text{if } \mathcal{R}(\mathbf{T}) = H(\mathbf{T}) \end{cases}$
- 5: **Sinkhorn-scaling:**  $\mathbf{T}^{(r+1)} = \text{SINKHORN}(\mathbf{a}, \mathbf{b}, \mathbf{K}^{(r)}, H)$   $O(Hmn)$ 
  - a) Initialize  $\mathbf{u}^{(0)} = \mathbf{1}_m, \mathbf{v}^{(0)} = \mathbf{1}_n$
  - b) **For**  $h = 0$  **to**  $H - 1$ :  $\mathbf{u}^{(h+1)} = \mathbf{a} \oslash (\mathbf{K}^{(r)} \mathbf{v}^{(h)})$ ,  $\mathbf{v}^{(h+1)} = \mathbf{b} \oslash (\mathbf{K}^{(r)\top} \mathbf{u}^{(h+1)})$
  - c)  $\mathbf{T}^{(r+1)} = \text{diag}(\mathbf{u}^{(H)}) \mathbf{K}^{(r)} \text{diag}(\mathbf{v}^{(H)})$
- 6: **Output:**  $\text{GW} = \langle \mathcal{C}(\mathbf{T}^{(R)}), \mathbf{T}^{(R)} \rangle$   $O(m^2 n^2)$

or (2). Without  $\mathcal{R}(\mathbf{T})$ , the problem in (4) becomes a constrained linear programming given  $\mathbf{T}^{(r)}$ , and this is often solved via the conditioned gradient followed by line-search (Titouan et al. 2019a). To improve the efficiency of solving the problem, Xu et al. (2019) implements  $\mathcal{R}(\mathbf{T})$  as a Bregman proximal term, that is, the KL-divergence  $\text{KL}(\mathbf{T} \parallel \mathbf{T}^{(r)}) = \langle \mathbf{T}, \log \mathbf{T} - \log \mathbf{T}^{(r)} \rangle$ , which leads to a proximal gradient algorithm (PGA) and improves the smoothness of  $\mathbf{T}$ 's update. When the regularizer is implemented as the entropy of  $\mathbf{T}$ , that is,  $\mathcal{R}(\mathbf{T}) := H(\mathbf{T})$ , the problem in (4) becomes an entropic optimal transport problem and solving it iteratively leads to the approximation of GW distance or entropic GW distance (Peyré, Cuturi, and Solomon 2016). Note that, when using the Bregman proximal term or the entropic regularizer, the problem in (4) can be solved via the Sinkhorn-scaling algorithm (Sinkhorn and Knopp 1967; Cuturi 2013), and accordingly, Algorithm 1 shows the computational scheme of the GW distance, where  $\odot$  and  $\oslash$  represent element-wise multiplication and division, respectively. When  $\mathcal{R}(\mathbf{T}) = H(\mathbf{T})$ , Algorithm 1 can also be used to compute the entropic GW distance by modifying the output to  $\text{GW}_\varepsilon = \langle \mathcal{C}(\mathbf{T}^{(R)}), \mathbf{T}^{(R)} \rangle + \varepsilon H(\mathbf{T}^{(R)})$ .

## 2.3. Problem Statement

The computational bottleneck of Algorithm 1 is the computation of the cost matrix  $\mathcal{C}(\mathbf{T})$ , which involves a tensor-matrix multiplication (i.e., the weighted summation of  $mn$  matrices of size  $m \times n$ ) with time complexity  $O(m^2 n^2)$ . Although the complexity can be reduced to  $O(n^2 m + m^2 n)$  when the ground cost  $\mathcal{L}$  is decomposable, that is,  $\mathcal{L}$  can be decomposed as  $\mathcal{L}(x_1, x_2) = f_1(x_1) + f_2(x_2) - h_1(x_1)h_2(x_2)$  for functions  $(f_1, f_2, h_1, h_2)$ , like the  $\ell_2$  loss or the KL-divergence (Peyré, Cuturi, and Solomon 2016). However, this setting restricts the choice of the ground cost and thus is inapplicable for GW distances in more general scenarios.



(a) The principle of SPAR-GW.

 (b) The computation of  $\tilde{\mathcal{C}}(\tilde{\mathbf{T}}^{(r)})$ .

**Figure 1.** (a) An illustration of our Spar-GW. The nonzero elements of each matrix are labeled with colors. The arrows between  $\tilde{\mathbf{T}}^{(r)}$ 's are applied when using the proximal gradient algorithm. (b) An illustration of the computation of  $\tilde{\mathcal{C}}(\tilde{\mathbf{T}}^{(r)})$  in the  $r$ th iteration.

### 3. Importance Sparsification for GW Distance

#### 3.1. Importance Sparsification

According to the analysis above, to approximate the GW distance efficiently, the key point is constructing a sparse  $\tilde{\mathcal{C}}(\mathbf{T})$  with low complexity as a surrogate for  $\mathcal{C}(\mathbf{T})$ , which motivates us to propose the SPAR-GW algorithm. Replacing the  $\mathcal{C}(\mathbf{T})$  with a sparse  $\tilde{\mathcal{C}}(\mathbf{T})$  results in two benefits. First,  $\tilde{\mathcal{C}}(\mathbf{T})$  is associated with a sparse kernel matrix  $\tilde{\mathbf{K}}$ , which enables us to use sparse matrix multiplications to accelerate the Sinkhorn-scaling algorithm, and the output is a sparse transport plan  $\tilde{\mathbf{T}}$  with the same sparsity structure as  $\tilde{\mathbf{K}}$ , that is,  $\tilde{T}_{ij} = 0$  if  $\tilde{K}_{ij} = 0$ , as shown in Figure 1(a). Second, when  $\tilde{\mathbf{T}}$  is sparse,  $\tilde{\mathcal{C}}(\tilde{\mathbf{T}})$  can be calculated by summing  $s < mn$  sparse matrices instead of  $mn$  dense ones, and each of these sparse matrices contains at most  $s$  nonzero elements, as shown in Figure 1(b). Therefore, the principle of our SPAR-GW algorithm is leveraging a simple but effective importance sparsification mechanism (i.e., constructing the sampling probability matrix  $\mathbf{P}$  in Figure 1(a)) to derive an informative sparse  $\tilde{\mathcal{C}}(\tilde{\mathbf{T}})$ , and accordingly, achieving an asymptotically unbiased estimate of the GW distance.

Recall that the GW distance in (3) can be rewritten as a summation  $\text{GW} = \sum_{i,j} T_{ij}^* C_{ij}^*$ , where  $C_{ij}^*$  is the  $(i, j)$ th element of  $\mathcal{C}(\mathbf{T}^*)$ . According to the idea of importance sampling (Liu 1996, 2008), this summation can be approximated by a weighted sum of  $s$  components, that is,  $\text{GW} \approx \sum_{(i,j) \in \mathcal{S}} T_{ij}^* C_{ij}^* / (sp_{ij})$ , where  $\mathcal{S} := \{(i, j)\}_{j=1}^s$  represents the set of  $s$  pairs of indices selected by the sampling probability  $\{p_{ij}\}_{(i,j) \in [m] \times [n]}$ . Ideally, the optimal sampling probability, which leads to the minimum estimation variance, satisfies  $p_{ij}^* \propto T_{ij}^* C_{ij}^*$ . Because neither the optimal  $T_{ij}^*$  nor  $C_{ij}^*$  is known beforehand, we propose to use a proper upper bound for  $T_{ij}^* C_{ij}^*$  as a surrogate. In particular, for the cost  $C_{ij}^*$ , we impose a mild assumption on it:  $\exists c_0 > 0$  such that  $\forall i, j, C_{ij}^* \leq c_0$ . Moreover, based on the constraint that  $\mathbf{T}^* \in \Pi(\mathbf{a}, \mathbf{b})$ , we have  $T_{ij}^* \leq a_i$  and  $T_{ij}^* \leq b_j$ , and thus  $T_{ij}^* \leq \sqrt{a_i b_j}$ . Combining these inequalities, we use the sampling probability as

$$T_{ij}^* C_{ij}^* \leq c_0 \sqrt{a_i b_j} \quad \Rightarrow \quad p_{ij} = \frac{\sqrt{a_i b_j}}{\sum_{i,j} \sqrt{a_i b_j}}, \quad 1 \leq i \leq m, \quad 1 \leq j \leq n. \quad (5)$$

Intuitively,  $T_{ij}^*$  can be large when both  $a_i$  and  $b_j$  are relatively large; otherwise,  $T_{ij}^*$  should be small if either  $a_i$  or  $b_j$  is small. Therefore, from the perspective of importance sampling, it is natural to take the geometric mean of  $a_i$  and  $b_j$  as our sampling probability.

#### 3.2. Proposed Algorithm

Let  $\mathbf{P}$  be the sampling probability matrix such that the  $(i, j)$ th element equals the  $p_{ij}$  in (5). Given  $\mathbf{P}$ , we first construct the index set  $\mathcal{S}$  by sampling  $s$  pairs of indices, and then, build  $s$  matrices  $\{\tilde{\mathbf{L}}_{i,j}\}_{(i,j) \in \mathcal{S}}$ , whose elements are

$$\tilde{L}_{i'j'} = \begin{cases} \mathcal{L}(C_{i'j'}^X, C_{i'j'}^Y) & \text{if } (i', j') \in \mathcal{S} \\ 0 & \text{otherwise} \end{cases} \quad \text{for } (i, j) \in \mathcal{S}. \quad (6)$$

As shown in Figure 1(b), in the  $r$ th iteration, we construct a sparse coupling matrix  $\tilde{\mathbf{T}}^{(r)}$ , with  $\tilde{T}_{ij}^{(r)} = 0$  if  $(i, j) \notin \mathcal{S}$ , and compute the sparse cost matrix  $\tilde{\mathcal{C}}(\tilde{\mathbf{T}}^{(r)}) = \sum_{(i,j) \in \mathcal{S}} \tilde{\mathbf{L}}_{i,j} \cdot \tilde{T}_{ij}^{(r)}$ . Accordingly, we derive the sparse kernel matrix  $\tilde{\mathbf{K}}^{(r)}$  with nonzero elements

$$\tilde{K}_{ij}^{(r)} = \exp(-\tilde{\mathcal{C}}_{ij} / \varepsilon) \tilde{T}_{ij}^{(r)} / (sp_{ij}) \quad \text{or} \quad \exp(-\tilde{\mathcal{C}}_{ij} / \varepsilon) / (sp_{ij})$$

only for  $(i, j) \in \mathcal{S}$ , where the adjustment factor  $sp_{ij}$  ensures the unbiasedness of the estimation. We then calculate the coupling matrix  $\tilde{\mathbf{T}}^{(r+1)}$  via applying the Sinkhorn-scaling algorithm to the sparse  $\tilde{\mathbf{T}}^{(r)}$  and  $\tilde{\mathbf{K}}^{(r)}$ . Algorithm 2 summarizes the SPAR-GW algorithm.

*Computational cost.* Generating the sampling probability matrix  $\mathbf{P}$  requires  $O(mn)$  time. In each iteration, calculating  $\tilde{\mathcal{C}}(\tilde{\mathbf{T}}^{(r)})$  involves the summation of  $s$  matrices, and each of them contains only  $s$  nonzero elements, resulting in  $O(s^2)$  time. For Step 7, the Sinkhorn-scaling algorithm requires  $O(Hs)$  time by using sparse matrix multiplications. Calculating GW distance using the sparse  $\tilde{\mathbf{T}}^{(R)}$  requires  $O(s^2)$  operations. Therefore, for Algorithm 2, its overall time complexity is  $O(mn + Rs^2 + RHs + s^2)$ , which becomes  $O(mn + s^2)$  when  $R$  and  $H$  are constants, and its memory cost is  $O(mn)$ . When  $m = O(n)$ , we obtain the complexity shown in Table 1.

*Applicability for entropic GW distance and fused GW distance.* As shown in Algorithm 2, our algorithm can approximate the

**Algorithm 2** SPAR-GW algorithm

- 1: **Input:** Sample distributions  $\mathbf{a}, \mathbf{b}$ , relation matrices  $\mathbf{C}^X, \mathbf{C}^Y$ , ground cost function  $\mathcal{L}$ , regularization parameter  $\varepsilon$ , number of selected elements  $s$ , number of outer/inner iterations  $R, H$
- 2: Construct the sampling probability  $\mathbf{P}$  defined by (5)  $O(mn)$
- 3: Generate an iid subsample of size  $s$  using  $\mathbf{P}$ , let  $\mathcal{S} = \{(i'_l, j'_l)\}_{l=1}^s$  be the index set
- 4: Initialize  $\tilde{\mathbf{T}}^{(0)} = \mathbf{0}_{m \times n}$  and  $\tilde{T}_{ij}^{(0)} = a_i b_j$  if  $(i, j) \in \mathcal{S}$
- 5: **For**  $r = 0$  to  $R - 1$ :
- 6: **Construct a sparse kernel matrix:**  $O(s^2)$ 
  - a) Compute  $\tilde{\mathcal{C}}(\tilde{\mathbf{T}}^{(r)}) = \sum_{(i,j) \in \mathcal{S}} \tilde{\mathbf{L}}_{i,j} \tilde{T}_{ij}^{(r)}$  using (6), and replace its 0's at  $\mathcal{S}$  with  $\infty$ 's
  - b)  $\tilde{\mathbf{K}}^{(r)} = \begin{cases} \exp(-\frac{\tilde{\mathcal{C}}(\tilde{\mathbf{T}}^{(r)})}{\varepsilon}) \odot \tilde{\mathbf{T}}^{(r)} \odot (s\mathbf{P}) & \text{if } \mathcal{R}(\mathbf{T}) = \text{KL}(\mathbf{T} \parallel \mathbf{T}^{(r)}) \\ \exp(-\frac{\tilde{\mathcal{C}}(\tilde{\mathbf{T}}^{(r)})}{\varepsilon}) \odot (s\mathbf{P}) & \text{if } \mathcal{R}(\mathbf{T}) = H(\mathbf{T}) \end{cases}$
- 7: **Sinkhorn-scaling (with sparse inputs):**  $\tilde{\mathbf{T}}^{(r+1)} = \text{SINKHORN}(\mathbf{a}, \mathbf{b}, \tilde{\mathbf{K}}^{(r)}, H)$   $O(Hs)$
- 8: **Output:**  $\widehat{\text{GW}} = \sum_{(i,j) \cap (i',j') \in \mathcal{S}} \mathcal{L}(C_{ii'}^X, C_{jj'}^Y) \tilde{T}_{ij}^{(R)} \tilde{T}_{i'j'}^{(R)}$   $O(s^2)$

entropic GW distance as well. Moreover, it is natural to extend the algorithm to approximate the fused Gromov-Wasserstein (FGW) distance (Titouan et al. 2019a; Vayer et al. 2020). In particular, when computing the FGW distance, the cost matrix  $\mathcal{C}(\mathbf{T})$  takes the direct comparison among the samples into account, and our importance sparsification mechanism is still applicable. Details for the modified algorithm are relegated to supplementary material.

#### 4. Theoretical Results

This section shows the convergence and consistency of  $\widehat{\text{GW}}$  obtained by Algorithm 2 under  $\mathcal{R}(\mathbf{T}) = H(\mathbf{T})$ . To ease the conversation, we focus on the case that  $m = n$ , and the extension to unequal cases is straightforward. Let  $(\mathbf{T}, \mathbf{T}') \in \Pi^2(\mathbf{a}, \mathbf{b})$ . We define

$$\mathcal{E}(\mathbf{T}, \mathbf{T}') := \sum_{i, i', j, j'} \mathcal{L}(C_{ii'}^X, C_{jj'}^Y) T_{ij} T'_{i'j'},$$

$$G(\mathbf{T}) := \mathcal{E}(\mathbf{T}, \mathbf{T}) - \min_{\mathbf{T}' \in \Pi(\mathbf{a}, \mathbf{b})} \mathcal{E}(\mathbf{T}, \mathbf{T}').$$

For notation simplicity, we overload  $\mathcal{E}(\mathbf{T}, \mathbf{T})$  as  $\mathcal{E}(\mathbf{T})$ . Following Kerdoncuff, Emonet, and Sebban (2021), our goal is to provide a guarantee on the convergence of  $G(\mathbf{T})$ , because  $\mathbf{T}$  is a stationary point of  $\mathcal{E}(\mathbf{T})$  if and only if  $G(\mathbf{T}) = 0$  (Reddi et al. 2016). In addition, we define  $\tilde{\mathbf{K}}^{(r)} = \exp(-\mathcal{C}(\tilde{\mathbf{T}}^{(r)})/\varepsilon)$ , which is the unsampled counterpart to  $\tilde{\mathbf{K}}^{(r)}$  at the  $r$ th iteration. Consider the following regularity conditions.

- (H.1) The relation matrices  $\mathbf{C}^X, \mathbf{C}^Y$  are symmetric;
- (H.2) The ground cost is bounded, that is,  $0 \leq \mathcal{L}(C_{ii'}^X, C_{jj'}^Y) \leq 2B$  for a constant  $B > 0$ ;
- (H.3)  $\|\tilde{\mathbf{K}}^{(r)}\|_2 \geq n^\alpha/c_1$  for constants  $1/2 < \alpha \leq 1$  and  $c_1 > 0$ , and the condition number of  $\tilde{\mathbf{K}}^{(r)}$  is positive and bounded by  $c_2 > 0$ , for any  $r \leq R - 1$ ;
- (H.4)  $p_{ij} \geq c_3/n^2$  for a constant  $c_3 > 0$ ;
- (H.5)  $s \geq c_4 n^{3-2\alpha} \log^4(n)$  for constants  $c_4 = 8 \log^4(2)/(c_3 \log^4(1 + \varepsilon))$  and  $\varepsilon > 0$ .

Conditions (H.1)–(H.3) are natural. Condition (H.4) indicates that the sampling probabilities could not be too small, requiring  $p_{ij}$  to be at the order of  $O(1/n^2)$ . The order can always be satisfied by linear interpolating between the proposed sampling probability and the uniform sampling probability. Such a shrinkage strategy is commonly used in subsampling literature (Ma, Mahoney, and Yu 2015; Yu et al. 2022). Condition (H.5) requires  $s$  to be large enough. For a general case that  $\|\tilde{\mathbf{K}}^{(r)}\|_2 = O(n)$ , that is,  $\alpha = 1$ , condition (H.5) can be achieved when  $s = O(n^{1+\delta})$  for an arbitrary small  $\delta > 0$ . Such order indicates we only need to compute around  $n^2$  elements from the entire tensor that contains  $n^4$  elements.

We now provide our main convergence result, whose proofs are provided in supplementary material.

**Theorem 1.** Under the conditions (H.1)–(H.5), assume that  $R \exp(-16 \log^4(n)/\varepsilon^4) \rightarrow 0$  for some  $\varepsilon > 0$ , and  $n > 76$ . The following bound holds in probability

$$G(\tilde{\mathbf{T}}^{(R-1)}) \leq \frac{\mathcal{E}(\tilde{\mathbf{T}}^{(R)}) - \mathcal{E}(\tilde{\mathbf{T}}^{(R-1)})}{2} + 6\sqrt{2}\varepsilon(2 + \varepsilon)c_1c_2\sqrt{\frac{n^{3-2\alpha}}{c_3s}} + \varepsilon \log(n) + Bn^2 \|\tilde{\mathbf{T}}^{(R)} - \tilde{\mathbf{T}}^{(R-1)}\|_F^2. \quad (7)$$

Consider the upper bound in Theorem 1. The second term on the right-hand side of (7) results from importance sparsification. Under the common condition that  $\alpha = 1$ , it tends to zero when  $s = O(n^{1+\delta})$  as  $n \rightarrow \infty$ . The third term  $\varepsilon \log(n)$  is caused by the regularization mechanism, which goes to zero when  $\varepsilon = o(\log^{-1}(n))$ . The remaining terms are due to the iterative scheme in Algorithm 2. Theorem 1 indicates that the estimation error of the proposed estimator decreases when the regularization parameter  $\varepsilon$  decreases or the subsample size  $s$  increases. We provide the following corollary to show the consistency of the proposed estimator.

**Corollary 1.** Suppose the conditions of Theorem 1 hold with  $\alpha = 1$ . Further suppose Algorithm 2 converges with  $\|\tilde{\mathbf{T}}^{(R)} - \tilde{\mathbf{T}}^{(R-1)}\|_F \leq c_5/n^{3/2+\eta}$  for some  $c_5, \eta > 0$ . When  $s = O(n^{1+\delta})$  for any  $\delta > 0$  and  $\varepsilon = o(\log^{-1}(n))$ ,  $G(\tilde{\mathbf{T}}^{(R-1)}) \rightarrow 0$  in probability, as  $n \rightarrow \infty$ .

The local stationary convergence of Algorithm 2 implies  $\tilde{\mathbf{T}}^{(R)}$  and  $\tilde{\mathbf{T}}^{(R-1)}$  will get closer and closer with the increase of  $R$ . Therefore, for any given  $n$ , we can set the assumption  $\|\tilde{\mathbf{T}}^{(R)} - \tilde{\mathbf{T}}^{(R-1)}\|_F \leq c_5/n^{3/2+\eta}$  as the stopping criterion, which can be naturally satisfied for a relatively large  $R$ .

## 5. Importance Sparsification for UGW Distance

### 5.1. Unbalanced GW Distance

In this section, we extend the SPAR-GW algorithm to approximate the unbalanced Gromov-Wasserstein (UGW) distance. Similar to the unbalanced optimal transport (UOT) (Liero, Mielke, and Savaré 2016; Chizat et al. 2018a, 2018b, 2018c), UGW distance is able to compare metric-measure spaces endowed with arbitrary positive distributions  $\mathbf{a} \in \mathbb{R}_+^m, \mathbf{b} \in \mathbb{R}_+^n$  (Séjourné, Vialard, and Peyré 2021; Kawano and Mason 2021; Luo et al. 2022). Following the definition in Séjourné,

Vialard, and Peyré (2021), UGW distance relaxes the marginal constraints via the quadratic KL-divergence  $\text{KL}^\otimes(\mu\|\nu) = \text{KL}(\mu \otimes \mu\|\nu \otimes \nu)$ , where  $\mu \otimes \nu$  is the tensor product measure defined by  $d(\mu \otimes \nu)(x, y) = d\mu(x)d\nu(y)$ . In particular, UGW distance takes the form

$$\begin{aligned} \text{UGW}((\mathbf{C}^X, \mathbf{a}), (\mathbf{C}^Y, \mathbf{b})) & := \min_{\mathbf{T} \in \mathbb{R}_+^{m \times n}} \langle \mathcal{L}(\mathbf{C}^X, \mathbf{C}^Y) \otimes \mathbf{T}, \mathbf{T} \rangle + \lambda \text{KL}^\otimes(\mathbf{T}\mathbf{1}_n\|\mathbf{a}) \\ & \quad + \lambda \text{KL}^\otimes(\mathbf{T}^\top \mathbf{1}_m\|\mathbf{b}) \\ & = \min_{\mathbf{T} \in \mathbb{R}_+^{m \times n}} \langle \mathcal{C}_{\text{un}}(\mathbf{T}), \mathbf{T} \rangle + \lambda m(\mathbf{T}) \text{KL}(\mathbf{T}\mathbf{1}_n\|\mathbf{a}) \\ & \quad + \lambda m(\mathbf{T}) \text{KL}(\mathbf{T}^\top \mathbf{1}_m\|\mathbf{b}). \end{aligned}$$

Here,  $m(\mathbf{T}) = \sum_{i,j} T_{ij}$  is the total mass of  $\mathbf{T}$ ,  $\mathcal{C}_{\text{un}}(\mathbf{T}) := \mathcal{L}(\mathbf{C}^X, \mathbf{C}^Y) \otimes \mathbf{T} + E(\mathbf{T})$  is the cost matrix with  $E(\mathbf{T}) := \lambda \sum_i \log(\sum_j T_{ij}/a_i) \sum_j T_{ij} + \lambda \sum_j \log(\sum_i T_{ij}/b_j) \sum_i T_{ij}$ , and  $\lambda > 0$  is the marginal regularization parameter balancing the tradeoff between mass transportation and mass variation. Note that when  $\mathbf{a} \in \Delta^{m-1}$  and  $\mathbf{b} \in \Delta^{n-1}$ , UGW distance degenerates to the classical GW distance as  $\lambda \rightarrow \infty$ .

### 5.2. Proposed Algorithm

To approximate the UGW distance, we update the coupling matrix via proximal gradient algorithm (PGA) by adding a Bregman proximal term (Xie et al. 2020; Xu et al. 2019; Kerdoncuff, Emonet, and Sebban 2021). Specifically,  $\mathbf{T}$  is updated as

$$\begin{aligned} \mathbf{T}^{(r+1)} & = \text{argmin}_{\mathbf{T} \in \mathbb{R}_+^{m \times n}} \langle \mathcal{C}_{\text{un}}(\mathbf{T}^{(r)}), \mathbf{T} \rangle + \lambda m(\mathbf{T}^{(r)}) \text{KL}(\mathbf{T}\mathbf{1}_n\|\mathbf{a}) \\ & \quad + \lambda m(\mathbf{T}^{(r)}) \text{KL}(\mathbf{T}^\top \mathbf{1}_m\|\mathbf{b}) + \varepsilon m(\mathbf{T}^{(r)}) \text{KL}(\mathbf{T}\|\mathbf{T}^{(r)}). \end{aligned} \quad (8)$$

The subproblem (8) can be solved using unbalanced Sinkhorn-scaling algorithm (Chizat et al. 2018b; Pham et al. 2020) with the kernel matrix  $\mathbf{K} = \exp\{-\mathcal{C}_{\text{un}}(\mathbf{T}^{(r)})/(\varepsilon m(\mathbf{T}^{(r)}))\} \odot \mathbf{T}^{(r)}$ ; see Step 9 in Algorithm 3 for details. Now the convergent scaling vectors  $\mathbf{u} \in \mathbb{R}_+^m, \mathbf{v} \in \mathbb{R}_+^n$  satisfies that

$$(u_i)^{\frac{\lambda+\varepsilon}{\lambda}} \left( \sum_j K_{ij} v_j \right) = a_i \quad \text{and} \quad (v_j)^{\frac{\lambda+\varepsilon}{\lambda}} \left( \sum_i K_{ij} u_i \right) = b_j.$$

As a result, it holds that

$$\begin{aligned} (u_i)^{\frac{\lambda+\varepsilon}{\lambda}} K_{ij} v_j & \leq a_i, \\ u_i K_{ij} (v_j)^{\frac{\lambda+\varepsilon}{\lambda}} & \leq b_j \quad \Rightarrow \quad (u_i)^{\frac{2\lambda+\varepsilon}{\lambda}} K_{ij}^2 (v_j)^{\frac{2\lambda+\varepsilon}{\lambda}} \leq a_i b_j, \end{aligned}$$

which follows that  $T_{ij}^* = u_i K_{ij} v_j \leq (a_i b_j)^{\frac{\lambda}{2\lambda+\varepsilon}} K_{ij}^{\frac{\varepsilon}{2\lambda+\varepsilon}}$ . Such an inequality motivates us to sample with probability

$$p_{ij} = \frac{(a_i b_j)^{\frac{\lambda}{2\lambda+\varepsilon}} K_{ij}^{\frac{\varepsilon}{2\lambda+\varepsilon}}}{\sum_{i,j} (a_i b_j)^{\frac{\lambda}{2\lambda+\varepsilon}} K_{ij}^{\frac{\varepsilon}{2\lambda+\varepsilon}}}, \quad 1 \leq i \leq m, \quad 1 \leq j \leq n. \quad (9)$$

Unfortunately, such a probability involves the kernel matrix  $\mathbf{K}$ , which requires the unknown coupling matrix  $\mathbf{T}$ . To bypass the obstacle, we propose to replace the unknown  $\mathbf{T}$  with the initial value  $\tilde{\mathbf{T}}^{(0)} = \mathbf{a}\mathbf{b}^\top / \sqrt{m(\mathbf{a})m(\mathbf{b})}$ , where  $m(\mathbf{a}) = \sum_i a_i$  and  $m(\mathbf{b}) = \sum_j b_j$  are the total mass of  $\mathbf{a}$  and  $\mathbf{b}$ , respectively. Algorithm 3 details the proposed SPAR-UGW algorithm for approximating UGW distances. Note that when  $m(\mathbf{a}) = m(\mathbf{b})$

### Algorithm 3 SPAR-UGW algorithm

- 1: **Input:** Sample distributions  $\mathbf{a}, \mathbf{b}$ , relation matrices  $\mathbf{C}^X, \mathbf{C}^Y$ , ground cost function  $\mathcal{L}$ , regularization parameters  $\lambda, \varepsilon$ , number of selected elements  $s$ , number of outer/inner iterations  $R, H$
- 2: Initialize  $\tilde{\mathbf{T}}^{(0)} = \mathbf{a}\mathbf{b}^\top / \sqrt{m(\mathbf{a})m(\mathbf{b})}$
- 3:  $\mathbf{K} = \exp\{-\mathcal{C}_{\text{un}}(\tilde{\mathbf{T}}^{(0)})/(\varepsilon m(\tilde{\mathbf{T}}^{(0)}))\} \odot \tilde{\mathbf{T}}^{(0)} \quad O(m^2 n^2)$  or  $O(mn)$
- 4: Construct the sampling probability  $\mathbf{P}$  defined by (9)  $O(mn)$
- 5: Generate an iid subsample of size  $s$  using  $\mathbf{P}$ , let  $\mathcal{S} = \{(i'_l, j'_l)\}_{l=1}^s$  be the index set
- 6: **For**  $r = 0$  **to**  $R - 1$ :
- 7:  $\bar{\varepsilon} = \varepsilon m(\tilde{\mathbf{T}}^{(r)}), \bar{\lambda} = \lambda m(\tilde{\mathbf{T}}^{(r)})$
- 8: **Construct a sparse kernel matrix:**  $O(s^2)$ 
  - (a) Compute the cost matrix  $\tilde{\mathcal{C}}_{\text{un}}(\tilde{\mathbf{T}}^{(r)}) = \sum_{(i,j) \in \mathcal{S}} \tilde{\mathbf{I}}_{i,j} \tilde{T}_{ij}^{(r)} + E(\tilde{\mathbf{T}}^{(r)})$  using formula (6) and replace its 0's at  $\mathcal{S}$  with  $\infty$ 's
  - (b)  $\tilde{\mathbf{K}}^{(r)} = \exp(-\frac{\tilde{\mathcal{C}}_{\text{un}}(\tilde{\mathbf{T}}^{(r)})}{\bar{\varepsilon}}) \odot \tilde{\mathbf{T}}^{(r)} \odot (s\mathbf{P})$
- 9: **Unbalanced Sinkhorn-scaling:**  $\tilde{\mathbf{T}}^{(r+1)} = \text{SINKHORN}_{\text{UOT}}(\mathbf{a}, \mathbf{b}, \tilde{\mathbf{K}}^{(r)}, \bar{\lambda}, \bar{\varepsilon}, H) \quad O(Hs)$ 
  - (a) Initialize  $\mathbf{u}^{(0)} = \mathbf{1}_m, \mathbf{v}^{(0)} = \mathbf{1}_n$
  - (b) **For**  $h = 0$  **to**  $H - 1$ :
$$\mathbf{u}^{(h+1)} = (\mathbf{a} \odot (\tilde{\mathbf{K}}^{(r)} \mathbf{v}^{(h)}))^{\bar{\lambda}/(\bar{\lambda}+\bar{\varepsilon})}, \quad \mathbf{v}^{(h+1)} = (\mathbf{b} \odot (\tilde{\mathbf{K}}^{(r)\top} \mathbf{u}^{(h+1)}))^{\bar{\lambda}/(\bar{\lambda}+\bar{\varepsilon})}$$
  - (c)  $\tilde{\mathbf{T}}^{(r+1)} = \text{diag}(\mathbf{u}^{(H)}) \tilde{\mathbf{K}}^{(r)} \text{diag}(\mathbf{v}^{(H)})$
- 10:  $\tilde{\mathbf{T}}^{(r+1)} = \sqrt{m(\tilde{\mathbf{T}}^{(r)})/m(\tilde{\mathbf{T}}^{(r+1)})} \cdot \tilde{\mathbf{T}}^{(r+1)}$
- 11: **Output:**  $\widehat{\text{UGW}} = \langle \mathcal{L}(\mathbf{C}^X, \mathbf{C}^Y) \otimes \tilde{\mathbf{T}}^{(R)}, \tilde{\mathbf{T}}^{(R)} \rangle + \lambda \text{KL}^\otimes(\tilde{\mathbf{T}}^{(R)}\mathbf{1}_n\|\mathbf{a}) + \lambda \text{KL}^\otimes(\tilde{\mathbf{T}}^{(R)\top}\mathbf{1}_m\|\mathbf{b}) \quad O(s^2)$

and  $\lambda \rightarrow \infty$ , SPAR-UGW degenerates to SPAR-GW. Such an observation is consistent with the relationship between GW and UGW.

*Computational cost.* In Algorithm 3, although calculating  $\mathbf{K}$  in Step 3 requires  $O(m^2 n^2)$  time, we only need to calculate it once. Moreover, when  $\mathcal{L}$  is decomposable, the complexity of calculating  $\mathbf{K}$  can be reduced to  $O(mn)$  by using the fact that  $\mathbf{T}^{(0)}$  is a rank-one matrix. Therefore, the total time complexity of SPAR-UGW is  $O(mn + s^2)$  when  $R$  and  $H$  are constants, and  $\mathcal{L}$  is decomposable.

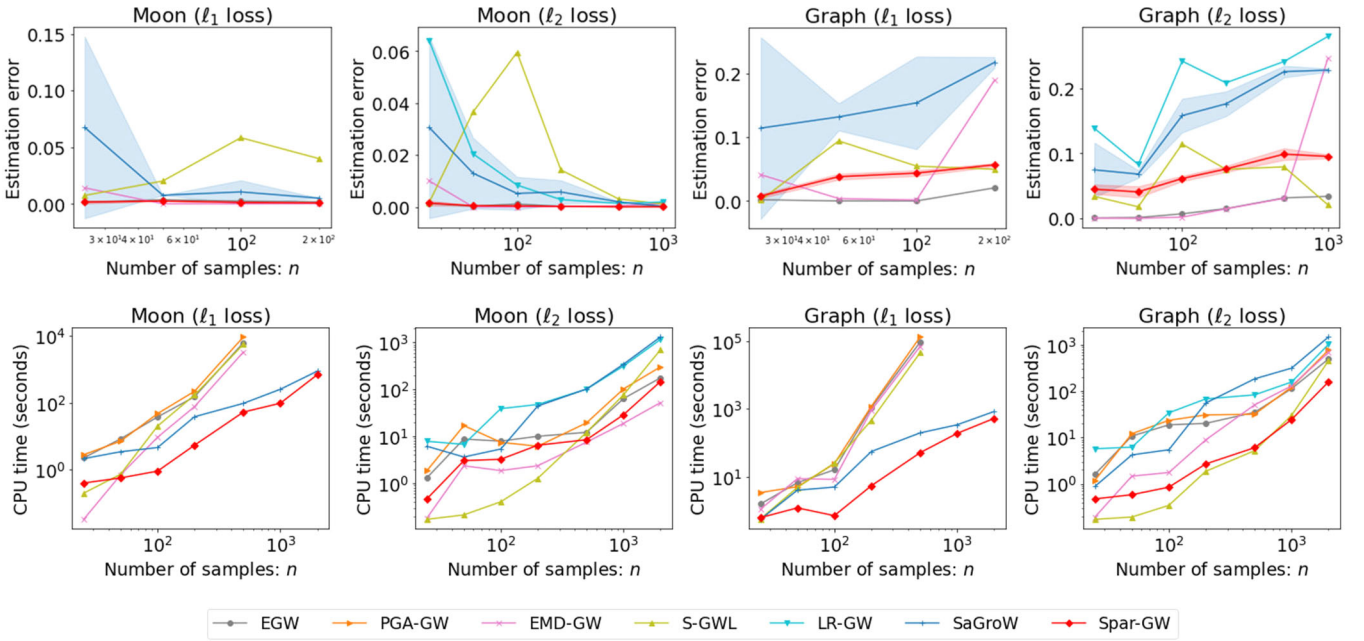
## 6. Experiments

In this section, we evaluate the performance of SPAR-GW and its variants in both distance estimation and graph analysis.

### 6.1. Synthetic Data Analysis

#### 6.1.1. Approximation of GW and UGW Distances

We compare the proposed SPAR-GW with main competitors including: (i) EGW, Algorithm 1 with entropic regularization (Peyré, Cuturi, and Solomon 2016); (ii) PGA-GW, Algorithm 1 with proximal regularization (Xu et al. 2019); (iii) EMD-GW, that is, EGW with  $\varepsilon = 0$ , but replacing Sinkhorn-scaling algorithm in EGW with the solver for unregularized OT prob-



**Figure 2.** Comparison on estimation error (Top) and CPU time (Bottom). The mean and standard deviation are reported for sampling-based methods.

lems (Bonneel et al. 2011); (iv) S-GWL (Xu, Luo, and Carin 2019), adapted for arbitrary ground cost following Kerdoncuff, Emonet, and Sebban (2021); (v) LR-GW, the quadratic approach in Scetbon, Peyré, and Cuturi (2022); (vi) SaGroW (Kerdoncuff, Emonet, and Sebban 2021). Other methods in Table 1, that is, Sliced GW (Titouan et al. 2019b), AE (Sato et al. 2020) and FlowAlign (Le, Ho, and Yamada 2021), are not included as they fail to approximate the original GW distance. We adopted the proximal term, that is, KL-divergence, as  $\mathcal{R}(\mathbf{T})$  in SaGroW and SPAR-GW. The other choice of regularization term yields similar results. The regularization parameter  $\varepsilon$  is chosen among  $\{1, 10^{-1}, 10^{-2}, 10^{-3}\}$  and the result with the smallest distance w.r.t. each method is presented. For LR-GW, the nonnegative rank of the coupling matrix is set to  $\lceil n/20 \rceil$ . For SPAR-GW and SPAR-UGW, we set the subsample size  $s = 16n$ . For a fair comparison, we set the subsample size  $s' = s^2/n^2$  for SaGroW to ensure that it has the same sampling budget (i.e., samples the same number of elements) to SPAR-GW (or SPAR-UGW). Other parameters of the methods mentioned above are set by default. To take into account the randomness of sampling-based methods, that is, SaGroW, SPAR-GW, and SPAR-UGW, their estimations are averaged over 10 runs. All the experiments are performed on a server with 8-core CPUs and 30GB RAM.

We consider two popular synthetic datasets called Moon following Séjourné, Vialard, and Peyré (2021), Muzellec et al. (2020), and Graph following Xu et al. (2019), Xu, Luo, and Carin (2019). We also consider two other widely used datasets including the case where the source and target are distributed in heterogeneous spaces. The results have a similar pattern to those of Moon and are relegated to supplementary material. Specifically, for the Moon dataset, marginals are two Gaussian distributions,  $N(n/3, n/20)$  and  $N(n/2, n/20)$ , supported on  $n$  points in  $\mathbb{R}^2$ . The source and target supported points are respectively generated from two interleaving half circles by sklearn toolbox (Pedregosa et al. 2011). The matrices  $\mathbf{C}^X, \mathbf{C}^Y$  are defined using pairwise Euclidean distances in  $\mathbb{R}^2$ . For the Graph dataset,

we first generate one graph with  $n$  nodes and power-law degree distribution from NetworkX library (Hagberg, Schult, and Swart 2008), and then generate the other graph by adding extra edges randomly with probability 0.2 on the first one. Their degree distributions are used as two marginals, and the adjacency matrix of each graph is used as  $\mathbf{C}^X, \mathbf{C}^Y$ . Both  $\ell_1$  and  $\ell_2$  losses are considered for the ground cost. LR-GW is only capable of the  $\ell_2$  loss, and thus its result w.r.t. the  $\ell_1$  loss is omitted. To compare the estimation accuracy w.r.t. different methods, we take PGA-GW as a benchmark and calculate the absolute error between its estimation and other estimations of GW distance.

Figure 2 shows the estimation error (top row) and the CPU time (bottom row) versus increasing sample size  $n$ . We observe that the proposed SPAR-GW method yields almost the smallest estimation error for the Moon dataset and reasonable errors for the Graph dataset. Such a difference is because Gaussian distributions in Moon are more concentrated and thus are easier to sketch by subsamples; while the structure of graphs in Graph is more complicated, and the transportation between their degree distributions is also more difficult to approximate by the subsampling technique. As for computational efficiency, SPAR-GW requires less CPU time than most of the competitors, and such an advantage is more obvious for the indecomposable  $\ell_1$  loss. These observations indicate SPAR-GW is capable of dealing with large-scale GW problems with arbitrary ground cost.

For unbalanced problems, we set the total mass of  $\mathbf{a}, \mathbf{b}$  to be units and the marginal relaxation parameter to be  $\lambda = 1$ . We compare SPAR-UGW with: (i) Naive transport plan  $\mathbf{T} = \mathbf{a}\mathbf{b}^\top$ ; (ii) EUGW, entropic regularization in Séjourné, Vialard, and Peyré (2021); (iii) PGA-UGW; (iv) SaGroW, adapted for unbalanced problems. We calculate the estimation error w.r.t. the PGA-UGW benchmark. Other settings are the same as the aforementioned. From Figure 3, we observe that SPAR-UGW consistently achieves the best accuracy for the former dataset and a relatively small estimation error for the latter one, requiring the least amount of time for the  $\ell_2$  loss. Although the



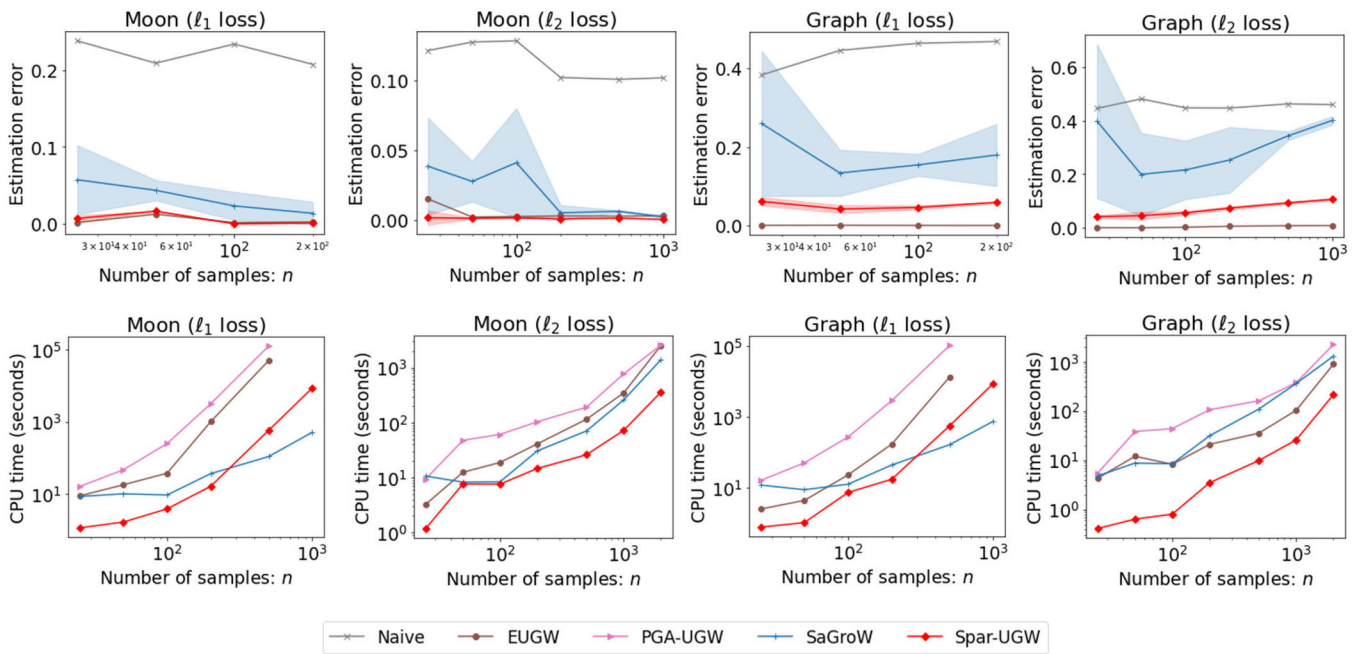


Figure 3. Comparison on estimation error (Top) and CPU time (Bottom) w.r.t. UGW distance. The mean and standard deviation are reported for sampling-based methods.

computational cost of SPAR-UGW becomes more considerable for the indecomposable  $\ell_1$  loss, it still computes much faster than the classical EUGW and PGA-UGW methods.

### 6.1.2. Sensitivity Analysis

We now show that our method is robust to hyperparameters by analyzing its sensitivity to the subsample size  $s$  and the regularization parameter  $\varepsilon$ . Specifically, for synthetic datasets with fixed sample size  $n = 200$ , the hyperparameters are considered among  $s \in \{2^1, 2^2, \dots, 2^5\} \times n$  and  $\varepsilon \in \{5^0, 5^{-1}, \dots, 5^{-4}\}$ .

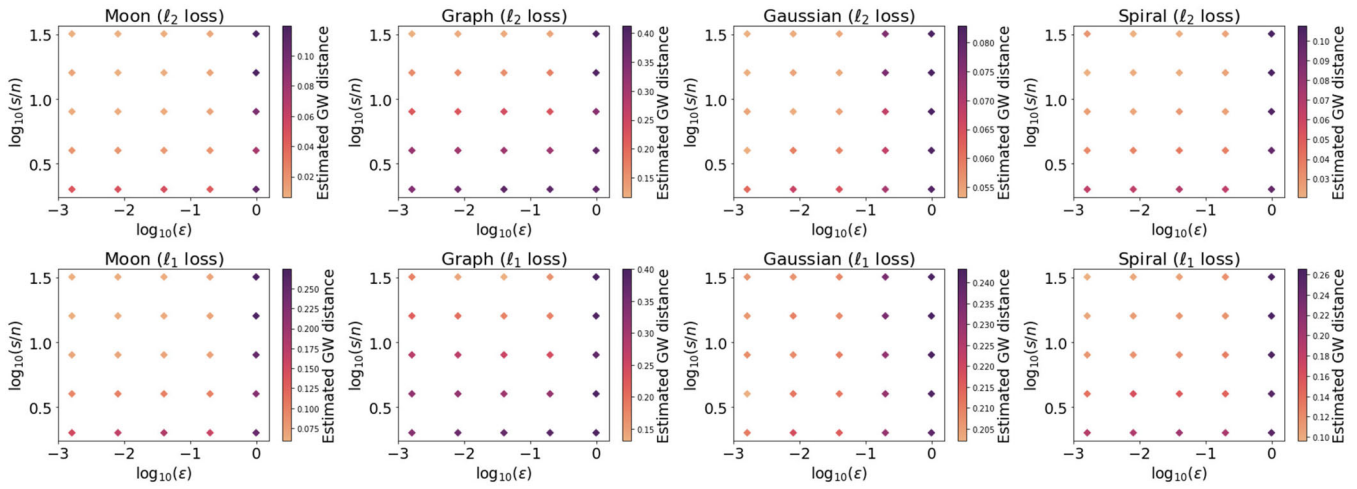
From the results in Figure 4, we find that a large number of selected elements  $s$  and/or a small value of  $\varepsilon$  is associated with a small GW distance estimation and a long CPU time. Such a finding is consistent with our theoretical results. We also observe that SPAR-GW can yield a satisfactory estimation for a large range of hyperparameters. More precisely, as long as  $s = O(n)$  and  $\varepsilon$  is not too large, the estimated GW distance is approximately in the same order, which implies SPAR-GW is not sensitive to hyperparameters and can cover a wide range of tradeoffs between accuracy and speed. This observation supports the key assumption that only a few important elements in kernel and coupling matrices are required to approximate the GW distance effectively. Moreover, our method is largely free from numerical instability because  $\varepsilon$  need not be extremely small, which is in good agreement with the statements in Xie et al. (2020) and Xu et al. (2019).

## 6.2. Real-World Applications

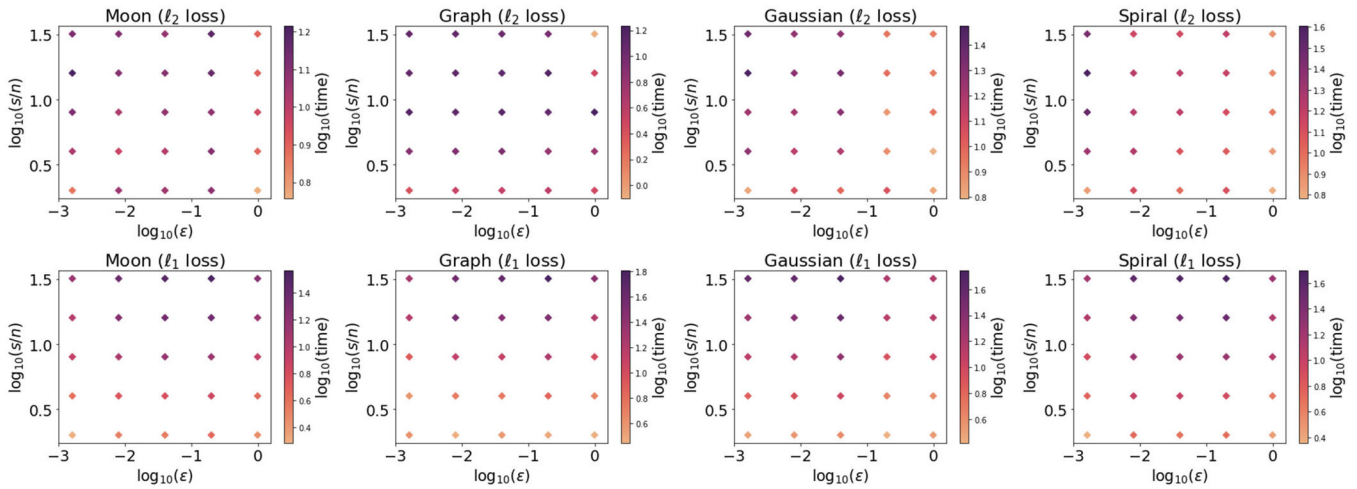
We consider two applications, graph clustering and graph classification, to demonstrate the effectiveness of our method in applications. Six widely-used benchmark datasets are considered: BZR, COX2 (Sutherland, O’Brien, and Weaver 2003), CUNEIFORM (Kriege et al. 2018), SYNTHETIC (Feragen et al.

2013) with vector node attributes; FIRSTMM\_DB (Neumann et al. 2013) with discrete attributes; and IMDB-B (Yanardag and Vishwanathan 2015) with no attributes. All these datasets are available in PyTorch Geometric library (Fey and Lenssen 2019). Given  $N$  graphs, we first compute the pairwise GW distance matrix  $\mathbf{D} \in \mathbb{R}^{N \times N}$  and then construct the similarity matrix  $\mathbf{S} = \exp(-\mathbf{D}/\gamma)$  for  $\gamma > 0$ . For methods that can directly extend to approximate the fused GW (FGW) distance, including EGW, PGA-GW, EMD-GW, SaGroW, and SPAR-GW, we obtain the pairwise FGW distance matrix when the graphs have attributes. We set the tradeoff parameter  $\alpha = 0.6$ . Empirical results show the performance is not sensitive to  $\alpha$ . For the graph clustering task, we apply spectral clustering to the similarity matrix. We replicate the experiment 10 times with different random initialization and assess the clustering performance by average Rand index (RI) (Rand 1971). For the classification task, we train a classifier based on kernel SVM using the similarity matrix and test the classifier via nested 10-fold cross-validation following Titouan et al. (2019a). The performance is evaluated by average classification accuracy. To examine the effect of different loss functions, we consider both  $\ell_1$  loss and  $\ell_2$  loss for AE, SaGroW, and SPAR-GW. Other methods are mainly designed for the decomposable loss, and thus only the  $\ell_2$  loss is implemented. For all methods,  $\gamma$  is cross validated within  $\{2^{-10}, 2^{-9}, \dots, 2^{10}\}$ . Other settings are the same as those in the previous section.

Tables 2 and 3 report the average RI and average classification accuracy with the corresponding standard deviation, respectively. PGA-GW and EMD-GW are excluded for clarity as their results are similar to EGW. Sliced GW and FlowAlign are also not included since they cannot handle graphs. From Tables 2 and 3, we observe the proposed SPAR-GW approach is superior or at least comparable to other methods in all cases. We also observe the SPAR-GW with  $\ell_1$  cost almost consistently outperforms the one with  $\ell_2$  cost. This observation is consistent



(a) Sensitivity on the estimated GW distance.



(b) Sensitivity on the CPU time.

**Figure 4.** Impact of the subsample size  $s$  and the regularization parameter  $\epsilon$  for Spar-GW on the GW distance estimation (panel (a)) and computational time (panel (b)). The mean over 10 runs are reported.

**Table 2.** Comparison on clustering performance w.r.t. RI (%).

Dataset	SYNTHETIC	BZR	Cuneiform	COX2	FIRSTMM_DB	IMDB-B
# graphs: $N$	300	405	267	467	41	1000
Ave. # nodes: $n$	100.00	35.75	21.27	41.22	1377.27	19.77
Subsample size: $s$	$2^5 \times n$	$2^3 \times n$	$2^3 \times n$	$2^3 \times n$	$2^7 \times n$	$2^3 \times n$
EGW	<b>100.00</b> $\pm 0.00$	67.18 $\pm 0.44$	<b>94.90</b> $\pm 0.08$	64.81 $\pm 0.58$	<b>92.51</b> $\pm 0.15$	50.79 $\pm 0.00$
S-GWL	<b>100.00</b> $\pm 0.00$	66.84 $\pm 0.73$	94.32 $\pm 0.07$	65.02 $\pm 0.23$	81.42 $\pm 0.16$	<b>51.30</b> $\pm 0.00$
LR-GW	50.13 $\pm 0.02$	65.34 $\pm 4.31$	26.47 $\pm 0.57$	64.99 $\pm 0.10$	45.93 $\pm 3.14$	<b>51.54</b> $\pm 0.01$
AE ( $\ell_2$ loss)	50.17 $\pm 0.59$	67.04 $\pm 0.00$	82.51 $\pm 3.24$	62.36 $\pm 0.03$	84.63 $\pm 0.00$	50.79 $\pm 0.03$
AE ( $\ell_1$ loss)	50.17 $\pm 0.59$	67.04 $\pm 0.00$	82.64 $\pm 3.22$	62.36 $\pm 0.00$	84.67 $\pm 0.11$	50.79 $\pm 0.03$
SaGroW ( $\ell_2$ loss)	52.41 $\pm 0.00$	67.24 $\pm 0.26$	94.56 $\pm 0.20$	<b>65.94</b> $\pm 0.92$	92.07 $\pm 0.00$	50.45 $\pm 0.00$
SaGroW ( $\ell_1$ loss)	54.15 $\pm 0.19$	<b>67.33</b> $\pm 0.47$	94.54 $\pm 0.14$	<b>65.97</b> $\pm 1.03$	92.09 $\pm 0.35$	50.45 $\pm 0.00$
Spar-GW ( $\ell_2$ loss)	<b>98.67</b> $\pm 0.00$	<b>68.22</b> $\pm 0.00$	<b>94.66</b> $\pm 0.05$	65.54 $\pm 0.00$	<b>92.24</b> $\pm 0.30$	<b>50.82</b> $\pm 0.00$
Spar-GW ( $\ell_1$ loss)	<b>98.67</b> $\pm 0.00$	<b>68.22</b> $\pm 0.00$	<b>94.64</b> $\pm 0.06$	<b>66.29</b> $\pm 1.09$	<b>92.41</b> $\pm 0.33$	<b>50.82</b> $\pm 0.00$

\* The top-3 results of each dataset are in bold, and the best result is in italics.

with the observation in Kerdoncuff, Emonet, and Sebban (2021), which stated that the  $\ell_1$  cost tends to yield better performance than the  $\ell_2$  cost in graphical data analysis. Such observation also justifies the essence of developing a computational tool that can handle arbitrary ground costs in GW distance approximation.

Considering the CPU time, SPAR-GW computes much faster than EGW, S-GWL, and AE when the number of nodes is relatively large. Take the FIRSTMM\_DB dataset as an example, in which each graph has an average of 1,377 nodes, the average CPU time for these methods are 414.82s (EGW), 1059.95s

**Table 3.** Comparison on classification performance w.r.t. accuracy (%).

Dataset	SYNTHETIC	BZR	Cuneiform	COX2	FIRSTMM_DB	IMDB-B
EGW	<b>100.00</b> $\pm$ 0.00	<b>85.92</b> $\pm$ 0.50	<b>25.66</b> $\pm$ 2.01	<b>80.21</b> $\pm$ 0.59	<b>53.75</b> $\pm$ 2.24	66.01 $\pm$ 0.64
S-GWL	<b>100.00</b> $\pm$ 0.00	<b>87.67</b> $\pm$ 0.41	9.05 $\pm$ 1.13	78.23 $\pm$ 0.22	17.50 $\pm$ 4.03	64.54 $\pm$ 0.55
LR-GW	55.83 $\pm$ 1.46	79.12 $\pm$ 0.34	3.77 $\pm$ 0.46	78.06 $\pm$ 0.20	15.25 $\pm$ 3.05	58.47 $\pm$ 0.35
AE ( $\ell_2$ loss)	43.47 $\pm$ 1.18	81.48 $\pm$ 0.20	4.90 $\pm$ 0.67	78.11 $\pm$ 0.19	10.00 $\pm$ 3.49	62.98 $\pm$ 0.40
AE ( $\ell_1$ loss)	44.73 $\pm$ 1.69	81.65 $\pm$ 0.34	5.28 $\pm$ 0.60	78.19 $\pm$ 0.25	14.50 $\pm$ 3.72	63.54 $\pm$ 0.49
SaGroW ( $\ell_2$ loss)	66.33 $\pm$ 1.52	79.47 $\pm$ 0.32	17.84 $\pm$ 1.43	78.06 $\pm$ 0.37	47.50 $\pm$ 4.10	<b>67.40</b> $\pm$ 0.37
SaGroW ( $\ell_1$ loss)	68.97 $\pm$ 1.31	80.17 $\pm$ 0.76	16.98 $\pm$ 1.44	78.27 $\pm$ 0.54	50.00 $\pm$ 2.79	<b>67.69</b> $\pm$ 0.55
Spar-GW ( $\ell_2$ loss)	98.79 $\pm$ 0.16	83.65 $\pm$ 0.22	<b>18.87</b> $\pm$ 0.99	<b>78.92</b> $\pm$ 0.11	<b>54.25</b> $\pm$ 3.17	66.70 $\pm$ 0.46
Spar-GW ( $\ell_1$ loss)	<b>99.00</b> $\pm$ 0.22	<b>84.19</b> $\pm$ 0.33	<b>22.26</b> $\pm$ 1.35	<b>78.49</b> $\pm$ 0.69	<b>62.50</b> $\pm$ 3.81	<b>67.00</b> $\pm$ 0.41

\* The top-3 results of each dataset are in bold, and the best result is in italics.

(S-GWL), 22.41s (LR-GW), 501.06s/530.12s (AE under  $\ell_2$  loss/ $\ell_1$  loss), 196.18s/189.57s (SaGroW under  $\ell_2$  loss/ $\ell_1$  loss), and 82.45s/147.33s (SPAR-GW under  $\ell_2$  loss/ $\ell_1$  loss). Such results indicate that SPAR-GW achieves a decent tradeoff between speed and accuracy.

## 7. Conclusion

We developed a novel importance sparsification strategy, achieving the approximation of GW, FGW, and UGW distances in a unified framework with theoretical convergence guarantees. Experiments show that our SPAR-GW method outperforms state-of-the-art approaches in various tasks and attains a decent accuracy-speed tradeoff.

We plan to further investigate the theoretical properties of the specific proposed sampling probability, and we are also interested in theoretically deriving the optimal sampling probability. The proposed importance sparsification mechanism can also be applied to more complex OT problems, for example, the multi-marginal optimal transport problem. Further methodological and theoretical analyses are left to our future work.

## Supplementary Materials

**Appendix:** contains the importance sparsification algorithm for approximating the fused Gromov-Wasserstein distance; complete proofs of theoretical results; and additional experiments to evaluate the approximation accuracy, time cost, and memory consumption of the proposed method. (appendix.pdf, a pdf file)

**Code:** contains Python code that implements the proposed method and reproduces the numerical results. A readme file is included describing the contents. (code.zip, a zip file)

## Acknowledgments

We appreciate the Editor, Associate Editor, and two anonymous reviewers for their constructive comments that helped improve the work. This research was supported by Public Computing Cloud, Renmin University of China.

## Disclosure Statement

The authors report there are no competing interests to declare.

## Funding

Mengyu Li is supported by the Outstanding Innovative Talents Cultivation Funded Programs 2021 of Renmin University of China. The authors would like to acknowledge the support from National Natural Science Foundation

of China grant no.12101606, No.12001042, No.12271522, and Renmin University of China research fund program for young scholars, and Beijing Institute of Technology research fund program for young scholars.

## ORCID

Mengyu Li  <https://orcid.org/0000-0002-5286-7525>

Jun Yu  <https://orcid.org/0000-0001-6068-8415>

Hongteng Xu  <https://orcid.org/0000-0003-4192-5360>

Cheng Meng  <https://orcid.org/0000-0002-7111-0966>

## References

- Alaux, J., Grave, E., Cuturi, M., and Joulin, A. (2019), "Unsupervised Hyper-Alignment for Multilingual Word Embeddings," in *7th International Conference on Learning Representations*, New Orleans, LA, USA. [1512]
- Alvarez-Melis, D., and Jaakkola, T. (2018), "Gromov-Wasserstein Alignment of Word Embedding Spaces," in *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pp. 1881–1890, Brussels, Belgium. ACL. [1512]
- Benamou, J.-D., Brenier, Y., and Guittet, K. (2002), "The Monge-Kantorovich Mass Transfer and its Computational Fluid Mechanics Formulation," *International Journal for Numerical Methods in Fluids*, 40, 21–30. [1513]
- Blumberg, A. J., Carriere, M., Mandell, M. A., Rabadan, R., and Villar, S. (2020), "MREC: A Fast and Versatile Framework for Aligning and Matching Point Clouds with Applications to Single Cell Molecular Data," arXiv preprint arXiv:2001.01666. [1512]
- 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. [1514]
- Bonneel, N., van de Panne, M., Paris, S., and Heidrich, W. (2011), "Displacement Interpolation Using Lagrangian Mass Transport," *ACM Transactions on Graphics*, 30, 1–12. [1518]
- Brenier, Y. (1997), "A Homogenized Model for Vortex Sheets," *Archive for Rational Mechanics and Analysis*, 138, 319–353. [1513]
- Brogat-Motte, L., Flamary, R., Brouard, C., Rousu, J., and d'Alché Buc, F. (2022), "Learning to Predict Graphs with Fused Gromov-Wasserstein Barycenters," in *International Conference on Machine Learning*, pp. 2321–2335. PMLR. [1512]
- Bunne, C., Alvarez-Melis, D., Krause, A., and Jegelka, S. (2019), "Learning Generative Models Across Incomparable Spaces," in *International Conference on Machine Learning*, pp. 851–861. PMLR. [1512]
- Chapel, L., Alaya, M. Z., and Gasso, G. (2020), "Partial Optimal Transport with Applications on Positive-Unlabeled Learning," in *Advances in Neural Information Processing Systems*, (Vol. 33), 2903–2913. [1512]
- Chizat, L., Peyré, G., Schmitzer, B., and Vialard, F.-X. (2018a), "An Interpolating Distance between Optimal Transport and Fisher-Rao Metrics," *Foundations of Computational Mathematics*, 18, 1–44. [1516]
- (2018b), "Scaling Algorithms for Unbalanced Optimal Transport Problems," *Mathematics of Computation*, 87, 2563–2609. [1516,1517]
- (2018c), "Unbalanced Optimal Transport: Dynamic and Kantorovich Formulations," *Journal of Functional Analysis*, 274, 3090–3123. [1516]

- Chowdhury, S., and Mémoli, F. (2019), “The Gromov-Wasserstein Distance between Networks and Stable Network Invariants,” *Information and Inference: A Journal of the IMA*, 8, 757–787. [1512]
- Chowdhury, S., Miller, D., and Needham, T. (2021), “Quantized Gromov-Wasserstein,” in *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, pp. 811–827, Springer. [1512]
- Chowdhury, S., and Needham, T. (2021), “Generalized Spectral Clustering via Gromov-Wasserstein Learning,” in *Proceedings of the 24th International Conference on Artificial Intelligence and Statistics* (Vol. 130), pp. 712–720. PMLR. [1512]
- Cuturi, M. (2013), “Sinkhorn Distances: Lightspeed Computation of Optimal Transport,” in *Advances in Neural Information Processing Systems* (Vol. 26), pp. 2292–2300. [1512,1513,1514]
- 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*, pp. 10648–10656. IEEE. [1514]
- Ezuz, D., Solomon, J., Kim, V. G., and Ben-Chen, M. (2017), “GWCNN: A Metric Alignment Layer for Deep Shape Analysis,” *Computer Graphics Forum*, 36, 49–57. [1512]
- Feragen, A., Kasenburg, N., Petersen, J., de Bruijne, M., and Borgwardt, K. (2013), “Scalable Kernels for Graphs with Continuous Attributes,” in *Advances in Neural Information Processing Systems* (Vol. 26), pp. 216–224. [1519]
- Fey, M., and Lenssen, J. E. (2019), “Fast Graph Representation Learning with PyTorch Geometric,” in *ICLR Workshop on Representation Learning on Graphs and Manifolds*. [1519]
- Genevay, A., Chizat, L., Bach, F., Cuturi, M., and Peyré, G. (2019), “Sample Complexity of Sinkhorn Divergences,” in *22nd International Conference on Artificial Intelligence and Statistics*, pp. 1574–1583. PMLR. [1514]
- Gong, F., Nie, Y., and Xu, H. (2022), “Gromov-Wasserstein Multi-Modal Alignment and Clustering,” in *Proceedings of the 31st ACM International Conference on Information & Knowledge Management*, pp. 603–613. [1512]
- Hagberg, A. A., Schult, D. A., and Swart, P. J. (2008), “Exploring Network Structure, Dynamics, and Function using NetworkX,” in *Proceedings of the 7th Python in Science Conference*, pp. 11–15, Pasadena, CA, USA. [1518]
- Kantorovich, L. (1942), “On the Transfer of Masses,” *Doklady Akademii Nauk*, 37, 227–229 (in Russian). [1513]
- Kawano, S., and Mason, J. K. (2021), “Classification of Atomic Environments via the Gromov-Wasserstein Distance,” *Computational Materials Science*, 188, 110144. [1516]
- Kerdoncuff, T., Emonet, R., and Sebban, M. (2021), “Sampled Gromov Wasserstein,” *Machine Learning*, 110, 2151–2186. [1513,1516,1517,1518,1520]
- Kriege, N. M., Fey, M., Fissler, D., Mutzel, P., and Weichert, F. (2018), “Recognizing Coneiform Signs Using Graph based Methods,” in *International Workshop on Cost-Sensitive Learning*, pp. 31–44. PMLR. [1519]
- Le, T., Ho, N., and Yamada, M. (2021), “Flow-based Alignment Approaches for Probability Measures in Different Spaces,” in *International Conference on Artificial Intelligence and Statistics*, pp. 3934–3942. PMLR. [1512,1513,1518]
- Li, T., Meng, C., Yu, J., and Xu, H. (2022), “Hilbert Curve Projection Distance for Distribution Comparison,” arXiv preprint arXiv:2205.15059. [1514]
- Liao, Q., Chen, J., Wang, Z., Bai, B., Jin, S., and Wu, H. (2022a), “Fast Sinkhorn I: An  $O(N)$  Algorithm for the Wasserstein-1 Metric,” arXiv preprint arXiv:2202.10042. [1514]
- Liao, Q., Wang, Z., Chen, J., Bai, B., Jin, S., and Wu, H. (2022b), “Fast Sinkhorn II: Collinear Triangular Matrix and Linear Time Accurate Computation of Optimal Transport,” arXiv preprint arXiv:2206.09049. [1514]
- Liero, M., Mielke, A., and Savaré, G. (2016), “Optimal Transport in Competition with Reaction: The Hellinger-Kantorovich Distance and Geodesic Curves,” *SIAM Journal on Mathematical Analysis*, 48, 2869–2911. [1516]
- Lin, T., Ho, N., and Jordan, M. I. (2019), “On the Acceleration of the Sinkhorn and Greenkhorn Algorithms for Optimal Transport,” arXiv preprint arXiv:1906.01437. [1514]
- Liu, J. S. (1996), “Metropolized Independent Sampling with Comparisons to Rejection Sampling and Importance Sampling,” *Statistics and Computing*, 6, 113–119. [1515]
- Liu, J. S. (2008), *Monte Carlo Strategies in Scientific Computing*, New York: Springer. [1515]
- Luo, D., Wang, Y., Yue, A., and Xu, H. (2022), “Weakly-Supervised Temporal Action Alignment Driven by Unbalanced Spectral Fused Gromov-Wasserstein Distance,” in *Proceedings of the 30th ACM International Conference on Multimedia*, pp. 728–739. [1516]
- Ma, P., Mahoney, M. W., and Yu, B. (2015), “A Statistical Perspective on Algorithmic Leveraging,” *The Journal of Machine Learning Research*, 16, 861–911. [1516]
- Mémoli, F. (2011), “Gromov-Wasserstein Distances and the Metric Approach to Object Matching,” *Foundations of Computational Mathematics*, 11, 417–487. [1512,1514]
- Meng, C., Ke, Y., Zhang, J., Zhang, M., Zhong, W., and Ma, P. (2019), “Large-Scale Optimal Transport Map Estimation Using Projection Pursuit,” in *Advances in Neural Information Processing Systems* (Vol. 32), pp. 8118–8129. [1514]
- Muzellec, B., Josse, J., Boyer, C., and Cuturi, M. (2020), “Missing Data Imputation Using Optimal Transport,” in *International Conference on Machine Learning*, pp. 7130–7140. PMLR. [1518]
- Nadjahi, K. (2021), “Sliced-Wasserstein Distance for Large-Scale Machine Learning: Theory, Methodology and Extensions,” PhD thesis, Institut polytechnique de Paris. [1514]
- Neumann, M., Moreno, P., Antanas, L., Garnett, R., and Kersting, K. (2013), “Graph Kernels for Object Category Prediction in Task-Dependent Robot Grasping,” in *Online Proceedings of the Eleventh Workshop on Mining and Learning with Graphs*, pp. 10–6, Chicago, Illinois, USA. ACM. [1519]
- Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., Vanderplas, J., Passos, A., Cournapeau, D., Brucher, M., Perrot, M., and Duchesnay, E. (2011), “Scikit-Learn: Machine Learning in Python,” *The Journal of Machine Learning Research*, 12, 2825–2830. [1518]
- Peyré, G., and Cuturi, M. (2019), “Computational Optimal Transport: With Applications to Data Science,” *Foundations and Trends® in Machine Learning*, 11, 355–607. [1513]
- Peyré, G., Cuturi, M., and Solomon, J. (2016), “Gromov-Wasserstein Averaging of Kernel and Distance Matrices,” in *International Conference on Machine Learning*, pp. 2664–2672. PMLR. [1512,1513,1514,1517]
- Pham, K., Le, K., Ho, N., Pham, T., and Bui, H. (2020), “On Unbalanced Optimal Transport: An Analysis of Sinkhorn Algorithm,” in *International Conference on Machine Learning*, pp. 7673–7682. PMLR. [1517]
- Rand, W. M. (1971), “Objective Criteria for the Evaluation of Clustering Methods,” *Journal of the American Statistical Association*, 66, 846–850. [1519]
- Reddi, S. J., Sra, S., Póczos, B., and Smola, A. (2016), “Stochastic Frank-Wolfe Methods for Nonconvex Optimization,” in *2016 54th Annual Allerton Conference on Communication, Control, and Computing (Allerton)*, pp. 1244–1251, IEEE. [1516]
- Sato, R., Cuturi, M., Yamada, M., and Kashima, H. (2020), “Fast and Robust Comparison of Probability Measures in Heterogeneous Spaces,” arXiv preprint arXiv:2002.01615. [1512,1513,1518]
- Scetbon, M., and Cuturi, M. (2020), “Linear Time Sinkhorn Divergences Using Positive Features,” *Advances in Neural Information Processing Systems* (Vol. 33), pp. 13468–13480. [1514]
- Scetbon, M., Peyré, G., and Cuturi, M. (2022), “Linear-Time Gromov Wasserstein Distances Using Low Rank Couplings and Costs,” in *International Conference on Machine Learning*, pp. 19347–19365. PMLR. [1512,1513,1518]
- Séjourné, T., Vialard, F.-X., and Peyré, G. (2021), “The Unbalanced Gromov Wasserstein Distance: Conic Formulation and Relaxation,” in *Advances in Neural Information Processing Systems* (Vol. 34), pp. 8766–8779. [1516,1517,1518]
- Sinkhorn, R., and Knopp, P. (1967), “Concerning Nonnegative Matrices and Doubly Stochastic Matrices,” *Pacific Journal of Mathematics*, 21, 343–348. [1512,1513,1514]
- Solomon, J., Peyré, G., Kim, V. G., and Sra, S. (2016), “Entropic Metric Alignment for Correspondence Problems,” *ACM Transactions on Graphics*, 35, 1–13. [1512]

- Sturm, K.-T. (2006), “On the Geometry of Metric Measure Spaces,” *Acta Mathematica*, 196, 65–131. [1512]
- Sutherland, J. J., O’Brien, L. A., and Weaver, D. F. (2003), “Spline-fitting with a Genetic Algorithm: A Method for Developing Classification Structure-activity Relationships,” *Journal of Chemical Information and Computer Sciences*, 43, 1906–1915. [1519]
- Titouan, V., Courty, N., Tavenard, R., and Flamary, R. (2019a), “Optimal Transport for Structured Data with Application on Graphs,” in *International Conference on Machine Learning*, pp. 6275–6284. PMLR. [1512,1514,1516,1519]
- Titouan, V., Flamary, R., Courty, N., Tavenard, R., and Chapel, L. (2019b), “Sliced Gromov-Wasserstein,” in *Advances in Neural Information Processing Systems* (Vol. 32), pp. 4753–14763. [1512,1513,1518]
- Vayer, T., Chapel, L., Flamary, R., Tavenard, R., and Courty, N. (2020), “Fused Gromov-Wasserstein Distance for Structured Objects,” *Algorithms*, 13, 212–244. [1516]
- Villani, C. (2009), *Optimal Transport: Old and New* (Vol. 338), Berlin: Springer. [1513]
- Vincent-Cuaz, C., Flamary, R., Corneli, M., Vayer, T., and Courty, N. (2022), “Semi-relaxed Gromov-Wasserstein Divergence with Applications on Graphs,” in *10th International Conference on Learning Representations*. [1512]
- Xie, Y., Wang, X., Wang, R., and Zha, H. (2020), “A Fast Proximal Point Method for Computing Exact Wasserstein Distance,” in *Uncertainty in Artificial Intelligence*, pp. 433–453. PMLR. [1517,1519]
- Xu, H., Liu, J., Luo, D., and Carin, L. (2023), “Representing Graphs via Gromov-Wasserstein Factorization,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45, 999–1016. [1512]
- Xu, H., Luo, D., and Carin, L. (2019), “Scalable Gromov-Wasserstein Learning for Graph Partitioning and Matching,” in *Advances in Neural Information Processing Systems* (Vol. 32), pp. 3052–3062. [1512,1513,1518]
- Xu, H., Luo, D., Zha, H., and Carin, L. (2019), “Gromov-Wasserstein Learning for Graph Matching and Node Embedding,” in *International Conference on Machine Learning*, pp. 6932–6941. PMLR. [1512,1514,1517,1518,1519]
- Yan, Y., Li, W., Wu, H., Min, H., Tan, M., and Wu, Q. (2018), “Semi-supervised Optimal Transport for Heterogeneous Domain Adaptation,” in *Proceedings of the 27th International Joint Conference on Artificial Intelligence*, pp. 2969–2975. [1512]
- Yanardag, P., and Vishwanathan, S. (2015), “Deep Graph Kernels,” in *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pp. 1365–1374. [1519]
- Yu, J., Wang, H., Ai, M., and Zhang, H. (2022), “Optimal Distributed Subsampling for Maximum Quasi-Likelihood Estimators with Massive Data,” *Journal of the American Statistical Association*, 117, 265–276. [1516]
- Zhang, J., Ma, P., Zhong, W., and Meng, C. (2022), “Projection-Based Techniques for High-Dimensional Optimal Transport Problems,” *Wiley Interdisciplinary Reviews: Computational Statistics*, e1587. [1514]