
Orthogonal Gromov-Wasserstein Discrepancy with Efficient Lower Bound

Hongwei Jin¹

Zishun Yu¹

Xinhua Zhang¹

¹Computer Science Department, University of Illinois at Chicago, Chicago, Illinois, USA

Abstract

Comparing structured data from possibly different metric-measure spaces is a fundamental task in machine learning, with applications in, e.g., graph classification. The Gromov-Wasserstein (GW) discrepancy formulates a coupling between the structured data based on optimal transportation, tackling the incomparability between different structures by aligning the intra-relational geometries. Although efficient *local* solvers such as conditional gradient and Sinkhorn are available, the inherent non-convexity still prevents a tractable evaluation, and the existing lower bounds are not tight enough for practical use. To address this issue, we take inspirations from the connection with the quadratic assignment problem, and propose the orthogonal Gromov-Wasserstein (OGW) discrepancy as a surrogate of GW. It admits an efficient and *closed-form* lower bound with $\mathcal{O}(n^3)$ complexity, and directly extends to the fused Gromov-Wasserstein (FGW) distance, incorporating node features into the coupling. Extensive experiments on both the synthetic and real-world datasets show the tightness of our lower bounds, and both OGW and its lower bounds efficiently deliver accurate predictions and satisfactory barycenters for graph sets.

1 INTRODUCTION

Similarity based learning has been a popular approach in many machine learning applications. Instead of directly modeling each individual object which may pose challenge in some areas, it resorts to models based on pairwise similarity, possibly across different domains. The most common example is the kernels used in support vector machines and Gaussian process, including RBF kernel and covariance matrices that measure similarity in the Euclidean space [Chang

et al., 2010], and string or tree kernels that compare discrete objects [Lodhi et al., 2002]. More generally, featured graphs have been a useful tool for capturing similarities and relations in the structured data that are commonly not Euclidean. Examples include social network [Fan et al., 2019], recommendation systems [Wu et al., 2020], fraud detection [Li et al., 2020], quantum chemistry [Coley et al., 2019] and topology-aware IoT applications [Abusnaina et al., 2019].

Despite the inherent challenge in comparing graphs from possibly different metric-measure spaces, there has been a wealth of refined discrepancy measures between graphs, including kernels [Vishwanathan et al., 2010, Shervashidze et al., 2011] and GCNs based approaches [Bronstein et al., 2017, Defferrard et al., 2016]. Recently, the Gromov-Wasserstein discrepancy [GW, Peyre et al., 2016], which extends the Gromov-Wasserstein distance [Memoli, 2011], has emerged as an effective transportation distance between structured data, alleviating the incomparability issue between different structures by aligning the *intra*-relational geometries. Thanks to its favorable properties such as efficiency and isometry-awareness, GW has been applied to domain adaptation [Yan et al., 2018], word embedding [Alvarez-Melis and Jaakkola, 2018], graph classification [Vayer et al., 2019a], metric alignment [Ezuz et al., 2017], generative modeling [Cohen and Sejdinovic, 2019], and graph matching and node embedding [Xu et al., 2019b,a].

However, different from the standard Wasserstein distance which is a linear program, GW is unfortunately intractable to evaluate. Despite the practical success of non-convex optimization techniques such as conditional gradient method and entropic regularization [Peyre et al., 2016, Gold and Rangarajan, 1996], it remains NP-hard to find the global optimum. Hence, existing practice settles with local solutions, lacking an analyzable guarantee. This significantly challenges the trustworthiness of the GW discrepancy.

Towards a tractable approximation, Memoli [2011] proposed three lower bounds of GW, which cost $\mathcal{O}(n^2)$, $\mathcal{O}(n^4)$ and $\mathcal{O}(n^5)$. They can be useful for branch-and-bound based

global optimization, as well as recent algorithms for certifying the robustness of nonconvex models. However, these lower bounds are rarely used in practice, and as we will show later in Figure 3, even the most expensive lower bound can be quite loose, raising the concerns on their effectiveness.

Instead of developing yet another lower bound or tight approximation of GW, our goal in this paper is to design a *surrogate* of it (namely orthogonal GW, or OGW) such that:

- (a) It retains GW’s desirable properties such as permutation invariance, non-negativity, triangle inequality, and good performance in machine learning tasks such as classification and barycenter. We stress that OGW does not need to be either an upper bound or a lower bound of GW. We are also *not* concerned about the gap between OGW and GW because what matters is the desirable mathematical properties and the performance in learning, instead of how close OGW is to GW.
- (b) It does *not* have to be tractable, but it must admit a *tight and efficient* lower bound – upper bound is easier from local optimization as in GW. The tightness will potentially contribute to global optimization such as certification, a topic that is beyond this paper. Ideally, such a lower bound should also possess the aforementioned good properties of the surrogate itself.

Our inspiration, as unrolled in Section 2, stems from the connection of GW to quadratic assignment problems (QAPs), which was tapped into by sliced GW [Vayer et al., 2019b] and Gromov-Monge problems [Memoli and Needham, 2021]. It paved the way for approximating the set of doubly stochastic matrices (used in GW) by orthonormal matrices under marginal constraints [Rendl and Wolkowicz, 1992, Hadley et al., 1992, Anstreicher and Brixius, 2001]. The resulting problem is well known to admit tight approximate solutions, and accommodates fused GW to account for node features [FGW, Vayer et al., 2019a]. Experiments on classification and barycenter demonstrate the effectiveness of our proposed OGW and its lower bounds.

2 OGW-DISCREPANCY WITH TRACTABLE LOWER BOUNDS

We represent an undirected graph G with n nodes by an adjacency matrix $\mathbf{A} \in \{0, 1\}^{n \times n}$, where $\mathbf{A}_{ii} = 0$, and $\mathbf{A}_{ij} = 1$ if there is an edge between nodes i and j ($i \neq j$), and 0 otherwise. To start with, we consider the discrepancy between two graphs with the same order (*i.e.*, number of nodes). A detailed generalization to any graph orders will be addressed in Section 2.3. Associated with the nodes is a distribution, encoding some prior information about their importance, *e.g.*, the normalized degree of each node [Xu et al., 2019b]. However, many applications lack such natural normalization [Vayer et al., 2019a,b, Peyre et al., 2016].

Therefore, we will stick with a uniform distribution over all nodes and include our extension of non-uniform distribution in Appendix A.5. Letting $\mathbf{1} = (1, \dots, 1)^\top$ whose dimensionality can be implicitly induced from the context, the standard GW distance between graph G and H based on ℓ_2 distance can be formulated as (the square root of)

$$\text{GW}(G, H) := \min_{\mathbf{P} \in \mathcal{E} \cap \mathcal{N}} \sum_{i, j, k, l} [c_G(i, j) - c_H(k, l)]^2 \mathbf{P}_{ik} \mathbf{P}_{jl}, \quad (1)$$

where $\mathcal{E} := \{\mathbf{P} \in \mathbb{R}^{n \times n} : \mathbf{P}\mathbf{1} = \mathbf{P}^\top \mathbf{1} = \mathbf{1}\}$ and $\mathcal{N} := \mathbb{R}_+^{n \times n}$ [Memoli, 2011]. Here $c_G(i, j)$ represents a distance measure between node i and j on G , and common choices include their shortest-path distance, or simply $1 - \mathbf{A}_{ij}$ (the complement of adjacency). When both c_G and c_H are a metric, (1) is a squared metric on isomorphism classes of measurable metric spaces. However, as pointed out by Peyre et al. [2016], c_G and c_H do not have to be restricted to metrics and ℓ_2 can be extended to other asymmetric or non-subadditive losses such as f -divergence. They call it GW discrepancy, which broadens its applicability in machine learning. We will refer to it as GW, without even taking the square root of (1) just like in Peyre et al. [2016].

Remark 1. *It is noteworthy that although the original GW requires $c_G(i, j)$ to be a distance metric [Memoli, 2011], it can be relaxed in (1) where ℓ_2 loss is used. Indeed, $c_G(i, j)$ can also be served by similarities between nodes instead of distance, *e.g.*, by simply flipping the sign of the distance measure. This also opens up the use of non-metric dissimilarity measures such as constrained shortest path [Lozano and Medaglia, 2013].*

Define two n -by- n symmetric matrices \mathbf{C} and \mathbf{D} whose (i, j) -th elements are $\mathbf{C}_{ij} = c_G(i, j)$ and $\mathbf{D}_{ij} = c_H(i, j)$, respectively. For example, the complement of adjacency can be written as $\mathbf{C} = \mathbf{1}\mathbf{1}^\top - \mathbf{A}$. The above GW can be compactly rewritten in the Koopmans-Beckmann form [Koopmans and Beckmann, 1957]:

$$\text{GW}(G, H) = \frac{1}{n^2} \left(\|\mathbf{C}\|_F^2 + \|\mathbf{D}\|_F^2 - 2 \max_{\mathbf{P} \in \mathcal{E} \cap \mathcal{N}} \text{tr}(\mathbf{CPDP}^\top) \right). \quad (2)$$

Here $\|\cdot\|_F$ is the Frobenius norm. Obviously, GW is permutation invariant, nonnegative, and equals 0 when G and H are isomorphic. The major drawback is that the maximization over \mathbf{P} is intractable, although efficient *local* algorithms are available such as conditional gradient [Vayer et al., 2019a] and Sinkhorn [Peyre et al., 2016].

2.1 CONNECTING OGW WITH THE QUADRATIC ASSIGNMENT PROBLEM

Rewriting GW with quadratic optimization over $\mathcal{E} \cap \mathcal{N}$ as in (2) reveals an innate connection to the quadratic assignment



Figure 1: Connection between QAP and GW, OGW

problem (QAP). Noting that by the Birkhoff–von Neumann theorem [Birkhoff, 1946], $\mathcal{E} \cap \mathcal{N}$ is the convex hull of the set of $n \times n$ permutation matrices (denoted as Π). Indeed, the connection with QAP has been used to formulate Gromov–Monge distances [Memoli and Needham, 2021], and to accelerate the evaluation of GW via projection (slicing) to 1-D [Vayer et al., 2019b]. Fortunately, a number of tractable relaxations of QAP are available, many of which are based on the following characterization of Π :

$$\Pi = \mathcal{E} \cap \mathcal{N} \cap \mathcal{O}, \quad (3)$$

where $\mathcal{O} := \{\mathbf{P} \in \mathbb{R}^{n \times n} : \mathbf{P}^\top \mathbf{P} = \mathbf{P} \mathbf{P}^\top = \mathbf{I}\}$. Here \mathbf{I} is the identity matrix. Whenever necessary, we will explicitize the dimensionality of \mathcal{O} by writing \mathcal{O}_n . Interestingly, $\text{tr}(\mathbf{C} \mathbf{P} \mathbf{D}^\top)$ can be maximized exactly by a simple eigen-decomposition if \mathbf{P} is restricted to \mathcal{O} [Umeyama, 1988]. Specifically, assume the eigen-decomposition of \mathbf{C} and \mathbf{D} are $\mathbf{C} = \mathbf{P}_\mathbf{C} \text{diag}(\boldsymbol{\lambda}_\mathbf{C}) \mathbf{P}_\mathbf{C}^\top$ and $\mathbf{D} = \mathbf{P}_\mathbf{D} \text{diag}(\boldsymbol{\lambda}_\mathbf{D}) \mathbf{P}_\mathbf{D}^\top$, respectively, and suppose the eigenvalues in $\boldsymbol{\lambda}_\mathbf{C}$ and $\boldsymbol{\lambda}_\mathbf{D}$ are both arranged in a descending order. Then

$$\mathbf{P}_1 \mathbf{P}_2^\top \in \arg \max_{\mathbf{P} \in \mathcal{O}} \text{tr}(\mathbf{C} \mathbf{P} \mathbf{D}^\top), \quad (4)$$

$$\text{and } \max_{\mathbf{P} \in \mathcal{O}} \text{tr}(\mathbf{C} \mathbf{P} \mathbf{D}^\top) = \boldsymbol{\lambda}_\mathbf{C}^\top \boldsymbol{\lambda}_\mathbf{D}. \quad (5)$$

Based on this result, Hadley et al. [1992] proposed tightening the domain approximation from \mathcal{O} to $\mathcal{O} \cap \mathcal{E}$, which, despite the original intention of approximating inhomogeneous QAPs, happens to be useful in our context too. Compared with $\mathcal{N} \cap \mathcal{E}$, $\mathcal{O} \cap \mathcal{E}$ offers more convenience in constructing upper and lower bounds that are tight and efficient. This substitution leads to our proposed new metric, named as *orthogonal Gromov–Wasserstein (OGW) discrepancy*:

$$\begin{aligned}
\text{OGW}(G, H) &:= \\
&\frac{1}{n^2} \left(\|\mathbf{C}\|_F^2 + \|\mathbf{D}\|_F^2 - 2 \max_{\mathbf{P} \in \mathcal{O} \cap \mathcal{E}} \text{tr}(\mathbf{C} \mathbf{P} \mathbf{D}^\top) \right). \quad (6)
\end{aligned}$$

Figure 1 illustrates how QAP is connected with GW and OGW through the different convex outer approximations of the domain of the permutation matrices.

2.2 UPPER AND LOWER BOUNDS OF OGW

The evaluation of OGW is hindered by the nonconvex objective and the nonconvex domain in the optimization of \mathbf{P} in (6). So it is natural to resort to its lower and upper bounds.

Upper bound of OGW. Obviously, any locally optimal \mathbf{P} in (6) yields an upper bound of OGW. To ease the local

optimization, we first leverage the characterization of $\mathcal{O} \cap \mathcal{E}$ [Hadley et al., 1992]:

$$\mathcal{O} \cap \mathcal{E} = \left\{ \frac{1}{n} \mathbf{1} \mathbf{1}^\top + \mathbf{V} \mathbf{Q} \mathbf{V}^\top : \mathbf{Q} \in \mathcal{O}_{n-1} \right\}, \quad (7)$$

where \mathbf{V} is any $n \times (n-1)$ matrix satisfying $\mathbf{V}^\top \mathbf{1} = \mathbf{0}$ and $\mathbf{V}^\top \mathbf{V} = \mathbf{I}_{n-1}$. An example is given in Appendix A.1. Plugging $\mathbf{P} = \frac{1}{n} \mathbf{1} \mathbf{1}^\top + \mathbf{V} \mathbf{Q} \mathbf{V}^\top$ into the optimization objective in (6) yields

$$\begin{aligned}
\max_{\mathbf{P} \in \mathcal{O} \cap \mathcal{E}} \text{tr}(\mathbf{C} \mathbf{P} \mathbf{D}^\top) &= \\
\frac{1}{n^2} s_{\mathbf{C}} s_{\mathbf{D}} + \max_{\mathbf{Q} \in \mathcal{O}} \underbrace{\{\text{tr}(\hat{\mathbf{C}} \mathbf{Q} \hat{\mathbf{D}} \mathbf{Q}^\top) + \text{tr}(\hat{\mathbf{E}}^\top \mathbf{Q})\}}_{=: \mathcal{Q}(\hat{\mathbf{C}}, \hat{\mathbf{D}}, \hat{\mathbf{E}})} & \quad (8)
\end{aligned}$$

where $\hat{\mathbf{X}} := \mathbf{V}^\top \mathbf{X} \mathbf{V}$ and $s_{\mathbf{X}} := \mathbf{1}^\top \mathbf{X} \mathbf{1}$ for any matrix \mathbf{X} , and $\mathbf{E} := \frac{2}{n} \mathbf{C} \mathbf{1} \mathbf{1}^\top \mathbf{D}$. Since $\mathcal{Q}(\hat{\mathbf{C}}, \hat{\mathbf{D}}, \hat{\mathbf{E}})$ involves both linear and quadratic terms in \mathbf{Q} , no closed-form solution remains available.

Clearly, any locally optimal \mathbf{Q} yields a *lower bound* for \mathcal{Q} (denoted as \mathcal{Q}_{lb}), i.e., an upper bound for OGW (denoted as OGW_{ub}). *Locally* optimizing \mathbf{Q} over \mathcal{O} (a.k.a. Stiefel manifold) has been very well studied [Absil et al., 2009, Wen and Yin, 2013, Arasu and Mohan, 2018], and we adopt a straightforward approach of projected quasi-Newton, noting that the projection of any matrix \mathbf{Q} on \mathcal{O} is simply $\mathbf{U}_\mathbf{Q} \mathbf{V}_\mathbf{Q}^\top$, where the singular value decomposition (SVD) of \mathbf{Q} is $\mathbf{U}_\mathbf{Q} \Lambda_\mathbf{Q} \mathbf{V}_\mathbf{Q}^\top$. With the locally optimal \mathbf{Q} in hand, the locally optimal \mathbf{P} for OGW can be recovered by plugging \mathbf{Q} into the formula in (7).

Similarly to the practice of GW which resorts to locally optimal solutions, we will use OGW_{ub} as a practical “evaluation” of OGW. Whenever there is no confusion (especially in empirical investigation), we will simply refer to the performance of OGW_{ub} as the performance of OGW.

Lower bounds of OGW. The simplest way to lower bound OGW is by relaxing the domain of \mathbf{P} into \mathcal{O} in (6):

$$\begin{aligned}
\text{OGW}_o(G, H) &:= \\
&\frac{1}{n^2} \left(\|\mathbf{C}\|_F^2 + \|\mathbf{D}\|_F^2 - 2 \max_{\mathbf{P} \in \mathcal{O}} \text{tr}(\mathbf{C} \mathbf{P} \mathbf{D}^\top) \right) \quad (9)
\end{aligned}$$

$$= \frac{1}{n^2} \|\boldsymbol{\lambda}_\mathbf{C} - \boldsymbol{\lambda}_\mathbf{D}\|^2, \quad (10)$$

where the last step is by (4). We note in passing that OGW_o embodies a different design principle from heat kernel signature [Sun et al., 2009] and wave kernel signature [Aubry et al., 2011], in that neither of the kernel signatures sort the kernel spectrum.

In practice, we found that completely dropping the constraint \mathcal{E} may lead to over relaxation. To bring back \mathcal{E} , we follow Hadley et al. [1992] and decompose \mathcal{Q} in (8) into

quadratic and linear terms by decoupling their \mathbf{Q} :

$$\begin{aligned} \mathcal{Q}_{ub}(\hat{\mathbf{C}}, \hat{\mathbf{D}}, \hat{\mathbf{E}}) \\ := \max_{\mathbf{Q}_1 \in \mathcal{O}} \text{tr}(\hat{\mathbf{C}}\mathbf{Q}_1\hat{\mathbf{D}}\mathbf{Q}_1^\top) + \max_{\mathbf{Q}_2 \in \mathcal{O}} \text{tr}(\hat{\mathbf{E}}^\top\mathbf{Q}_2). \end{aligned} \quad (11)$$

As a result, we obtain an *upper bound* of \mathcal{Q} (denoted as \mathcal{Q}_{ub}), which produces a *lower bound* of OGW:

$$\begin{aligned} \text{OGW}_{lb}(G, H) := \\ \frac{1}{n^2} \left(\|\mathbf{C}\|_F^2 + \|\mathbf{D}\|_F^2 - 2\mathcal{Q}_{ub}(\hat{\mathbf{C}}, \hat{\mathbf{D}}, \hat{\mathbf{E}}) - \frac{1}{n^2} s_{\mathbf{C}} s_{\mathbf{D}} \right). \end{aligned} \quad (12)$$

\mathcal{Q}_{ub} can be evaluated analytically. First, \mathbf{Q}_1 can be solved by (4). As for \mathbf{Q}_2 , the von Neumann's trace inequality implies that its optimal value is $\mathbf{U}_{\mathbf{E}}\mathbf{V}_{\mathbf{E}}^\top$, where $\mathbf{U}_{\mathbf{E}}\Lambda_{\mathbf{E}}\mathbf{V}_{\mathbf{E}}^\top$ is the SVD of $\hat{\mathbf{E}}$, and the maximum value of $\text{tr}(\hat{\mathbf{E}}^\top\mathbf{Q}_2)$ is $\|\hat{\mathbf{E}}\|_*$, the trace norm of $\hat{\mathbf{E}}$, which is the sum of the singular values of $\hat{\mathbf{E}}$. [Hadley et al. \[1992\]](#) showed that such an upper bound in (11) is often quite tight, which is also observed in our experiments. Indeed, we noticed that the magnitude of $\hat{\mathbf{C}}$ and $\hat{\mathbf{D}}$ in (11) is significantly larger than that of $\hat{\mathbf{E}}$. Therefore, although the optimal \mathbf{Q}_1 and \mathbf{Q}_2 are different, the resulting gap is small.

We next summarize the mathematical properties of OGW and its lower bounds as follows:

Theorem 1. *OGW, OGW_o , and OGW_{lb} are all nonnegative and symmetric. Their square root satisfies the triangle inequality. Their values are 0 if (but not only if) the two graphs are isomorphic.*

The proof is in Appendix A.2. Compared with the requirement of distance metric, OGW and its lower bounds only fall short of the “only if” part of the identity of indiscernibles. To see why “only if” cannot hold, consider OGW_o whose closed form in (9) shows that its value can be 0 as long as \mathbf{C} and \mathbf{D} are similar, *i.e.*, share the same eigenvalues. In general, however, \mathbf{C} and \mathbf{D} are derived from graphs with certain *discrete* properties, leaving permutation the most likely path to similarity.

Remark 2. *The coupling matrix \mathbf{P} in GW provides a useful matching between two sets of nodes. Although the \mathbf{P} in OGW and its lower bounds may contain negative entries, it optimizes over the orthonormal domain, which may still provide useful insights between the two groups of node. For example, invariance to orthogonal transformation is a longstanding pursuit in learning [[Kornblith et al., 2019](#)]. Despite the hardness of exactly optimizing \mathbf{P} for OGW, we can use \mathbf{Q}_1 from (11) to recover \mathbf{P} via the transformation in (7). This is reasonable because $\hat{\mathbf{E}}$ is generally much smaller in magnitude than $\hat{\mathbf{C}}$ and $\hat{\mathbf{D}}$.*

Computational complexity. The analytic solution for OGW_{lb} and OGW_o is achieved by singular value decomposition and eigen decomposition, whose computational

complexity is $\mathcal{O}(n^3)$. $\hat{\mathbf{C}}$, $\hat{\mathbf{D}}$, and $\hat{\mathbf{E}}$ can be computed in $\mathcal{O}(n^3)$ thanks to the structure of \mathbf{V} (see Appendix A.1). It is worth mentioning that [Memoli \[2011\]](#) also derived the lower bounds of GW by solving a set of linear assignment problems, named First Lower Bound (FLB), Second Lower Bound (SLB), and Third Lower Bound (TLB). The following inequalities provide the connections between different lower bounds of GW:

$$\text{GW} \geq \begin{cases} \text{GW}_{tlb} \geq \text{GW}_{flb} \\ \text{GW}_{slb}. \end{cases} \quad (13)$$

And the complexities of FLB, SLB and TLB are $\mathcal{O}(n^2)$, $\mathcal{O}(n^4)$, $\mathcal{O}(n^5)$, respectively. In general, TLB provides the tightest lower bound for GW, and SLB has resemblant performance compared with TLB. This is also observed in our experiments.

2.3 GRAPHS WITH DIFFERENT SIZES

So far, we have been restricting the two graphs to have the same number of nodes, primarily because the orthonormal domain \mathcal{O} only contains square matrices. In order to deal with the non-square matrix, *i.e.*, graphs of different order, we introduce the semi-orthogonal domain

$$\tilde{\mathcal{O}}_{m,n} := \{\mathbf{T} \in \mathbb{R}^{m \times n} : \mathbf{T}^\top \mathbf{T} = \mathbf{I}_n\}, \quad (14)$$

where $m > n$ without loss of generality. That is a domain of “tall” matrices, whose columns are orthonormal. When $m = n$, $\tilde{\mathcal{O}}$ recovers \mathcal{O}_n because $\mathbf{T}^\top \mathbf{T} = \mathbf{I}_n$ is equivalent to $\mathbf{T}\mathbf{T}^\top = \mathbf{I}_n$ for a square matrix \mathbf{T} .

For a non-square matrix $\mathbf{P} \in \mathbb{R}^{m \times n}$ and $m > n$, it is no longer feasible to impose the constraint of $\mathcal{E} := \{\mathbf{P} \in \mathbb{R}^{n \times n} : \mathbf{P}\mathbf{1} = \mathbf{P}^\top \mathbf{1} = \mathbf{1}\}$ because the dimensionality does not match. Instead, it can be generalized into

$$\tilde{\mathcal{E}} := \{\mathbf{P}\mathbf{1}_n = \sqrt{\frac{n}{m}}\mathbf{1}_m, \mathbf{P}^\top \mathbf{1}_m = \sqrt{\frac{m}{n}}\mathbf{1}_n\}. \quad (15)$$

To summarize, we can extend the OGW in (6) to graphs with different orders by replacing the domain of \mathbf{P} , amounting to

$$\text{OGW}(G, H) := \quad (16)$$

$$\frac{1}{m^2} \|\mathbf{C}\|_F^2 + \frac{1}{n^2} \|\mathbf{D}\|_F^2 - \frac{2}{mn} \max_{\mathbf{P} \in \tilde{\mathcal{O}} \cap \tilde{\mathcal{E}}} \text{tr}(\mathbf{C}\mathbf{P}\mathbf{D}\mathbf{P}^\top),$$

where G and H have the graph order of m and n , respectively. We reuse the symbol OGW because the constraints $\tilde{\mathcal{O}}$ and $\tilde{\mathcal{E}}$ recover \mathcal{O} and \mathcal{E} respectively when $m = n$. It is also easy to see that OGW is nonnegative because

$$\begin{aligned} \text{OGW}(G, H) \\ \geq \frac{1}{m^2} \|\mathbf{C}\|_F^2 + \frac{1}{n^2} \|\mathbf{D}\|_F^2 - \frac{2}{mn} \max_{\mathbf{P} \in \tilde{\mathcal{O}}} \text{tr}(\mathbf{C}\mathbf{P}\mathbf{D}\mathbf{P}^\top) \\ = \left\| \frac{1}{m} \boldsymbol{\lambda}_{\mathbf{C}} - \frac{1}{n} \boldsymbol{\lambda}_{\mathbf{D}} \right\|^2 \geq 0. \end{aligned}$$

In the similar spirit to (7), any $\mathbf{P} \in \tilde{\mathcal{O}} \cap \tilde{\mathcal{E}}$ can be reparameterized as follows

Theorem 2. Let $\mathbf{U} \in \mathbb{R}^{m \times (m-1)}$ and $\mathbf{V} \in \mathbb{R}^{n \times (n-1)}$ be arbitrary projection matrices satisfying

$$\mathbf{U}^\top \mathbf{1}_m = \mathbf{0}_{m-1}, \quad \mathbf{V}^\top \mathbf{1}_n = \mathbf{0}_{n-1}, \quad (17)$$

$$\mathbf{U}^\top \mathbf{U} = \mathbf{I}_{m-1}, \quad \mathbf{V}^\top \mathbf{V} = \mathbf{I}_{n-1}. \quad (18)$$

Then

$$\mathbf{P} \in \tilde{\mathcal{O}} \cap \tilde{\mathcal{E}} \iff \mathbf{P} = \frac{1}{\sqrt{mn}} \mathbf{1}_m \mathbf{1}_n^\top + \mathbf{U} \mathbf{Q} \mathbf{V}^\top, \quad (19)$$

where $\mathbf{Q} \in \tilde{\mathcal{O}}_{m-1, n-1}$.

The proof is relegated to Appendix A.4. Letting $\hat{\mathbf{C}} := \mathbf{U}^\top \mathbf{C} \mathbf{U}$ and $\hat{\mathbf{D}} := \mathbf{V}^\top \mathbf{D} \mathbf{V}$, the optimization over \mathbf{P} in (16) turns into a projected QAP (PQAP) with an additional linear term and some constant terms

$$\max_{\mathbf{Q} \in \tilde{\mathcal{O}}_{m-1, n-1}} \text{tr}(\hat{\mathbf{C}} \mathbf{Q} \hat{\mathbf{D}} \mathbf{Q}^\top) + \text{tr}(\hat{\mathbf{E}} \mathbf{Q}^\top) + \text{const}, \quad (20)$$

where $\hat{\mathbf{E}} = \frac{2}{\sqrt{mn}} \mathbf{U}^\top \mathbf{C} \mathbf{1}_m \mathbf{1}_n^\top \mathbf{D} \mathbf{V}$. Finally, in order to leverage the favorable properties of orthonormal matrices, we right-pad the matrix \mathbf{Q} by an $(m-1) \times (n-m)$ matrix \mathbf{J} , such that $[\mathbf{Q}, \mathbf{J}] \in \mathcal{O}_{m-1}$. Indeed

$$\mathbf{Q} \in \tilde{\mathcal{O}}_{(m-1), (n-1)} \iff \exists \mathbf{J} : [\mathbf{Q} \quad \mathbf{J}] \in \mathcal{O}_{m-1}. \quad (21)$$

Such a \mathbf{J} matrix only needs to be any basis of the kernel space of \mathbf{Q} . As a result, the quadratic term $\max_{\mathbf{Q} \in \tilde{\mathcal{O}}} \text{tr}(\hat{\mathbf{C}} \mathbf{Q} \hat{\mathbf{D}} \mathbf{Q}^\top)$ is equivalent to

$$\max_{[\mathbf{Q} \quad \mathbf{J}] \in \mathcal{O}} \text{tr} \left(\hat{\mathbf{C}} [\mathbf{Q} \quad \mathbf{J}] \begin{bmatrix} \hat{\mathbf{D}} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} \end{bmatrix} [\mathbf{Q} \quad \mathbf{J}]^\top \right), \quad (22)$$

and we can optimize $[\mathbf{Q} \quad \mathbf{J}]$ as a whole, utilizing the closed-form solution to the squared matrix case.

Similarly, for the linear term, we have

$$\max_{\mathbf{Q} \in \tilde{\mathcal{O}}} \text{tr}(\hat{\mathbf{E}} \mathbf{Q}^\top) = \max_{[\mathbf{Q} \quad \mathbf{K}] \in \mathcal{O}} \text{tr} \left([\hat{\mathbf{E}} \quad \mathbf{0}] [\mathbf{Q} \quad \mathbf{K}]^\top \right), \quad (23)$$

where \mathbf{K} serves the same role as \mathbf{J} .

To conclude, by padding zero on the non-square matrices $\hat{\mathbf{D}}$ and $\hat{\mathbf{E}}$ to square matrices, we can enjoy the analytic solutions to the problems in (22) and (23) in the same way as in the square case.

Remark 3. We refrain from the interpretation of adding dummy nodes with 0 distance because, as pointed out in Remark 1, \mathbf{C} and \mathbf{D} can represent similarity measures. In such cases, padding with 0 is still justified with the above derivation, but not amenable to dummy node interpretations.

Remark 4. Disconnected graphs can be modeled by any existing heuristic that is also required by GW. In (1), $c_G(i, j) \neq \infty$ because it would push GW to ∞ as long as all $c_H(k, l) < \infty$ and all nodes have nonzero marginals. A

simple heuristic is to employ a large distance value between two nodes that belong to two separate/disconnected subgraphs. Our experiments only involved connected graphs, because all the graphs from the real datasets are already connected – none was discarded.

We stress that the heuristic in Remark 3 is independent of padding 0 in (22) that tackles different graph sizes. The latter is on how to characterize the alignment of two matrices, which is orthogonal to the design of base measure itself within each graph.

3 BARYCENTER

We next study the application of OGW to the Barycenter problem, where given a set of sampled graphs and their associated weights, we aim to find their Fréchet mean by minimizing the weighted average discrepancy between the barycenter and the sampled graphs. Here the discrepancy is measured by our proposed OGW, and the sampled graphs are represented by a set of cost matrices \mathbf{D}_i , along with normalized weight λ_i for $i \in \{1, \dots, S\}$. \mathbf{D}_i can also encode pairwise similarities instead of dissimilarities, and our method accommodates both cases naturally. The barycenter problem can be formalized as minimizing the weighted average OGW discrepancy:

$$\min_{\mathbf{C}} \mathcal{B}(\mathbf{C}) := \min \sum_{i=1}^S \lambda_i \cdot \text{OGW}(\mathbf{C}, \mathbf{D}_i) \quad (24)$$

$$= \sum_{i=1}^S \lambda_i \left(\frac{1}{m^2} \|\mathbf{C}\|_F^2 - \frac{2}{mn_i} \max_{\mathbf{P}_i \in \mathcal{O} \cap \mathcal{E}} \text{tr}(\mathbf{C} \mathbf{P}_i \mathbf{D}_i \mathbf{P}_i^\top) \right) + \text{const.} \quad (25)$$

For simplicity, we specify the barycenter with a fixed order m , although the value of m can also be optimized.

We follow the block coordinate update proposed in [Peyre et al. \[2016\]](#), i.e., iteratively minimizing with respect to the couplings \mathbf{P}_i and updating the optimal cost matrix for the barycenter in a closed-form solution. Given a set of coupling matrices \mathbf{P}_i , we can retrieve the optimal \mathbf{C} by taking the partial derivative of (24)

$$\frac{\partial \mathcal{B}(\mathbf{C})}{\partial \mathbf{C}} = \sum_i \lambda_i \left(\frac{2}{m^2} \mathbf{C} - \frac{2}{mn_i} \mathbf{P}_i \mathbf{D}_i \mathbf{P}_i^\top \right) \quad (26)$$

$$= \frac{2}{m^2} \mathbf{C} - \sum_i \frac{2\lambda_i}{mn_i} \mathbf{P}_i \mathbf{D}_i \mathbf{P}_i^\top = \mathbf{0}. \quad (27)$$

So the optimal \mathbf{C} under the current $\{\mathbf{P}_i\}$ is $\mathbf{C}^* \leftarrow m \sum_{i=1}^S \frac{\lambda_i}{n_i} \mathbf{P}_i \mathbf{D}_i \mathbf{P}_i^\top$. Noting that OGW itself is still intractable, we resort to the tractable lower or upper bounds, replacing OGW in the definition of $\mathcal{B}(\mathbf{C})$ with its tractable bounds. In the case of OGW_{ub} (resp. OGW_o), we simply adopt the locally (resp. globally) optimal $\{\mathbf{P}_i\}$. For OGW_{lb} ,

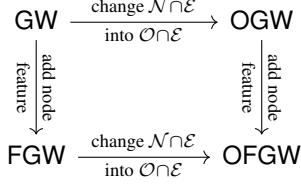


Figure 2: Steps of constructing the graph discrepancies

we can rewrite $\mathcal{B}(\mathbf{C})$ in terms of \mathbf{Q}_1 and \mathbf{Q}_2 from (11) and (12), and then \mathbf{C} can be updated using the optimal \mathbf{Q}_1 and \mathbf{Q}_2 . More details are provided in Appendix A.7.

It is noteworthy that any optimal solution \mathbf{C} for (24) leads to a set of optimal solutions $\{\mathbf{P}\mathbf{C}\mathbf{P}^\top : \mathbf{P} \in \mathcal{O} \cap \mathcal{E}\}$. This also resonates with the intuition retrievable from OGW_o in (9), where only the eigenvalues of \mathbf{C} matter. Therefore, additional ‘‘post-processing’’ is needed to pinpoint the optimal \mathbf{C} from the equivalent class. Furthermore, the optimal \mathbf{C}^* found over $\mathcal{O} \cap \mathcal{E}$ does not guarantee elementwise nonnegativity. Next, we present our method, named as *spectral reconstruction*, to find the appropriate \mathbf{C} .

3.1 SPECTRAL RECONSTRUCTION

To begin with, suppose \mathbf{C} and all \mathbf{D}_i are graphs of order m , and we consider their projections to $\mathbb{R}^{(m-1) \times (m-1)}$ via $\hat{\mathbf{C}}^* = \mathbf{V}^\top \mathbf{C}^* \mathbf{V}$ and $\hat{\mathbf{D}}_i = \mathbf{V}^\top \mathbf{D}_i \mathbf{V}$. Let their eigen-decomposition be $\hat{\mathbf{C}}^* = \mathbf{R}^\top \Sigma \mathbf{R}$ and $\hat{\mathbf{D}}_i = \mathbf{S}_i^\top \Delta_i \mathbf{S}_i$. Inspired by the above observation that $\text{OGW}(\mathbf{C}, \mathbf{D}_i)$ depends primarily on the eigenvalues of \mathbf{C} and \mathbf{D}_i , we rebuild $\hat{\mathbf{C}}^*$ by $\hat{\mathbf{C}}_{\text{recon}}^* = \sum_i \lambda_i \mathbf{S}_i^\top \Sigma \mathbf{S}_i$, *i.e.*, trusting and retaining the eigenvalues of the optimal solution $\hat{\mathbf{C}}^*$ while pairing them with the eigenvectors of the sampled graphs.

For graphs with different sizes, the trick in Section 2.3, *i.e.*, padding on smaller graph, helps us to assemble the $\hat{\mathbf{C}}^*$ from the top $m - 1$ eigen system. Noting that $\hat{\mathbf{C}}^*$ is still on the projected domain $\mathbb{R}^{m-1 \times m-1}$, *i.e.*, there exist \mathbf{V} such that $\hat{\mathbf{C}}^* = \mathbf{V}^\top \mathbf{C}^* \mathbf{V}$.

Next, we bring $\hat{\mathbf{C}}_{\text{recon}}^*$ back to its original space $\mathbb{R}^{m \times m}$ via

$$\mathbf{C}_{\text{recon}}^* = \mathbf{V} \hat{\mathbf{C}}_{\text{recon}}^* \mathbf{V}^\top + \mathbf{Y}, \quad (28)$$

where \mathbf{Y} satisfies $\mathbf{V}^\top \mathbf{Y} \mathbf{V} = 0$, ensuring that $\mathbf{V}^\top \mathbf{C}_{\text{recon}}^* \mathbf{V}$ recovers $\hat{\mathbf{C}}_{\text{recon}}^*$. In addition, when OGW operates on dissimilarity matrices, we require $\text{diag}(\mathbf{C}_{\text{recon}}^*) = \mathbf{0}$, *i.e.*, the dissimilarity between a node and itself is 0. A straightforward choice of \mathbf{Y} satisfying the two conditions is

$$\mathbf{Y} = -\frac{1}{2}(\mathbf{d} \mathbf{1}^\top + \mathbf{1} \mathbf{d}^\top), \quad (29)$$

$$\text{where } \mathbf{d} := \text{diag}(\mathbf{V} \hat{\mathbf{C}}_{\text{recon}}^* \mathbf{V}^\top). \quad (30)$$

To gain more intuition into the recipe, consider the barycenter problem with OGW_{lb} and only one sampled graph \mathbf{D}_1 .

By Theorem 1, $\mathcal{B}(\mathbf{C})$ can be driven to 0 and it is attained when $\hat{\mathbf{C}}$ shares the same eigenvalues as $\hat{\mathbf{D}}_1$, *i.e.*, $\Sigma = \Delta_1$. Then $\hat{\mathbf{C}}_{\text{recon}}^* = \hat{\mathbf{D}}_1$ by our construction. Furthermore, the proof of Theorem 1 indicates $\mathbf{1}^\top \mathbf{C} \mathbf{1} = \mathbf{1}^\top \mathbf{D}_1 \mathbf{1}$ and $\|\mathbf{V}^\top \mathbf{C} \mathbf{1}\| = \|\mathbf{V}^\top \mathbf{D}_1 \mathbf{1}\|$. It is then not hard to show that $\mathbf{C}_{\text{recon}}^*$ is exactly \mathbf{D}_1 . We provide our experiments on both synthetic and real dataset in Section 5.3.

4 EXTENSION TO FUSED GW

Most applications carry features for each node. To account for this important information, Vayer et al. [2019a] proposed the fused GW (FGW), employing an additional matrix \mathbf{M} whose (i, k) -th entry encodes the ℓ_2 distance between the features of node i in G and of node k in H . \mathbf{M} is asymmetric in general. Then the vanilla FGW-distance was formulated by Vayer et al. [2019a] as

$$\begin{aligned} \text{FGW}(G, H, \mathbf{M}) &:= \frac{\alpha}{n^2} \|\mathbf{C}\|_F^2 + \frac{\alpha}{n^2} \|\mathbf{D}\|_F^2 \\ &\quad - \frac{1}{n^2} \max_{\mathbf{P} \in \mathcal{E} \cap \mathcal{N}} \{2\alpha \text{tr}(\mathbf{C} \mathbf{P} \mathbf{D} \mathbf{P}^\top) - (1 - \alpha) \text{tr}(\mathbf{M}^\top \mathbf{P})\}, \end{aligned} \quad (31)$$

where $\alpha \in [0, 1]$ is a trade-off between structure and feature measure. For simplicity, we will only present the treatment for two graphs of the same size. The extension to different sizes can be easily derived in the same way as in Section 2.3. Similar to \mathcal{Q} , with the additional linear term $\text{tr}(\mathbf{M}^\top \mathbf{P})$, there is no closed-form solution even if we replace the domain of \mathbf{P} by \mathcal{O} . However, we can still tighten the domain from \mathcal{O} to $\mathcal{O} \cap \mathcal{E}$, leading to our new approximation

$$\begin{aligned} \text{OFGW}(G, H, \mathbf{M}) &:= \frac{\alpha}{n^2} \|\mathbf{C}\|_F^2 + \frac{\alpha}{n^2} \|\mathbf{D}\|_F^2 \\ &\quad - \frac{1}{n^2} \max_{\mathbf{P} \in \mathcal{E} \cap \mathcal{O}} \{2\alpha \text{tr}(\mathbf{C} \mathbf{P} \mathbf{D} \mathbf{P}^\top) - (1 - \alpha) \text{tr}(\mathbf{M}^\top \mathbf{P})\}. \end{aligned} \quad (32)$$

The pipeline of construction is illustrated in Figure 2. As α tends to zero, OFGW recovers OGW between only structures. A number of favorable properties are enjoyed by OFGW, which are summarized in Theorem 3 below (proof deferred to Appendix A.3). Although OGW must be nonnegative, OFGW is not guaranteed nonnegative for all \mathbf{M} . To see a counter-example, set $\mathbf{C} = \mathbf{D} = \mathbf{0}$ and $\mathbf{M} = \mathbf{I}$. Fortunately, for a large set of \mathbf{M} , it still enjoys nonnegativity.

Theorem 3. *Suppose \mathbf{M} satisfies $\mathbf{1}^\top \mathbf{M} \mathbf{1} \geq n \|\mathbf{M}\|_*$. Then $\text{OFGW}(G, H, \mathbf{M}) \geq 0$ for all G, H , and is invariant to the (different) permutations of G and H . When $\mathbf{M} = \mathbf{0}$, it degenerates to OGW.*

Since \mathbf{M} encodes the pairwise distance between two sets of node features $\{\mathbf{X}_i\}_{i=1}^n$ and $\{\mathbf{Y}_i\}_{i=1}^n$ with $\mathbf{M}_{ij} = \|\mathbf{X}_i - \mathbf{Y}_j\|^2$, we can confirm whether the above assumption holds a priori. Interestingly, this is the case in all the datasets considered in our experiment. For datasets without node attributes and labels, we take node degree as their features. In the sequel, we will make this assumption on \mathbf{M} .

Although OFGW is still intractable in general, local optimization can be performed very efficiently. As will be shown in Section 4.1, it also admits a tight *lower bound* using the same relaxation technique as for OGW. Now that OFGW is motivated by computational convenience, one may naturally wonder whether it captures as much graph structure as the original FGW does. We verified this in the affirmative by following Vayer et al. [2019a], where FGW is used as a kernel function served in a support vector machine. In Section 5.1, we will show that replacing FGW by OFGW achieves similar or better classification accuracy on a variety of datasets, corroborating the effectiveness of OFGW.

4.1 UPPER AND LOWER BOUNDS OF OFGW

Following (7), plugging $\mathbf{P} = \frac{1}{n}\mathbf{1}\mathbf{1}^\top + \mathbf{V}\mathbf{Q}\mathbf{V}^\top$ into the optimization objective in (32) yields

$$\begin{aligned} & \max_{\mathbf{P} \in \mathcal{O} \cap \mathcal{E}} \{2\alpha \text{tr}(\mathbf{C}\mathbf{P}\mathbf{D}\mathbf{P}^\top) - (1-\alpha) \text{tr}(\mathbf{M}^\top \mathbf{P})\} \quad (33) \\ &= \frac{1}{2n}s_{\mathbf{F}} - \frac{\alpha}{2n}s_{\mathbf{M}} + \underbrace{\max_{\mathbf{Q} \in \mathcal{O}} 2\alpha \text{tr}(\hat{\mathbf{C}}\mathbf{Q}\hat{\mathbf{D}}\mathbf{Q}^\top) + \text{tr}(\hat{\mathbf{F}}^\top \mathbf{Q})}_{=: \mathcal{Q}(\hat{\mathbf{C}}, \hat{\mathbf{D}}, \hat{\mathbf{F}})} \end{aligned}$$

where $\mathbf{F} := \frac{2\alpha}{n}\mathbf{C}\mathbf{1}\mathbf{1}^\top\mathbf{D} - \frac{1-\alpha}{2\alpha}\mathbf{M}$. It reveals that compared with the expression of GW in (8), the additional linear term \mathbf{M} in the FGW formulation does not change the structure. As a result, we can again derive the *lower bound* of \mathcal{Q} , *i.e.*, *upper bound* of OFGW, by using the projected quasi-Newton as before. And a *lower bound* of OFGW can also be obtained by decoupling the \mathbf{Q} in the two terms of (33). Specifically, by using the \mathcal{Q}_{ub} defined (11) via decoupling into \mathbf{Q}_1 and \mathbf{Q}_2 , we obtain

$$\begin{aligned} \text{OFGW}_{lb}(G, H, \mathbf{M}) &= \frac{\alpha}{n^2} \|\mathbf{C}\|_F^2 + \frac{\alpha}{n^2} \|\mathbf{D}\|_F^2 \\ &- \frac{1}{n^2} \left[2\alpha \mathcal{Q}_{ub}(\hat{\mathbf{C}}, \hat{\mathbf{D}}, \hat{\mathbf{F}}) + \frac{s_{\mathbf{F}}}{n} - \frac{1-\alpha}{2\alpha n} s_{\mathbf{M}} \right]. \quad (34) \end{aligned}$$

5 EXPERIMENTS

We now demonstrate the empirical effectiveness of OGW and OFGW via two applications: graph classification and barycenter problem. We will also illustrate the tightness of our lower bound OGW_{lb} . All the code and data are available at <https://github.com/cshjin/ogw>.

5.1 EFFECTIVENESS OF OGW/OFGW

Recall in Figure 2, $\mathcal{N} \cap \mathcal{E}$ was replaced by $\mathcal{O} \cap \mathcal{E}$ because the latter enjoys tight and efficiently computable upper and lower bounds. So it is important to validate the resulting OGW/OFGW as an equally good measure of comparing two graphs as the vanilla GW/FGW.

Datasets. We experimented on six graph classification datasets: BZR, COX2, MUTAG, PTC-MR, IMDB-Binary, and IMDB-Multi [TUDataset]. Their statistics are given in Appendix A.8. The first four datasets contain a collection of molecules (*e.g.*, chemical compound and ligands), where the vertices represent atoms and edges are chemical bonds. The class label represents a certain property of the molecules, *e.g.*, "mutagenic effect on a specific bacterium" (MUTAG) and carcinogenicity of compounds for male rats (PTC-MR). BZR and COX2 consist collections of ligands for the benzodiazepine receptor and cyclooxygenase-2 inhibitors, respectively. IMDB-Binary and IMDB-Multi are the movie collaboration dataset, where nodes represent actors/actresses who played roles in movies in IMDB, and one edge means two played in the same movie. We group the dataset into three categories according to their feature property: vectorized, discrete, and no features. For the datasets with no features, we take the node degrees as their features.

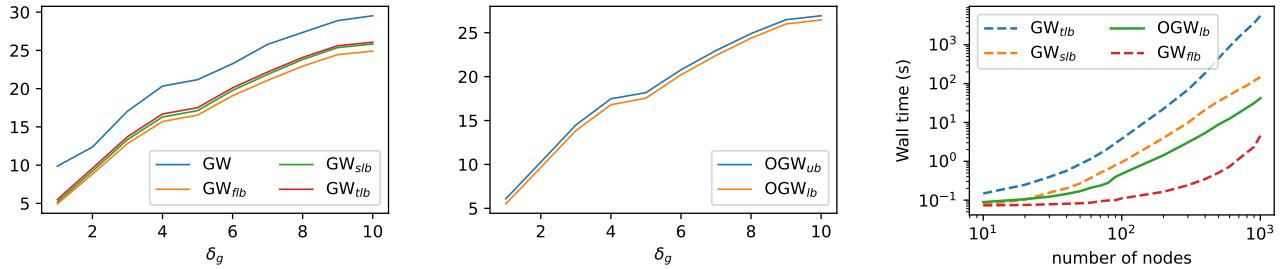
Settings. In order to evaluate a discrepancy measure d , we follow Vayer et al. [2019a] by studying the graph classification accuracy of an SVM, whose kernel $k(G, H)$ is computed by $\exp(-\gamma d(G, H))$. For both FGW and OFGW, the feature distance matrix \mathbf{M} employed the squared Euclidean distance. Since OFGW itself is intractable to evaluate, we resort to the lower bound of OFGW in (31) which has an analytic form and is nonnegative. For the vanilla GW/FGW, we adopt the implementation from POT package [Flamary et al., 2021], which initiates the transition matrix by the outer product of marginal distributions. And we instantiate the cost matrix c_G by all-pair shortest path for each graph in the datasets, knowing that the structures are all connected. We evaluate the models by cross-validation on the hyperparameters in SVM, setting γ from $\{2^{-10}, 2^{-9}, \dots, 2^{10}\}$ and C from 2^{-4} to 2^4 on evenly log scale with 15 steps. Moreover, for the FGW/OFGW, we cross-validate the value of α from $[0, 1]$ with grid search.

In addition, we consider the graph kernel methods as the baselines. More specifically, we adopt the implementation of shortest path (SP) kernel [Borgwardt and Kriegel, 2005] and graphlet sampling (GK) kernel [Przulj, 2007] from Sigidis et al. [2020]. SP kernel decomposes graphs into shortest paths and compares pairs of shortest paths according to their lengths and the labels of their endpoint, while GK kernel decomposes graphs into graphlets, *i.e.*, small subgraphs with k nodes where $k \in \{3, 4, 5, \dots\}$, and counts matching graphlets in the input graphs.

Results. The average accuracy achieved with 10-fold cross-validation is presented in Table 1. To see the effectiveness of OGW and its lower bounds (OGW_{lb} and OGW_o), we compare their accuracy against that of the vanilla GW and its first lower bound (GW_{flb}). In addition, we also include FGW and OFGW that incorporate node features. As the table shows, the tractable lower bounds of our proposed

Table 1: Graph classification

	Dataset	Graph kernel		GW-based SVM			OGW-based SVM			
		SP	GK ($k = 5$)	GW	GW _{flb}	FGW	OGW _{ub}	OGW _{lb}	OGW _o	OFGW _{lb}
Vec.	BZR	78.8 ± 3.3	78.8 ± 3.3	84.9 ± 1.8	78.8 ± 1.0	84.8 ± 3.2	78.8 ± 1.0	83.4 ± 3.4	83.4 ± 3.4	84.3 ± 3.8
Attr.	COX2	78.2 ± 0.4	78.2 ± 0.4	76.2 ± 2.1	78.8 ± 2.2	78.5 ± 1.9	78.4 ± 1.8	78.1 ± 1.8	78.2 ± 0.8	80.2 ± 2.4
Disc.	MUTAG	78.2 ± 4.1	66.5 ± 0.9	85.1 ± 3.4	60.5 ± 2.3	85.7 ± 2.4	66.5 ± 2.3	82.8 ± 3.0	82.8 ± 3.0	85.4 ± 1.7
Attr.	PTC-MR	57.3 ± 1.0	55.8 ± 0.7	53.4 ± 4.3	60.2 ± 5.1	51.8 ± 3.4	59.5 ± 10.1	57.9 ± 4.5	57.9 ± 4.5	57.1 ± 4.1
No	IMDB-B	57.5 ± 2.6	60.1 ± 2.4	63.4 ± 0.9	63.7 ± 4.0	65.6 ± 1.8	65.1 ± 0.3	68.3 ± 1.7	67.4 ± 1.1	67.3 ± 2.1
Attr.	IMDB-M	39.7 ± 1.8	38.2 ± 2.7	47.5 ± 2.3	43.2 ± 2.6	49.7 ± 1.7	48.1 ± 2.2	48.5 ± 1.9	47.9 ± 1.2	47.1 ± 2.3



(a) Average value of GW and its lower bounds under δ_g number of perturbations (b) Average value of OGW_{ub} and OGW_{lb} under δ_g number of perturbations (c) Running time of the lower bounds

Figure 3: Average value of GW and OGW as compared with their respective lower bounds to demonstrate the tightness of the lower bounds. The running time of these lower bounds is also compared.

metric OGW have comparable performance with GW and its variant. With the additional information from node features (except IMDB datasets), the OFGW provides higher accuracy in general. It corroborates that the chain in Figure 2 preserves important structures in graphs, making the tractable lower bounds of OGW/OFGW sound discrepancy measures for graphs. Moreover, OGW_{lb} achieves higher accuracy than OGW_o on two datasets, and performs similarly on the other four datasets.

5.2 TIGHTNESS OF THE LOWER BOUND

The tightness of our lower bound OGW_{lb}, with respect to OGW, can be evaluated through its difference to the upper bound OGW_{ub}, which is obtained by the projected quasi-Newton method. Such a gap will be compared with its counterpart in GW distance, where the upper bound is served by a local optimizer based on the conditional gradient method, and the lower bounds are proposed by Memoli [2011], including FLB, SLB, and TLB (the best known lower bound).

Synthetic data. To demonstrate the tightness in synthetic data, we generate a path graph with 20 nodes and randomly perturb $\delta_g \in [1, 10]$ edges for 50 times, so that we can measure the distance (dissimilarity) between the original graph and the perturbed graph under different measures. Only connected graphs are kept. Figure 3a and 3b provide, respectively for GW and OGW, the average distance as a

function of the number of perturbed edges. The gap between OGW_{ub} and OGW_{lb} is much tighter than the best gap for the GW case, *i.e.*, GW – GW_{tlb}.

Note that TLB of GW requires $\mathcal{O}(n^5)$ computational time, while OGW_{lb} only costs $\mathcal{O}(n^3)$. To verify it, we measure the running time by varying the graph sizes. For each graph size that ranges from 10 to 1000, 20 Erdős-Rényi random graphs are generated, ensuring they are connected. Then their average running time is reported in Figure 3c, which clearly matches the analyzed computational complexity.

Real-world data. We also examine the gap between lower and upper bounds on the real-world dataset MUTAG by evaluating pairwise distance. Figure 4a provides the distribution of the gaps from GW and OGW. For GW, we report the gap between GW local optimizer and GW_{tlb}. Clearly, the gap between OGW_{ub} and OGW_{lb} centers around 0.1, while that between the GW and its TLB concentrates around 1.

In addition, we also compare in Figure 4b the tightness of OGW_{lb} and OGW_o, both as a lower bound for OGW. Each point in the scatter plot represents a graph in the MUTAG dataset, and the horizontal (resp. vertical) axis is the gap between the upper bound of OGW and OGW_{lb} (resp. OGW_o). The fact that the vast majority of the points lie above the diagonal confirms the superior tightness achieved by OGW_{lb}. This demonstrates the benefit of employing \mathcal{E} in the constraint and separating \mathbf{Q}_1 and \mathbf{Q}_2 in (11). More tightness results on other datasets are provided in Appendix A.9.

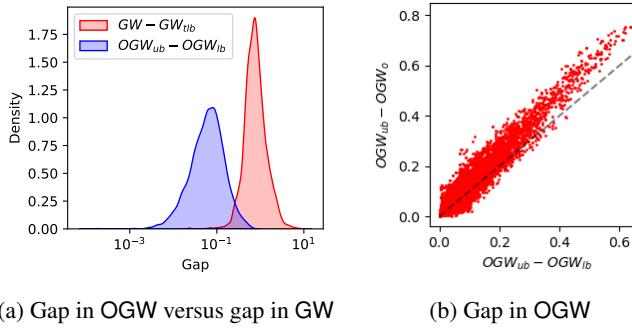


Figure 4: Tightness of lower and upper bounds for GW and OGW (MUTAG dataset)

5.3 BARYCENTER

In the last set of experiments, we evaluate the ability of OGW_{lb} and OGW_{ub} to solve the barycenter problem.

Synthetic data. To start with, we generate a set of cycle-like graphs with different sizes ranging from 15 to 25. In addition, we also explicitly add random structural noise to the graphs separately, and pre-compute their shortest path as cost matrices. All synthetic samples are plotted in Appendix A.10. To initialize the barycenter, we fix the number of nodes to be 20 and initiate a random symmetric C when starting the block coordinate descend to update P . For GW , we adopt the implementation from [Peyre et al. \[2016\]](#) to find the optimal cost matrix of the center. For the OGW_{lb} and OGW_{ub} , we solve the problem by our proposed eigen projection method in Section 3.

To better visualize the results, we also reconstruct the adjacency matrix following a standard heuristic [\[Vayer et al., 2019a\]](#). In particular, for a given threshold, a pair of nodes can be connected by an edge if and only if their corresponding entry in the cost matrix is below the threshold. Then we perform a line search on the threshold to minimize the difference between the optimal cost matrix and the one corresponding to the threshold based adjacency matrix. From Figure 5, we can see the reconstructed cost matrix for the barycenter from our OGW_{lb} and OGW_{ub} are well aligned to its node ordering. We also recognize that the result from GW is just one of the local minimizers in the context of permutation. Due to the tightness between the lower and upper bounds, it is also hard to differentiate the structures found by OGW_{lb} and OGW_{ub} (i.e., sub-figures e and f).

Point cloud data. We also explore the barycenter on point cloud dataset called MNIST-2D ¹. Refer to Appendix A.10 for the plots of the corresponding point cloud data. Clearly, the point clouds for 6 and 9 are quite similar up to rotation.

¹Point cloud MNIST-2D dataset: <https://www.kaggle.com/cristiangarcia/pointcloudmnist2d>

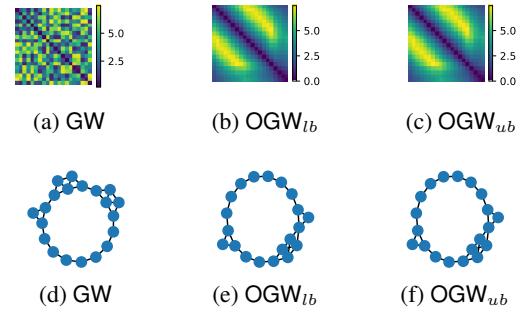


Figure 5: Optimized barycenter. Subfigures (a-c) show the C^* retrieved by optimizing the barycenter. Subfigures (d-f) show the reconstructed graph from the optimal cost matrix C^* .

Table 2: Barycenters from point cloud MNIST-2D samples

	GW	OGW_{lb}	OGW_{ub}
$y = 0$			
$y = 6$			
$y = 9$			

The point cloud is modeled by a graph whose nodes correspond to non-zero (non-black) pixels, represented by their 2D coordinates. Without constructing the explicit mesh connections between pixels, we take the Euclidean distance as their cost matrix D_i . After retrieving the optimal C^* of the barycenter, we further uncover the associated optimal coordinates. A BFGS optimizer was used to seek the locally optimal coordinates such that the resulting Euclidean distance between pixels best reconstructs C^* . Clearly, such a recovery can only be up to the standard invariant transformations such as rotation and shift, and they are determined by the random initialization of BFGS.

Table 2 illustrates the optimal point clouds with labels to be 0, 6, and 9. We sample 5 different point clouds for each digit and set the weight λ_i to the uniform distribution. Moreover, we fix the number of nodes on the barycenter as the minimum size from the samples. Compared with GW , OGW_{lb} and OGW_{ub} clearly find better point clouds of the barycenter in the cases of digit 6 and 9. Moreover, thanks to the tight gap between the lower bounds, the performance of OGW_{lb} and OGW_{ub} differ only indistinguishably.

Acknowledgements

We thank the reviewers for their constructive comments. This work is supported by NSF grant RI:1910146.

References

P.-A. Absil, R. Mahony, and Rodolphe Sepulchre. *Optimization Algorithms on Matrix Manifolds*. Princeton University Press, 2009.

Ahmed Abusnaina, Aminollah Khormali, Hisham Alasmari, Jeman Park, Afsah Anwar, and Aziz Mohaisen. Adversarial learning attacks on graph-based iot malware detection systems. In *ICDCS*, pages 1296–1305. IEEE, 2019.

David Alvarez-Melis and Tommi Jaakkola. Gromov-Wasserstein alignment of word embedding spaces. In *Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2018.

Kurt M. Anstreicher and Nathan W. Brixius. A new bound for the quadratic assignment problem based on convex quadratic programming. *Mathematical Programming*, 89:341–357, 2001.

K. T. Arasu and Manil T. Mohan. Optimization problems with orthogonal matrix constraints. *Numerical Algebra, Control & Optimization*, 8(4):413–440, 2018.

Mathieu Aubry, Ulrich Schlickewei, and Daniel Cremers. The wave kernel signature: A quantum mechanical approach to shape analysis. In *2011 IEEE International Conference on Computer Vision Workshops (ICCV Workshops)*, pages 1626–1633, 2011.

Garrett Birkhoff. Three observations on linear algebra. *Univ. Nac. Tacuman, Rev. Ser. A*, 5:147–151, 1946.

Karsten M. Borgwardt and Hans-Peter Kriegel. Shortest-path kernels on graphs. In *Proc. Intl. Conf. Data Mining*, pages 74–81, 2005.

Michael M. Bronstein, Joan Bruna, Yann LeCun, Arthur Szlam, and Pierre Vandergheynst. Geometric deep learning: Going beyond euclidean data. *IEEE Signal Processing Magazine*, 34(4):18–42, 2017.

Yin-Wen Chang, Cho-Jui Hsieh, Kai-Wei Chang, Michael Ringgaard, and Chih-Jen Lin. Training and testing low-degree polynomial data mappings via linear svm. *Journal of Machine Learning Research*, 11(4), 2010.

Samuel Cohen and Dino Sejdinovic. On the gromov-wasserstein distance and coupled deep generative models. In *NeurIPS 2019 Workshop on Optimal Transport & Machine Learning*, 2019.

Connor W Coley, Wengong Jin, Luke Rogers, Timothy F Jamison, Tommi S Jaakkola, William H Green, Regina Barzilay, and Klavs F Jensen. A graph-convolutional neural network model for the prediction of chemical reactivity. *Chemical science*, 10(2):370–377, 2019.

Michaël Defferrard, Xavier Bresson, and Pierre Vandergheynst. Convolutional neural networks on graphs with fast localized spectral filtering. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2016.

Danielle Ezuz, Justin Solomon, Vladimir G. Kim, and Mirela Ben-Chen. Gwcnn: A metric alignment layer for deep shape analysis. *Comput. Graph. Forum*, 36(5):49–57, 2017.

Wenqi Fan, Yao Ma, Qing Li, Yuan He, Eric Zhao, Jiliang Tang, and Dawei Yin. Graph neural networks for social recommendation. In *The World Wide Web Conference*, pages 417–426, 2019.

Rémi Flamary, Nicolas Courty, Alexandre Gramfort, Mokhtar Z. Alaya, Aurélie Boisbunon, Stanislas Chambon, Laetitia Chapel, Adrien Corenflos, Kilian Fatras, Nemo Fournier, Léo Gautheron, Nathalie T.H. Gayraud, Hicham Janati, Alain Rakotomamonjy, Ievgen Redko, Antoine Rolet, Antony Schutz, Vivien Seguy, Danica J. Sutherland, Romain Tavenard, Alexander Tong, and Titouan Vayer. Pot: Python optimal transport. *Journal of Machine Learning Research*, 22(78):1–8, 2021. URL <http://jmlr.org/papers/v22/20-451.html>.

Steven Gold and Anand Rangarajan. A graduated assignment algorithm for graph matching. *IEEE Transactions on pattern analysis and machine intelligence*, 18(4):377–388, 1996.

S. W. Hadley, F. Rendl, and H. Wolkowicz. A new lower bound via projection for the quadratic assignment problem. *Mathematics of Operations Research*, 17(3):727–739, 1992.

T. C. Koopmans and M. Beckmann. Assignment problems and the location of economic activities. In *Econometrica, Journal of the Econometric Society*, 1957.

Simon Kornblith, Mohammad Norouzi, Honglak Lee, and Geoffrey Hinton. Similarity of neural network representations revisited. In *International Conference on Machine Learning (ICML)*, pages 3519–3529, 2019.

Xiangfeng Li, Shenghua Liu, Zifeng Li, Xiaotian Han, Chuan Shi, Bryan Hooi, He Huang, and Xueqi Cheng. Flowscope: Spotting money laundering based on graphs. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 4731–4738, 2020.

H. Lodhi, C. Saunders, J. Shawe-Taylor, N. Cristianini, and C. Watkins. Text classification using string kernels. *Journal of Machine Learning Research*, 2:419–444, February 2002.

Leonardo Lozano and Andrés L. Medaglia. On an exact method for the constrained shortest path problem. *Computers & Operations Research*, 40(1):378–384, 2013.

Facundo Memoli. Gromov-Wasserstein distances and the metric approach to object matching. *Foundations of Computational Mathematics*, 11:417–487, 2011.

Facundo Memoli and Tom Needham. Distance distributions and inverse problems for metric measure spaces (former title: Gromov-Monge quasi-metrics and distance distributions). *arXiv:1810.09646v3*, 2021.

Gabriel Peyre, Marco Cuturi, and Justin Solomon. Gromov-Wasserstein averaging of kernel and distance matrices. In *International Conference on Machine Learning (ICML)*, 2016.

N. Przulj. Biological network comparison using graphlet degree distribution. *Bioinformatics*, 23(2):e177–e183, Jan 2007.

Franz Rendl and Henry Wolkowicz. Applications of parametric programming and eigenvalue maximization to the quadratic assignment problem. *Mathematical Programming: Series A and B*, 53(1–3):63–78, 1992.

Nino Shervashidze, Pascal Schweitzer, Erik Jan van Leeuwen, Kurt Mehlhorn, and Karsten M Borgwardt. Weisfeiler-lehman graph kernels. *Journal of Machine Learning Research*, 12:2539–2561, 2011.

Giannis Siglidis, Giannis Nikolentzos, Stratis Limnios, Christos Giatsidis, Konstantinos Skianis, and Michalis Vazirgiannis. Grakel: A graph kernel library in python. *J. Mach. Learn. Res.*, 21(54):1–5, 2020.

Jian Sun, Maks Ovsjanikov, and Leonidas Guibas. A concise and provably informative multi-scale signature based on heat diffusion. *Computer Graphics Forum*, 28(5):1383–1392, 2009.

TUDataset. Tudataset. <https://chrsmrrs.github.io/datasets/docs/home/>.

Shinji Umeyama. An eigendecomposition approach to weighted graph matching problems. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 10(5):695–703, 1988.

Titouan Vayer, Laetitia Chapel, Remi Flamary, Romain Tavenard, and Nicolas Courty. Optimal transport for structured data with application on graphs. In *International Conference on Machine Learning (ICML)*, 2019a.

Titouan Vayer, Remi Flamary, Romain Tavenard, Laetitia Chapel, and Nicolas Courty. Sliced Gromov-Wasserstein. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2019b.

S. V. N. Vishwanathan, Nicol N. Schraudolph, Imre Risi Kondor, and Karsten M. Borgwardt. Graph kernels. *Journal of Machine Learning Research*, 11(40):1201–1242, 2010.

Zaiwen Wen and Wotao Yin. A feasible method for optimization with orthogonality constraints. *Mathematical Programming*, 142:397–434, 2013.

Shiwen Wu, Fei Sun, Wentao Zhang, and Bin Cui. Graph neural networks in recommender systems: a survey. *arXiv preprint arXiv:2011.02260*, 2020.

Hongteng Xu, Dixin Luo, and Lawrence Carin. Scalable Gromov-Wasserstein learning for graph partitioning and matching. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2019a.

Hongteng Xu, Dixin Luo, Hongyuan Zha, and Lawrence Carin. Gromov-Wasserstein learning for graph matching and node embedding. In *International Conference on Machine Learning (ICML)*, 2019b.

Yuguang Yan, Wen Li, Hanrui Wu, Huaqing Min, Mingkui Tan, and Qingyao Wu. Semi-supervised optimal transport for heterogeneous domain adaptation. In *International Joint Conference on Artificial Intelligence (IJCAI)*, 2018.

A PROOFS AND DETAILED ALGORITHMS

A.1 VALUE OF \mathbf{V}

As proposed by Hadley et al. [1992], let $x = -\frac{1}{n+\sqrt{n}}$ and $y = -\frac{1}{\sqrt{n}}$. Then $\mathbf{V} \in \mathbb{R}^{n \times (n-1)}$ can be assigned as follows. Set the first row to $y\mathbf{1}_{n-1}^\top$, and the remaining $n-1$ rows, as an $(n-1) \times (n-1)$ matrix, to $x\mathbf{1}_{n-1}\mathbf{1}_{n-1}^\top + \mathbf{I}_{n-1}$.

A.2 PROOF OF THEOREM 1

First, we prove the non-negativity of OGW and OGW_{lb} , as that for OGW_o is trivial from (9). Since OGW_{lb} lower bounds OGW , it suffices to prove the non-negativity of the former. Denote $\mathbf{E} = \mathbf{C}\mathbf{1}\mathbf{1}^\top\mathbf{D}$. Then

$$\text{OGW}_{lb}(\mathbf{C}, \mathbf{D}) = \frac{1}{n^2}\|\mathbf{C}\|_F^2 + \frac{1}{n^2}\|\mathbf{D}\|_F^2 - \frac{2}{n^2} \left[\max_{\mathbf{Q}_1} \text{tr}(\hat{\mathbf{C}}\mathbf{Q}_1\hat{\mathbf{D}}\mathbf{Q}_1^\top) + \frac{2}{n} \max_{\mathbf{Q}_2} \text{tr}(\hat{\mathbf{E}}^\top\mathbf{Q}_2) + \frac{1}{n^2}s_{\mathbf{C}}s_{\mathbf{D}} \right]. \quad (35)$$

Let $\mathbf{u} = \mathbf{1}/\sqrt{n} \in \mathbb{R}^n$. Then $\mathbf{V}^\top\mathbf{u} = \mathbf{0}$ and

$$\|\mathbf{C}\|_F^2 = \left\| \begin{pmatrix} \mathbf{V}^\top \\ \mathbf{u}^\top \end{pmatrix} \mathbf{C} (\mathbf{V}, \mathbf{u}) \right\|_F^2 = \left\| \begin{pmatrix} \mathbf{V}^\top \mathbf{C} \mathbf{V} & \mathbf{V}^\top \mathbf{C} \mathbf{u} \\ \mathbf{u}^\top \mathbf{C} \mathbf{V} & \mathbf{u}^\top \mathbf{C} \mathbf{u} \end{pmatrix} \right\|_F^2 = \|\hat{\mathbf{C}}\|_F^2 + 2\|\mathbf{V}^\top \mathbf{C} \mathbf{u}\|^2 + (\mathbf{u}^\top \mathbf{C} \mathbf{u})^2. \quad (36)$$

It is easy to see that

$$\max_{\mathbf{Q}_1 \in \mathcal{O}} \text{tr}(\hat{\mathbf{C}}\mathbf{Q}_1\hat{\mathbf{C}}\mathbf{Q}_1^\top) = \|\hat{\mathbf{C}}\|_F^2, \quad \max_{\mathbf{Q}_2 \in \mathcal{O}} \text{tr}(\hat{\mathbf{E}}\mathbf{Q}_2^\top) = \|\mathbf{V}^\top \mathbf{C} \mathbf{u}\| \|\mathbf{V}^\top \mathbf{D} \mathbf{u}\|. \quad (37)$$

Denoting $\lambda_i(\hat{\mathbf{C}})$ as the i -th largest eigenvalue of $\hat{\mathbf{C}}$, we have

$$\text{OGW}_{lb}(\mathbf{C}, \mathbf{D}) = \frac{1}{n^2} \left(\|\hat{\mathbf{C}}\|_F^2 + 2\|\mathbf{V}^\top \mathbf{C} \mathbf{u}\|^2 + (\mathbf{u}^\top \mathbf{C} \mathbf{u})^2 + \|\hat{\mathbf{D}}\|_F^2 + 2\|\mathbf{V}^\top \mathbf{D} \mathbf{u}\|^2 + (\mathbf{u}^\top \mathbf{D} \mathbf{u})^2 \right) \quad (38)$$

$$- \frac{1}{n^2} \left(2 \sum_i \lambda_i(\hat{\mathbf{C}}) \lambda_i(\hat{\mathbf{D}}) + \frac{4}{n} \|\mathbf{V}^\top \mathbf{C} \mathbf{1}\| \|\mathbf{V}^\top \mathbf{D} \mathbf{1}\| + \frac{2}{n^2} (\mathbf{1}^\top \mathbf{C} \mathbf{1})(\mathbf{1}^\top \mathbf{D} \mathbf{1}) \right). \quad (39)$$

The non-negativity can then be easily derived by summing up the following equalities, which are all implied by $a^2 + b^2 \geq 2ab$:

$$\|\hat{\mathbf{C}}\|_F^2 + \|\hat{\mathbf{D}}\|_F^2 = \sum_i \lambda_i(\hat{\mathbf{C}})^2 + \sum_i \lambda_i(\hat{\mathbf{D}})^2 \geq 2 \sum_i \lambda_i(\hat{\mathbf{C}}) \lambda_i(\hat{\mathbf{D}}), \quad (40)$$

$$\|\mathbf{V}^\top \mathbf{C} \mathbf{u}\|_F^2 + \|\mathbf{V}^\top \mathbf{D} \mathbf{u}\|_F^2 \geq \frac{2}{n} \|\mathbf{V}^\top \mathbf{C} \mathbf{1}\| \|\mathbf{V}^\top \mathbf{D} \mathbf{1}\|, \quad (41)$$

$$(\mathbf{u}^\top \mathbf{C} \mathbf{u})^2 + (\mathbf{u}^\top \mathbf{D} \mathbf{u})^2 \geq \frac{2}{n^2} (\mathbf{1}^\top \mathbf{C} \mathbf{1})(\mathbf{1}^\top \mathbf{D} \mathbf{1}). \quad (42)$$

Next, let us assume the two graphs are isomorphic, *i.e.*, there exists a permutation matrix $\mathbf{\Pi}$ such that $\mathbf{C} = \mathbf{\Pi}^\top \mathbf{D} \mathbf{\Pi}$. For OGW , simply pick $\mathbf{P} = \mathbf{\Pi}^\top$, which is in $\mathcal{O} \cap \mathcal{E}$. Then n^2 times of the right-hand side of (6) equals

$$\|\mathbf{\Pi}^\top \mathbf{D} \mathbf{\Pi}\|_F^2 + \|\mathbf{D}\|_F^2 - 2 \text{tr}(\mathbf{\Pi}^\top \mathbf{D} \mathbf{\Pi} \mathbf{\Pi}^\top \mathbf{D} \mathbf{\Pi}) = \|\mathbf{D}\|_F^2 + \|\mathbf{D}\|_F^2 - 2 \text{tr}(\mathbf{D} \mathbf{D}) = 0. \quad (43)$$

Since $\text{OGW} \geq 0$, 0 must be its value in this case. As OGW_{lb} and OGW_o lower bound OGW and they are nonnegative, they must also be 0.

Finally, we will show the triangle inequality. For OGW , let graphs G, H, J share the same size and have distance matrices $\mathbf{C}, \mathbf{D}, \mathbf{E}$, respectively. Denote $d(\mathbf{C}, \mathbf{D}) = \sqrt{\text{OGW}(G, H)}$ and its corresponding \mathbf{P} in (6) is \mathbf{P}_1 . Likewise, denote $d(\mathbf{D}, \mathbf{E}) = \sqrt{\text{OGW}(H, J)}$ with \mathbf{P}_2 . Then $d(\mathbf{C}, \mathbf{D}) = \|\mathbf{P}_1^\top \mathbf{C} \mathbf{P}_1 - \mathbf{D}\|_F$ and $d(\mathbf{D}, \mathbf{E}) = \|\mathbf{P}_2^\top \mathbf{D} \mathbf{P}_2 - \mathbf{E}\|_F$. Now let $\mathbf{P}_3 = \mathbf{P}_1 \mathbf{P}_2$, which is obviously in $\mathcal{O} \cap \mathcal{E}$. Finally,

$$d(\mathbf{C}, \mathbf{E}) \leq \|\mathbf{P}_3^\top \mathbf{C} \mathbf{P}_3 - \mathbf{E}\|_F = \|\mathbf{P}_1^\top \mathbf{C} \mathbf{P}_1 - \mathbf{P}_2^\top \mathbf{D} \mathbf{P}_2^\top\|_F \quad (44)$$

$$\leq \|\mathbf{P}_1^\top \mathbf{C} \mathbf{P}_1 - \mathbf{D}\|_F + \|\mathbf{P}_2^\top \mathbf{D} \mathbf{P}_2 - \mathbf{E}\|_F = d(\mathbf{C}, \mathbf{D}) + d(\mathbf{D}, \mathbf{E}). \quad (45)$$

For OGW_{lb} , (40) to (42) effectively decompose the values of OGW_{lb} as the sum of “left-hand side minus right-hand side” from the three equations. Therefore, OGW_{lb} equals the squared ℓ_2 distance between the feature representations of the two graphs, where the first graph (and likewise the second graph) is represented as $(\lambda_1(\hat{\mathbf{C}}), \dots, \lambda_{n-1}(\hat{\mathbf{C}}), \|\mathbf{V}^\top \mathbf{C} \mathbf{1}\|/\sqrt{n}, \mathbf{1}^\top \mathbf{C} \mathbf{1}/n)^\top$. This directly leads to the triangle inequality over graphs.

OGW_o given in (9) obviously satisfies the triangle inequality, because it represents each graph with its sorted eigenvalues, and computes their Euclidean distance.

A.3 PROOF OF THEOREM 3

Following the idea of Theorem 1, we prove the non-negativity of OFGW as follows:

$$\text{OFGW}_{lb}(G, H, \mathbf{M}) \quad (46)$$

$$= \frac{1}{n^2} \left[\alpha \|\mathbf{C}\|_F^2 + \alpha \|\mathbf{D}\|_F^2 - 2\alpha \max_{\mathbf{Q}_1 \in \mathcal{O}} \text{tr}(\hat{\mathbf{C}} \mathbf{Q}_1 \hat{\mathbf{D}} \mathbf{Q}_1^\top) - 2 \left\| \alpha \hat{\mathbf{E}} - \frac{1-\alpha}{2\alpha} \hat{\mathbf{M}} \right\|_* - \frac{\alpha}{n} s_{\mathbf{E}} + \frac{1-\alpha}{\alpha n} s_{\mathbf{M}} \right] \quad (47)$$

$$\geq \frac{1}{n^2} \left[\alpha \|\mathbf{C}\|_F^2 + \alpha \|\mathbf{D}\|_F^2 - 2\alpha \max_{\mathbf{Q}_1 \in \mathcal{O}} \text{tr}(\hat{\mathbf{C}} \mathbf{Q}_1 \hat{\mathbf{D}} \mathbf{Q}_1^\top) - 2\alpha \left\| \hat{\mathbf{E}} \right\|_* - \frac{\alpha}{n} s_{\mathbf{E}} - \frac{1-\alpha}{\alpha} \left\| \hat{\mathbf{M}} \right\|_* + \frac{1-\alpha}{\alpha n} s_{\mathbf{M}} \right] \quad (48)$$

$$\geq \frac{1}{n^2} \left[-\frac{1-\alpha}{\alpha} \left\| \hat{\mathbf{M}} \right\|_* + \frac{1-\alpha}{\alpha n} s_{\mathbf{M}} \right] \quad (49)$$

$$\geq 0. \quad (50)$$

Here the first \geq is by the triangle inequality of the trace norm, the second \geq is by $\text{OGW}_{lb}(G, H) \geq 0$ in Theorem 1, and the third is by the assumption on \mathbf{M} in Theorem 3, *i.e.*, $\|\mathbf{M}\|_* \geq \mathbf{1}^\top \mathbf{M} \mathbf{1}$.

When $\mathbf{M} = \mathbf{0}$, $\text{OFGW}(G, H, \mathbf{M}) = 0$ because the optimal \mathbf{P} in the definition of Eq. (32) is identity when $\mathbf{C} = \mathbf{D}$.

A.4 PROOF OF THEOREM 2 REGARDING THE PARAMETRIZATION OF NON-SQUARE MATRICES

Sufficiency. One can easily check that

$$\mathbf{P}^\top \mathbf{P} = \frac{1}{n} \mathbf{1}_n \mathbf{1}_n^\top + \mathbf{V} \mathbf{V}^\top = \mathbf{I}_n, \quad (51)$$

where the last equality is by $\mathbf{V}^\top \mathbf{1} = \mathbf{0}$ and $\mathbf{V}^\top \mathbf{V} = \mathbf{I}_{n-1}$.

Necessity. By Hadley et al. [1992, Lemma 3.1], we can rewrite \mathbf{P} as

$$\mathbf{P} = \mathbf{B}_m \begin{bmatrix} 1 & \mathbf{0} \\ \mathbf{0} & \mathbf{Q} \end{bmatrix} \mathbf{B}_n^\top, \quad (52)$$

where $\mathbf{B}_m = \begin{bmatrix} \frac{\mathbf{1}_m}{\sqrt{m}} & \mathbf{U} \end{bmatrix}$, $\mathbf{B}_n = \begin{bmatrix} \frac{\mathbf{1}_n}{\sqrt{n}} & \mathbf{V} \end{bmatrix}$. As a result,

$$\mathbf{P}^\top \mathbf{P} = [\mathbf{1} \quad \mathbf{V}] \begin{bmatrix} 1 & \mathbf{0}^\top \\ \mathbf{0} & \mathbf{Q}^\top \mathbf{Q} \end{bmatrix} \begin{bmatrix} \mathbf{1}^\top \\ \mathbf{V}^\top \end{bmatrix} = \mathbf{I}_n. \quad (53)$$

A.5 FROM UNIFORM DISTRIBUTION TO NON-UNIFORM DISTRIBUTION

In the setup of OGW , we set, for simplicity, the node distribution as uniform, *i.e.*,

$$\tilde{\mathcal{O}} \cap \tilde{\mathcal{E}} := \left\{ \mathbf{P}^\top \mathbf{P} = \mathbf{I}_n, \mathbf{P} \mathbf{1}_n = \sqrt{\frac{n}{m}} \mathbf{1}_m, \mathbf{P}^\top \mathbf{1}_m = \sqrt{\frac{n}{m}} \mathbf{1}_m, \mathbf{P} \in \mathbb{R}^{m \times n} \right\}. \quad (54)$$

To consider the non-uniform distribution, we can simply replace the uniform distribution with any arbitrary distribution \mathbf{p} and \mathbf{q} by

$$\tilde{\mathcal{O}} \cap \tilde{\mathcal{E}} := \left\{ \mathbf{T}^\top \mathbf{T} = \frac{\mathbf{I}_n}{mn}, \mathbf{T} \mathbf{1}_n = \mathbf{p}, \mathbf{T}^\top \mathbf{1}_m = \mathbf{q}, \mathbf{p} \in \mathbb{R}^m, \mathbf{q} \in \mathbb{R}^n \right\}. \quad (55)$$

Note that the semi-orthogonal domain is replaced by a scalar, without hurting the tractable algorithm to provide bounded results.

A.6 FROM ℓ_2 LOSS TO KL DIVERGENCE

Let us recall the definition GW in Eq. (1), replacing ℓ_2 loss with KL divergence $KL(a|b) = a \log(a/b) - a + b$ will lead the problem to be rewritten in the Koopmans-Beckmann form

$$\text{GW}_{KL}(\mathbf{C}, \mathbf{D}) = \frac{\text{tr}(\mathbf{C} \log(\mathbf{C}) - \mathbf{C})}{m^2} + \frac{\text{tr}(\mathbf{D})}{n^2} - \max_{\mathbf{T} \in \mathcal{E} \cap \mathcal{N}} \text{tr}(\mathbf{C}\mathbf{T} \log(\mathbf{D})\mathbf{T}^\top). \quad (56)$$

Note that we stick with the original domain of \mathcal{E} , *i.e.*, $\mathbf{T}\mathbf{1} = \mathbf{p}$, $\mathbf{T}^\top\mathbf{1} = \mathbf{q}$, where there is no constant scalar in front of the last term. Therefore, consider the partial derivative of GW w.r.t. \mathbf{C} is presented as

$$\frac{\partial \text{GW}_{KL}}{\partial \mathbf{C}} = \frac{\log \mathbf{C}}{m^2} - \mathbf{T}^\top \log(\mathbf{D})\mathbf{T} = \mathbf{0} \implies \mathbf{C}^* = e^{m^2 \mathbf{T}^\top \log(\mathbf{D})\mathbf{T}}, \quad (57)$$

where the optimal \mathbf{C}^* is used in the block coordinate relaxation when solving the barycenter problem. To consider the

$$\text{OGW}_{KL}(\mathbf{C}, \mathbf{D}) = \frac{\text{tr}(\mathbf{C} \log(\mathbf{C}) - \mathbf{C})}{m^2} + \frac{\text{tr}(\mathbf{D})}{n^2} - \frac{1}{mn} \max_{\mathbf{P} \in \mathcal{E} \cap \mathcal{O}} \text{tr}(\mathbf{C}\mathbf{P} \log(\mathbf{D})\mathbf{P}^\top), \quad (58)$$

where $\log(\mathbf{C}), \log(\mathbf{D})$ are evaluated as the inputs, we can still apply the analytical methods in Section 2.2 to achieve OGW_{lb} and OGW_{ub} .

For tractable lower bound / upper bound, Let $\mathbf{P} = \frac{1}{\sqrt{mn}} \mathbf{1}\mathbf{1}^\top + \mathbf{U}\mathbf{Q}\mathbf{V}^\top$, we have

$$\max_{\mathbf{P} \in \mathcal{O} \cap \mathcal{E}} \text{tr}(\mathbf{C}\mathbf{P} \log(\mathbf{D})\mathbf{P}^\top) = \max_{\mathbf{Q} \in \tilde{\mathcal{O}}} \text{tr} \left[\mathbf{C} \left(\frac{1}{\sqrt{mn}} \mathbf{1}\mathbf{1}^\top + \mathbf{U}\mathbf{Q}\mathbf{V}^\top \right) \log \mathbf{D} \left(\frac{1}{\sqrt{mn}} \mathbf{1}\mathbf{1}^\top + \mathbf{U}\mathbf{Q}\mathbf{V}^\top \right)^\top \right]. \quad (59)$$

This is exactly the same as the ℓ_2 loss with the replacement of \mathbf{D} with $\log(\mathbf{D})$. Therefore, the same lower bound and upper bound can be applied to the loss function with KL divergence.

For the barycenter problem, we can take partial derivative of OGW w.r.t. \mathbf{C} ,

$$\frac{\partial \text{OGW}_{KL}}{\partial \mathbf{C}} = \frac{\log \mathbf{C}}{m^2} - \frac{1}{mn} \mathbf{P}^\top \log(\mathbf{D})\mathbf{P} = \mathbf{0} \implies \mathbf{C}^* = e^{\frac{m}{n} \mathbf{P}^\top \log(\mathbf{D})\mathbf{P}}. \quad (60)$$

Consider the OGW_{ub} , we recover the \mathbf{P} from the optimal \mathbf{Q} from \mathcal{Q}_{lb} by joint optimize the quadratic and linear term. And for the OGW_{lb} , consider

$$\mathcal{Q}_{ub} = \max_{\mathbf{Q}_1 \in \mathcal{O}} \text{tr}(\hat{\mathbf{C}}\mathbf{Q}_1 \hat{\mathbf{D}}\mathbf{Q}_1^\top) + \max_{\mathbf{Q}_2 \in \mathcal{O}} \text{tr}(\hat{\mathbf{E}}^\top \mathbf{Q}_2) \quad (61)$$

$$= \max_{\mathbf{Q}_1} \text{tr}(\mathbf{V}^\top \mathbf{C}\mathbf{V}\mathbf{Q}_1 \mathbf{V}^\top \mathbf{D}\mathbf{V}\mathbf{Q}_1^\top) + \max_{\mathbf{Q}_2} \text{tr}(\frac{2}{n} \mathbf{D}^\top \mathbf{1}\mathbf{1}^\top \mathbf{C}\mathbf{Q}_2), \quad (62)$$

the partial derivative of OGW_{lb} w.r.t. \mathbf{C} goes to

$$\begin{aligned} \frac{\partial \text{OGW}_{KL,lb}}{\partial \mathbf{C}} &= \frac{\log \mathbf{C}}{m^2} \\ &- \frac{2}{mn} \left(\frac{\mathbf{1}_m \mathbf{1}_m^\top \log \mathbf{D}}{mn} + \frac{1}{\sqrt{mn}} (\mathbf{1}_m \mathbf{1}_n^\top \log \mathbf{D} \mathbf{V} \mathbf{Q}_2^\top \mathbf{U}^\top + \mathbf{U} \mathbf{Q}_2 \mathbf{V}^\top \log \mathbf{D} \mathbf{1}_n \mathbf{1}_m^\top) + \mathbf{U} \mathbf{Q}_1 \mathbf{V}^\top \log \mathbf{D} \mathbf{V} \mathbf{Q}_1^\top \mathbf{U}^\top \right), \end{aligned} \quad (63)$$

which serves the closed-form solution of \mathbf{C} when applying the block coordinate update in barycenter.

A.7 DERIVATIVE OF BARYCENTER W.R.T. COST MATRIX

Consider the quadratic term in (6), let's take the partial derivative respect to \mathbf{C} ,

$$\frac{\partial}{\partial \mathbf{C}} \text{tr}(\mathbf{C}\mathbf{P}\mathbf{D}\mathbf{P}^\top) = \mathbf{P}\mathbf{D}\mathbf{P}^\top. \quad (64)$$

Noting that the OGW_{ub} solves the \mathcal{Q} jointly by projected quasi-Newton method. Having

$$\mathbf{Q}^* \leftarrow \max_{\mathbf{Q} \in \mathcal{O}} \text{tr}(\hat{\mathbf{C}}\mathbf{Q}\hat{\mathbf{D}}\mathbf{Q}^\top) + \text{tr}(\hat{\mathbf{E}}^\top \mathbf{Q}), \quad (65)$$

Table 3: Statistics of the datasets

dataset	# graphs	# class	# features	ave. edge	min edge	max edge	avg. node	min node	max node
BZR	405	2	56	74.0	26	120	35.0	13	57
COX2	467	2	38	86.0	68	118	41.0	32	56
MUTAG	188	2	7	38	20	66	17.5	10	28
PTC-MR	344	2	18	25.0	2	142	13.0	2	64
IMDB-B	1000	2	-	130	52	2498	17	12	136
IMDB-M	1500	3	-	72	23	2934	10	7	89

we can retrieve the gradient information directly by replacing $\mathbf{P} = \frac{1}{\sqrt{mn}} \mathbf{1}_m \mathbf{1}_n^\top + \mathbf{U} \mathbf{Q}^* \mathbf{V}^\top$.

For the OGW_{lb} , by plug into the $\mathbf{P} \mathbf{D} \mathbf{P}^\top$, we have

$$\mathbf{P} \mathbf{D} \mathbf{P}^\top = \left(\frac{1}{\sqrt{mn}} \mathbf{1}_m \mathbf{1}_n^\top + \mathbf{U} \mathbf{Q} \mathbf{V}^\top \right) \mathbf{D} \left(\frac{1}{\sqrt{mn}} \mathbf{1}_n \mathbf{1}_m^\top + \mathbf{V} \mathbf{Q}^\top \mathbf{U}^\top \right) \quad (66)$$

$$= \frac{1}{mn} \mathbf{1}_m \mathbf{1}_n^\top \mathbf{D} \mathbf{1}_n \mathbf{1}_m^\top + \frac{1}{\sqrt{mn}} \mathbf{1}_m \mathbf{1}_n^\top \mathbf{D} \mathbf{V} \mathbf{Q}^\top \mathbf{U}^\top + \frac{1}{\sqrt{mn}} \mathbf{U} \mathbf{Q} \mathbf{V}^\top \mathbf{D} \mathbf{1}_n \mathbf{1}_m^\top + \mathbf{U} \mathbf{Q} \mathbf{V}^\top \mathbf{D} \mathbf{V} \mathbf{Q}^\top \mathbf{U}^\top \quad (67)$$

$$= \frac{s_D}{mn} \mathbf{1}_m \mathbf{1}_m^\top + \frac{1}{\sqrt{mn}} \left(\mathbf{1}_m \mathbf{1}_n^\top \mathbf{D} \mathbf{V} \mathbf{Q}^\top \mathbf{U}^\top + \mathbf{U} \mathbf{Q} \mathbf{V}^\top \mathbf{D} \mathbf{1}_n \mathbf{1}_m^\top \right) + \mathbf{U} \mathbf{Q} \mathbf{V}^\top \mathbf{D} \mathbf{V} \mathbf{Q}^\top \mathbf{U}^\top. \quad (68)$$

And solving the decoupled two terms will have

$$\mathbf{Q}_1^*, \mathbf{Q}_2^* \leftarrow \max_{\mathbf{Q}_1 \in \mathcal{O}} \text{tr}(\hat{\mathbf{C}} \mathbf{Q}_1 \hat{\mathbf{D}} \mathbf{Q}_1^\top) + \max_{\mathbf{Q}_2 \in \mathcal{O}} \text{tr}(\hat{\mathbf{E}}^\top \mathbf{Q}_2), \quad (69)$$

to update the partial derivative of OGW_{lb} w.r.t. \mathbf{C} accordingly. Afterwards, considering the barycenter problem, we will update the cost matrix of the barycenter by the weighted summation over the dissimilarity between sample and center. To be specific,

$$\frac{\partial \mathcal{B}(\mathbf{C})}{\partial \mathbf{C}} = \sum_i \lambda_i \frac{\partial \text{OGW}(\mathbf{C}, \mathbf{D}_i)}{\partial \mathbf{C}} = \sum_i \lambda_i \left[\frac{2\mathbf{C}}{m^2} - \frac{2}{mn} \frac{\partial \text{tr}(\mathbf{C} \mathbf{P} \mathbf{D}_i \mathbf{P}^\top)}{\partial \mathbf{C}} \right]. \quad (70)$$

Therefore, we can plug the partial derivative of the trace in terms of lower and upper bounds accordingly.

A.8 DATASET STATISTICS

Table 3 provides the statistics of the six datasets used in graph classification experiment.

A.9 SUPPLEMENTARY MEASURES OF TIGHTNESS

Similar to the tightness measures for MUTAG, we also provide the tightness for datasets BZR, COX2, PTC-MR, IMDB-B, IMDB-M from Figure 6 to 10. They confirm our observation that our lower bound of OGW approximates it significantly more tightly than the third lower bound does for GW .

A.10 SAMPLES FOR THE BARYCENTER PROBLEM

Figure 11 are the synthetic noised graphs generated by randomly adding structural noises. And Figure 12 are the real samples from point cloud MNIST dataset, and the rows correspond to labels 0, 6, and 9.

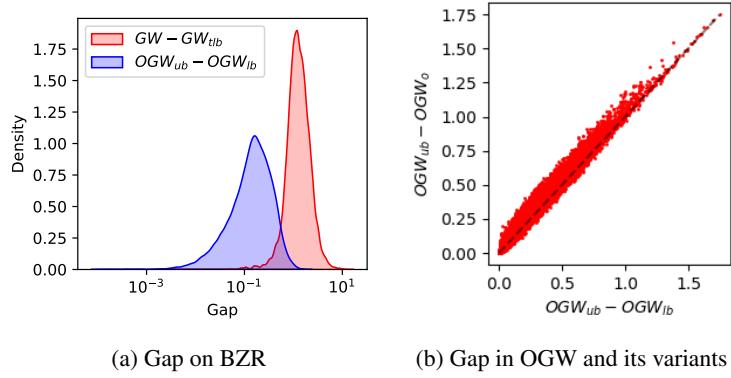


Figure 6: Tightness of lower and upper bounds for GW and OGW (BZR dataset)

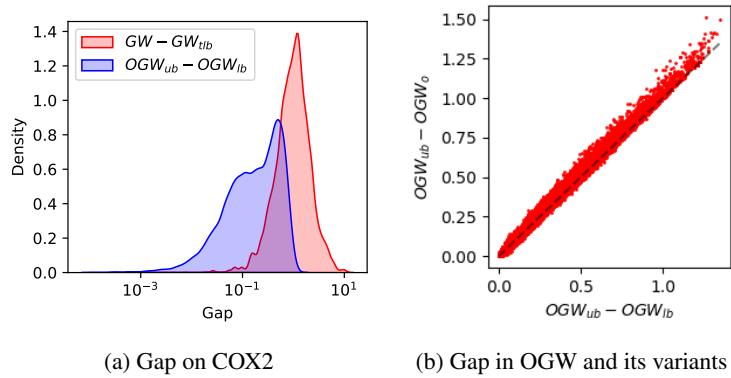


Figure 7: Tightness of lower and upper bounds for GW and OGW (COX2 dataset)

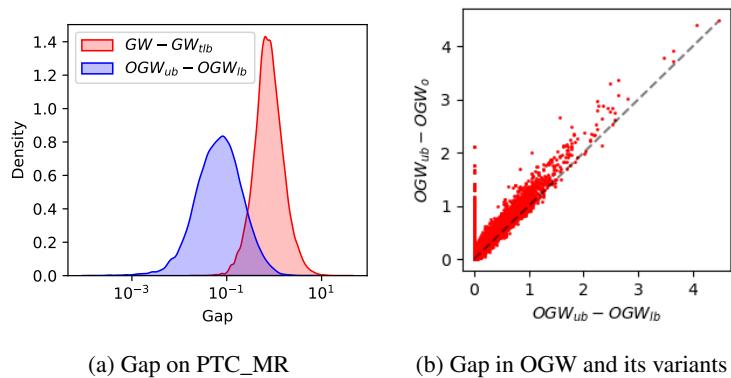


Figure 8: Tightness of lower and upper bounds for GW and OGW (PTC_MR dataset)

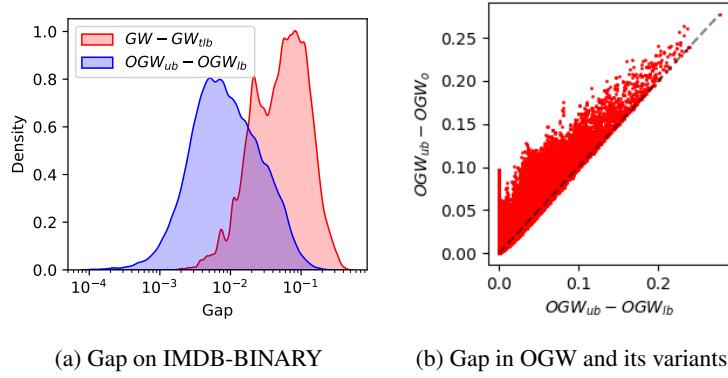


Figure 9: Tightness of lower and upper bounds for GW and OGW (IMDB-BINARY dataset)

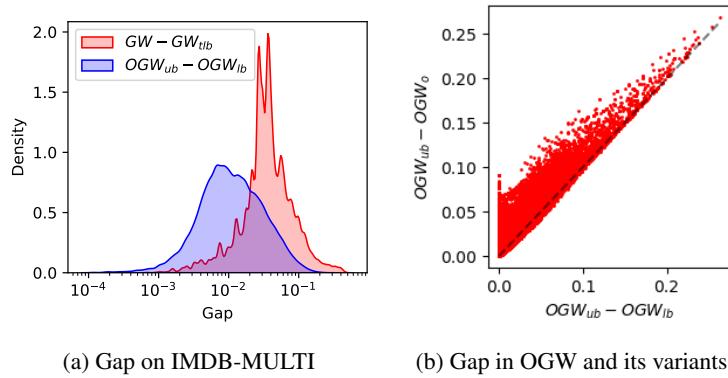


Figure 10: Tightness of lower and upper bounds for GW and OGW (IMDB-MULTI dataset)

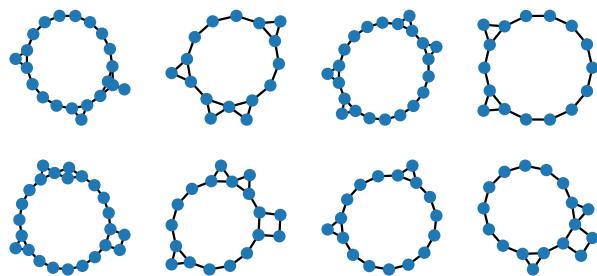


Figure 11: Generated sample cycle-like graphs with structural noises.

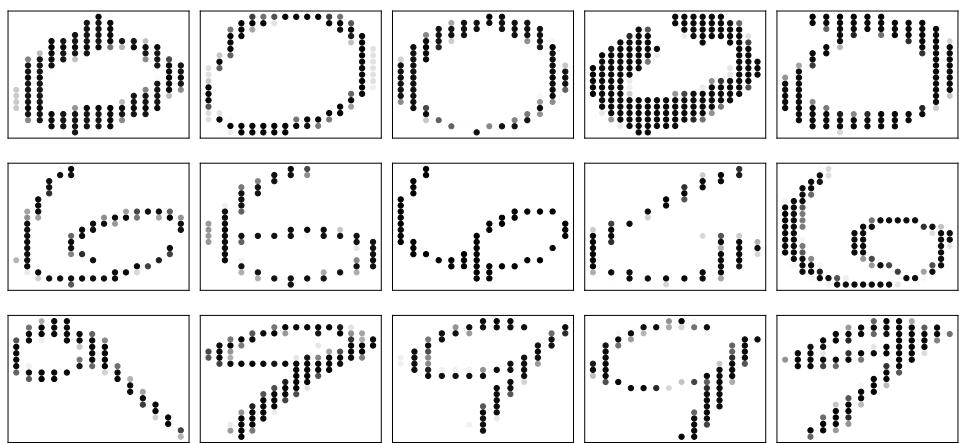


Figure 12: Samples from point cloud MNIST-2D