# QDOT: AN EFFICIENT QUANTILE-WEIGHTED DISTANCE METRIC FOR GEOMETRIC COMPARISON VIA OPTIMAL TRANSPORT

**Anonymous authors**Paper under double-blind review

#### **ABSTRACT**

Measuring the discrepancy between data distributions in heterogeneous metric spaces is a fundamental challenge. Existing methods, typically based on geometric structures, address this by embedding distributions into a shared space. However, these approaches face fundamental limitations, including the loss of geometric information, computationally intractable representations, and inability to preserve essential structural features. In this work, we introduce the Quantileweighted Distance Optimal Transport (ODOT), a novel and efficient metric for geometric comparison. QDOT constructs a family of isometry-invariant distance representations by leveraging distance quantiles as structural weights in Euclidean space, thereby preserving essential geometric characteristics and enabling optimal transport coupling within a common space. We prove that, under mild conditions, QDOT is a well-defined metric with a convergence rate no slower than the classical Wasserstein distance. Moreover, we present an integral version that computes the loss in complexity of  $\mathcal{O}(n \log n)$ . Extensive experiments demonstrate that our methods achieves strong performance across diverse applications, including cross-space comparison, transfer learning, and molecule generation, while also achieving state-of-the-art results on several key metrics.

#### 1 Introduction

The Wasserstein distance is a powerful metric for comparing probability distributions defined on a same metric space. It has found widespread adoption in a diverse range of machine learning tasks, such as generative models(Arjovsky et al., 2017; De Bortoli et al., 2021; Tong et al., 2024), language models(Kusner et al., 2015; Melnyk et al., 2024), multimodal learning(Xu & Chen, 2023; Alatkar & Wang, 2023; Shi et al., 2024) and reinforcement learning(Klink et al., 2022; Asadulaev et al., 2024). However, a notable limitation of the Wasserstein distance is its sensitivity to the spatial separation of the distributions' supports, meaning that two distributions with similar shapes but farapart supports can still have a large Wasserstein distance. Furthermore, the applicability of optimal transport is predicated on a pre-defined ground metric between the supports of the distributions. This requirement renders it intractable for comparing distributions in heterogeneous spaces where such a metric is not readily available.

To address these challenges, metrics based on shape features have emerged (Gromov, 1981; Sturm, 2006; Mémoli, 2011). These methods leverage the distribution of distances within a metric space to compare distributions. Among these, the Gromov-Wasserstein (GW) distance (Sturm, 2006; Mémoli, 2011) stands out as a canonical example. It resolves the issue of comparing distributions in disparate spaces by seeking an optimal coupling of points in a shared, albeit implicitly defined, metric space. However, this formulation is a non-convex optimization problem, the structure of the latent space is not explicitly constructed, and its computation is prohibitively expensive, hindering its use on practical datasets. A related approach, EMD under Transformation Sets (EMD<sup>G</sup>) (Cohen & Guibasm, 1999), extends the GW concept to Euclidean spaces of the same dimension by finding an optimal orthogonal transformation. However, it is not applicable for cross-space comparisons and is susceptible to converging to local optima. Faster approximations, such as the Invariant Sliced Gromov-Wasserstein (RISGW) (Titouan et al., 2019), achieve a favorable  $\mathcal{O}(n \log n)$  complexity, but at the cost of sacrificing theoretical guarantees, such as key metric properties.

Another line of work focuses on extracting shape-invariant features (Belongie et al., 2002; Sun et al., 2019; Yang et al., 2016). These methods map distributions from different spaces into a common feature space for comparison. However, they often suffer from unavoidable information loss due to the unidirectional nature of the feature mapping. To capture richer information, subsequent approaches have employed deep neural networks (Chen et al., 2019; Kim et al., 2020) or transformer based representations (Fuchs et al., 2020; Yu et al., 2023, TBR) for comparison. A key limitation is that these models are typically trained for specific tasks and thus lack generalizability.

To the best of our knowledge, the existing geometric comparison methods can be unified under a common paradigm: they map distributions into a shared space for comparison. However, the mappings employed by these methods often suffer from critical issues: (1) they lack an explicit form; (2) they lead to significant information

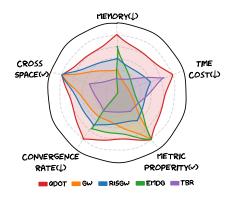


Figure 1: Comparison of various geometric metrics. Lower is better for memory, time cost, and convergence rate. We also evaluate metric properties and cross-space capabilities.

an explicit form; (2) they lead to significant information loss; or (3) they introduce non-geometric artifacts that can corrupt the comparison.

To overcome these challenges, we introduce QDOT, a novel framework that constructs an explicit, geometry-based representation derived from quantile-weighted distances. This principled construction guarantees the integrity of the intrinsic distance information. As illustrated in Figure 1, QDOT strikes a highly effective balance between theoretical properties, generalization capability, time complexity, and computational resource usage, addressing the key trade-offs that limit existing methods. The key contributions of QDOT can be summarized as follows:

**Theoretical Guarantees.** We prove that QDOT and its integral version constitute a metric on the space of isometry classes under certain conditions. Furthermore, we establish that their sample convergence rate is at least as fast as that of the Wasserstein distance.

Computational Efficiency. We propose a highly efficient algorithm. The computation of our representations has a complexity of only  $\mathcal{O}(n\log n)$ . The subsequent comparison step for the full QDOT requires an additional  $\mathcal{O}(n^2\log n)$  for a standard Wasserstein distance calculation. In contrast, its integral variant, IQDOT, leverages the closed-form solution of one-dimensional OT to achieve an overall complexity of just  $\mathcal{O}(n\log n)$ . Consequently, our family of methods offers a highly efficient, quasi-linear time solution.

**Versatility and Strong Performance.** Our method achieved strong performance across a diverse range of experiments. In point cloud comparison tasks, it accurately captured the geometric dissimilarities between distributions and produced high-quality alignments, showcasing its powerful cross-space capabilities. Furthermore, leveraging its computational efficiency, our method achieved excellent results in a large-scale transfer learning scenario. Finally, when integrated as a loss function for a molecular generation model, the QDOT loss significantly enhanced model generalization and achieved state-of-the-art performance on multiple key metrics.

### 2 Preliminary

To establish a rigorous framework, we begin with the fundamental definitions. For any metric space  $(X,d_X)$ , we can define its Borel  $\sigma$ -algebra, denoted  $\mathcal{B}(X)$ . For any probability measure  $\mu_X$  defined on the measurable space  $(X,\mathcal{B}(X))$ , its support, written as  $\mathrm{supp}(\mu_X)$ , is the smallest closed set  $C\subseteq X$  such that  $\mu_X(C)=1$ . For a measurable map  $f:X\to Y$ , the push-forward measure  $f_\#\mu_X$  on  $(Y,\mathcal{B}(Y))$  is defined by  $f_\#\mu_X(A):=\mu_X(f^{-1}(A))$  for any set  $A\in\mathcal{B}(Y)$ . For two measures  $\mu_X$  and  $\mu_Y$  in spaces X and Y, a joint probability  $\pi$  on the product space  $X\times Y$  is a coupling if its marginals satisfy  $(\mathrm{proj}_X)_\#\pi=\mu_X$  and  $(\mathrm{proj}_Y)_\#\pi=\mu_Y$ . When  $\mu_X$  and  $\mu_Y$  are defined on the same metric space  $(\Omega,d)$ , they can be compared using the Wasserstein distance (Villani et al., 2008), which is defined as:

$$\mathcal{W}_p(\mu_X, \mu_Y) := \left( \inf_{\pi \in \Pi(\mu_X, \mu_Y)} \int_{\Omega \times \Omega} d(x, y)^p \, d\pi(x, y) \right)^{\frac{1}{p}}. \tag{1}$$

The primary challenge arises when comparing distributions defined on different metric spaces. To handle this scenario, the measurable metric space(Mémoli, 2011, mm-space) is defined as the triplet  $\mathcal{X}=(X,d_X,\mu_X)$ . For two distinct mm-spaces,  $\mathcal{X}=(X,d_X,\mu_X)$  and  $\mathcal{Y}=(Y,d_Y,\mu_Y)$ , Sturm (2006) proposed the Gromov-Wasserstein distance, extending the concept from equation 1:

$$\mathfrak{S}_p(\mathcal{X}, \mathcal{Y}) := \inf_{\mathcal{Z}} \mathcal{W}_p(f(\mathcal{X}), g(\mathcal{Y})), \tag{2}$$

where the maps  $f: X \to Z$  and  $g: Y \to Z$  are isometric embeddings, satisfying  $d_Z(f(x),f(x'))=d_X(x,x'), \forall x,x'\in \operatorname{supp}(\mu_X)$  and  $d_Z(g(y),g(y'))=d_Y(y,y'), \forall y,y'\in \operatorname{supp}(\mu_Y)$ . Intuitively, the definition in equation 2 can be understood as finding an optimal "lossless" projection of the two disparate spaces X and Y into a common latent space Z, within which their Wasserstein distance can be computed. A significant contribution of the Gromov-Wasserstein distance is that it provides a complete metric on the space of mm-spaces. This metric is formally defined as follows.

**Definition 1 (Metric on Isometry Classes of mm-spaces)** For two mm-spaces  $\mathcal{X} = (X, d_X, \mu_X)$  and  $\mathcal{Y} = (Y, d_Y, \mu_Y)$ , a map  $f : X \to Y$  is called an isometry if it is a surjection satisfying  $d_Y(f(x), f(x')) = d_X(x, x'), \forall x, x' \in \operatorname{supp}(\mu_X)$ . Two mm-spaces are considered isometric if such an isometry exists between them. A function  $\mathcal{L}$  that measures the dissimilarity between two mm-spaces is a metric on the isometry classes of mm-spaces if it satisfies: (1) **Identity of Indiscernibles:**  $\mathcal{L}(\mathcal{X}, \mathcal{Y}) \geq 0$ , and  $\mathcal{L}(\mathcal{X}, \mathcal{Y}) = 0$  iff  $\mathcal{X}$  and  $\mathcal{Y}$  are isometric; (2) **Symmetry:**  $\mathcal{L}(\mathcal{X}, \mathcal{Y}) = \mathcal{L}(\mathcal{Y}, \mathcal{X})$ ; (3) **Triangle Inequality:**  $\mathcal{L}(\mathcal{X}, \mathcal{Y}) \leq \mathcal{L}(\mathcal{X}, \mathcal{Z}) + \mathcal{L}(\mathcal{Y}, \mathcal{Z})$ .

Inspired by the structure of the Gromov-Wasserstein distance in equation 2, we summarized those methods for comparing two mm-spaces can be expressed in a general form:

$$\mathcal{L}_p(\mathcal{X}, \mathcal{Y}) = \mathcal{W}_p(f(\mathcal{X}), g(\mathcal{Y})) \tag{3}$$

where  $f: X \to Z$  and  $g: Y \to Z$  are mapping functions into a common space Z. However, constructing such maps f and g that yield a valid metric satisfying Definition 1 imposes two key requirements. First, the maps must be **isometry-invariant**. That is, if mm-spaces  $\mathcal{X}_1$  and  $\mathcal{X}_2$  are isometric, their representations must be identically distributed, i.e.,  $f_{\#\mu_{X_1}} = f_{\#\mu_{X_2}}$ . Second, to satisfy the identity of indiscernibles property, the representation must be **information-preserving**; it must uniquely encode the metric structure of the original space such that non-isometric spaces map to distinct distributions.

However, the requirement of isometry invariance often entails an unavoidable loss of structural information, and existing methods conforming to the structure of equation 3 fail to meet both requirements simultaneously, thereby falling short of constituting a well-defined metric.

Our work introduces a novel framework designed to explicitly resolve this conflict. To achieve isometry-invariance, we introduce the Quantile-weighted Distance Mean (QDM). Its quantile-based weighting scheme is inherently invariant to isometries, ensuring that any two isometric mm-spaces will produce the same canonical mean representation. To preserve information, , we draw inspiration from the principle of trilateration. As illustrated in Figure 2, this principle dictates that the location of a target point can be uniquely determined from its distances to a sufficient number of known anchor points. We leverage this concept in our framework by treating the QDMs as a set of canonical, isometry-invariant anchors. Consequently, the collection of distances from all points in the support to these QDMs forms a new representation that effectively captures the intrinsic distance information of the entire distribution. In the following sections, we will formally define the QDM and QDOT and detail their theoretical properties.

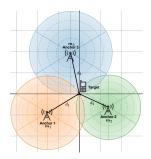


Figure 2: Illustration of trilateration

#### 3 Proposed method

#### 3.1 QUANTILE DISTANCE-WEIGHTED OPTIMAL TRANSPORT

When comparing the metric properties of distributions, their absolute positions in space are typically irrelevant. Therefore, throughout our analysis, we will assume that any mm-space  $\mathcal{X}$  and  $\mathcal{Y}$ 

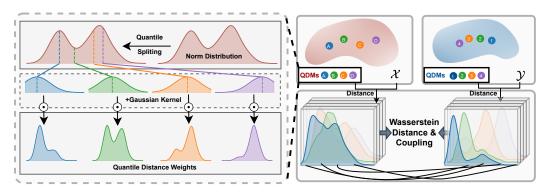


Figure 3: Illustration of the QDOT framework. The method first computes canonical QDM anchors for each distribution, then generates QDMD vector representations based on distances to these anchors. The final QDOT distance is the Wasserstein distance between these two resulting representation distributions.

embedded in a Euclidean space has been centered (i.e., has a zero mean). A key intuition follows from this: if two centered mm-spaces  $\mathcal X$  and  $\mathcal Y$  are isometric, then the distributions of their norms,  $\|\cdot\|_{2\#\mu_X}$  and  $\|\cdot\|_{2\#\mu_X}$ , must be identical.

Motivated by this observation, we use a Gaussian kernel to construct a weighting scheme based on the norm-quantiles, which leads to the definition of the Quantile-weighted Distance Mean:

**Definition 2 (Quantile-weighted Distance Mean (QDM))** Given a centered mm-space  $\mathcal{X} = (X, d_X, \mu_X)$ , where  $X \subseteq \mathbb{R}^p$  and the metric is the standard Euclidean distance,  $d_X(\boldsymbol{x}_1, \boldsymbol{x}_2) = \|\boldsymbol{x}_1 - \boldsymbol{x}_2\|_2$ , we can derive the distribution of its norms, denoted by  $\mu_{\|X\|_2} = \|\cdot\|_{2\#\mu_X}$ . Let  $F_{\|X\|_2}$  denote the cumulative distribution function (CDF) of  $\mu_{\|X\|_2}$ , and let  $F_{\|X\|_2}^{-1}$  be its associated quantile function. For any quantile level  $q \in (0,1)$ , we define the Quantile Distance Weight function  $w: X \times (0,1) \to \mathbb{R}$  as

$$w^{X}(\mathbf{x},q) := e^{-\sigma \left(\|\mathbf{x}\|_{2} - F_{\|X\|_{2}}^{-1}(q)\right)^{2}},$$
(4)

where  $\sigma$  is a bandwidth parameter. Based on these weights, the corresponding QDM  $\mathbf{m}^X:(0,1)\to\mathbb{R}^p$  is defined as the weighted mean:

$$\boldsymbol{m}^{X}(q) := \frac{\mathbb{E}_{\mu_{X}}[w^{X}(X,q)X]}{\mathbb{E}_{\mu_{X}}[w^{X}(X,q)]}.$$
 (5)

Intuitively, for a given quantile level q, the weight function  $w^X(x,q)$  is concentrated on points whose norms are close to the q-th quantile of the norm distribution. Therefore, the QDM can be served as a canonical, isometry-invariant anchor point. We use it to define the Quantile Distance-weighted Mean Distance (QDMD), which is a function  $\phi: X \times (0,1) \to \mathbb{R}$  that measures the distance from a point to the QDM:

$$\phi^X(\boldsymbol{x},q) = d_X(\boldsymbol{x},\boldsymbol{m}^X(q)). \tag{6}$$

Given a quantile level vector  $\mathbf{q} = (q_1, \dots, q_k) \in (0, 1)^k$ . The corresponding QDMD maps each point  $\mathbf{x}$  to a feature vector in  $\mathbb{R}^{k+1}$ :

$$\phi^X(x, \mathbf{q}) := [\phi_0^X(x), \phi^X(x, q_1), \phi^X(x, q_2), \dots, \phi^X(x, q_k)],$$

where  $\phi_0^X(x) = \|x\|_2$  is the original norm, included as a fundamental reference distance to the origin. This mapping transforms the measure  $\mu_X$  into a push-forward measure on  $\mathbb{R}^{k+1}$ . Based on this transformation, we define the Quantile-weighted Distance Optimal Transport as follows:

**Definition 3 (Quantile-weighted Distance Optimal Transport (QDOT))** Let  $\mathcal{X} = (X, d_X, \mu_X)$  and  $\mathcal{Y} = (Y, d_Y, \mu_Y)$  be two mm-spaces, where  $X \subseteq \mathbb{R}^d$  and  $Y \subseteq \mathbb{R}^q$ , and their metrics  $d_X$  and  $d_Y$  are the standard Euclidean distances. For a given vector of quantile levels  $\mathbf{q} \in (0, 1)^k$ , we compute their corresponding QDMD representations,  $\phi^X$  and  $\phi^Y$ . The QDOT distance is then defined as the Wasserstein distance between the resulting push-forward measures:

$$QD_p(\mathcal{X}, \mathcal{Y}) = \mathcal{W}_p(\phi_{\#\mu_X}^X, \phi_{\#\mu_Y}^Y). \tag{7}$$

The overall procedure of the QDOT framework is illustrated in Figure 3.

#### 3.2 Theoretical Results

We will now establish the key theoretical properties of our proposed QDOT framework, demonstrating that it is a well-defined metric with favorable sample convergence guarantees.

**Theorem 1 (Metric Property on Isometry Classes)** Let  $\mathcal{X} = (X, d_X, \mu_X)$  and  $\mathcal{Y} = (Y, d_Y, \mu_Y)$  be two mm-spaces embedded in Euclidean spaces. For a given quantile level vector  $\mathbf{q} \in (0, 1)^k$ , suppose their corresponding QDMs satisfy the dimensionality condition  $\dim(\{\mathbf{m}^X(q_i)\}_{1 \le i \le k}) = \dim(X)$  or  $\dim(\{\mathbf{m}^Y(q_i)\}_{1 \le i \le k}) = \dim(Y)$ . Then, the p-QDOT distance,  $\mathcal{QD}_p$ , defines a metric on the isometry classes of these mm-spaces.

The condition on dimensionality is inspired by the principle of trilateration (Thomas & Ros, 2005) and is generally satisfied in practical scenarios. The proof for Theorem 1 is provided in Appendix B.1. In addition to its metric properties, we also prove that QDOT has an empirical convergence rate.

**Theorem 2 (Empirical Convergence Rate)** Let  $\mathcal{X} = (X, d_X, \mu_X)$  be an mm-space where  $X \subseteq \mathbb{R}^d$ . Let  $\{x_i\}_{i=1}^n$  be i.i.d. samples drawn from  $\mu_X$ , and let the empirical measure be defined as  $\mu_n := \frac{1}{n} \sum_{i=1}^n \delta_{x_i}$ . If we denote the corresponding empirical mm-space as  $\mathcal{X}_n = (\{x_1, \dots, x_n\}, d_X, \mu_n)$ , then we have  $\mathcal{QD}_p(\mathcal{X}, \mathcal{X}_n) \to 0$ , as  $n \to \infty$ .

Moreover, under additional regularity conditions, namely (1) the distribution has a finite fourth moment and (2) its CDF is continuous at the chosen quantile levels and has a strictly positive density, the convergence rate of  $QD_p$  is at least as fast as the standard rate for the Wasserstein distance (Fournier & Guillin, 2015). Specifically, for the common case where  $d \geq 2$ , the expected error is bounded by:

$$\mathbb{E}[\mathcal{QD}_p(\mathcal{X}, \mathcal{X}_n)] = O(n^{-1/d}).$$

The finite moment requirement is a standard condition inherited from Fournier & Guillin (2015), while the continuity condition is a classical assumption that ensures the consistency of empirical quantiles. Both of these conditions are mild and hold for the vast majority of practical data distributions. For a detailed discussion and rates under other conditions, we refer the reader to Fournier & Guillin (2015) and our proof in Appendix B.2.

Based on the preceding theoretical results, we can summarize several key properties of the QDOT.

**Corollary 1** Let  $\mathcal{X} = (X, d_X, \mu_X)$  and  $\mathcal{Y} = (Y, d_Y, \mu_Y)$  be two mm-spaces embedded in Euclidean spaces, where  $X \subseteq \mathbb{R}^d$  and  $Y \subseteq \mathbb{R}^q$ . The QDOT distance exhibits the following properties:

- (1) **Location Invariance.** For any translation vector  $z \in \mathbb{R}^d$ , let the translated mm-space be denoted by  $\mathcal{X}_z = (X + z, d_X, \mu_{X+z})$ , we have  $\mathcal{QD}_p(\mathcal{X}_z, \mathcal{Y}) = \mathcal{QD}_p(\mathcal{X}, \mathcal{Y})$ .
- (2) Rotation and Reflection Invariance. For any orthogonal transformation  $R : \mathbb{R}^d \to \mathbb{R}^d$ , let  $\mathcal{X}_R = (R(X), d_{R(X)}, R_{\#\mu_X})$ , we have  $\mathcal{QD}_p(\mathcal{X}_R, \mathcal{Y}) = \mathcal{QD}_p(\mathcal{X}, \mathcal{Y})$ .
- (3) Numerical Convergence. Let  $\mathcal{X}_n$  and  $\mathcal{Y}_n$  be the empirical mm-spaces constructed from n i.i.d. samples drawn from  $\mathcal{X}$  and  $\mathcal{Y}$ , respectively. The empirical QDOT distance converges as  $n \to \infty$ :  $\mathcal{QD}_p(\mathcal{X}_n, \mathcal{Y}_n) \to \mathcal{QD}_p(\mathcal{X}, \mathcal{Y})$ . Furthermore, under the condition in2, its expectation convergence rate is given by  $\mathbb{E}[\mathcal{QD}_p(\mathcal{X}_n, \mathcal{Y}_n) \mathcal{QD}_p(\mathcal{X}, \mathcal{Y})] = \mathcal{O}(n^{-\frac{1}{\max(d,q)}})$ .

#### 3.3 Numerical Implementation

In the numerical setting, we consider discrete probability measures. Let  $\Delta^{n-1}$  denote the (n-1)-simplex. Given two sets of samples,  $X = \{x_i\}_{1 \leq i \leq n}$  in  $\mathbb{R}^d$  and  $Y = \{y_j\}_{1 \leq j \leq m}$  in  $\mathbb{R}^q$ , with corresponding probability vectors  $\mathbf{p}^X \in \Delta^{n-1}$  and  $\mathbf{p}^Y \in \Delta^{m-1}$ , their associated discrete measures are  $\mu_X = \sum_{i=1}^n p_i^X \delta_{x_i}$  and  $\mu_Y = \sum_{j=1}^m p_j^Y \delta_{y_j}$ , respectively. We represent the sample sets as data matrices  $\mathbf{X} = (x_1, x_2, \cdots, x_n)^\top \in \mathbb{R}^{n \times d}$  and  $\mathbf{Y} = (y_1, y_2, \cdots, y_m)^\top \in \mathbb{R}^{m \times q}$ . For a given quantile level vector  $\mathbf{q} \in (0, 1)^k$ , the procedure for computing the QDOT distance is detailed in Algorithm 1.

#### **Algorithm 1 QDOT**

```
271
                                              Require: X, Y, p^X, p^Y, q
272
                                                    1: Initialize the data matrices X and Y.
                                                 2: Compute the sample norms of \mathbf{X} and \mathbf{Y}: \phi_0^{\mathbf{X}} \leftarrow (\|\boldsymbol{x}_i\|)_{1 \leq i \leq n}, \phi_0^{\mathbf{Y}} \leftarrow (\|\boldsymbol{y}_j\|)_{1 \leq j \leq m}.
3: Compute the \boldsymbol{q}-quantiles of the sample norms: \qquad \qquad \triangleright \mathcal{O}(n \log n + m \log m)
\boldsymbol{r}^{\mathbf{X}} \leftarrow (F_{\|\mathbf{X}\|}^{-1}(q_1), \dots, F_{\|\mathbf{X}\|}^{-1}(q_k)), \quad \boldsymbol{r}^{\mathbf{Y}} \leftarrow (F_{\|\mathbf{Y}\|}^{-1}(q_1), \dots, F_{\|\mathbf{Y}\|}^{-1}(q_k)).
4: for i \leftarrow 1 to k do \qquad \qquad \triangleright \mathcal{O}(knd + kmq)
273
274
275
276
277
                                                                                    Compute the quantile weights:  w_{ij}^{\mathbf{X}} \leftarrow \frac{p_j^X \exp\{-\sigma(\boldsymbol{d}_j^{\mathbf{X}} - \boldsymbol{r}_i^{\mathbf{X}})^2\}}{\sum_{j'=1}^{n} p_{j'}^X \exp\{-\sigma(\boldsymbol{d}_j^{\mathbf{Y}} - \boldsymbol{r}_i^{\mathbf{X}})^2\}}, \quad w_{ij}^{\mathbf{Y}} \leftarrow \frac{p_j^Y \exp\{-\sigma(\boldsymbol{d}_j^{\mathbf{Y}} - \boldsymbol{r}_i^{\mathbf{Y}})^2\}}{\sum_{j'=1}^{m} p_{j'}^X \exp\{-\sigma(\boldsymbol{d}_j^{\mathbf{Y}} - \boldsymbol{r}_i^{\mathbf{X}})^2\}}, \\ \text{Compute the $q_i$-quantile means: } m_i^X \leftarrow \mathbf{X}^{\top} w_i^{\mathbf{X}}, \quad m_i^Y \leftarrow \mathbf{Y}^{\top} w_i^{\mathbf{Y}} \\ \text{Compute the distances to quantile means: } \phi_i^X \leftarrow (\|\boldsymbol{x}_j - \boldsymbol{m}_i^X\|)_{1 \leq j \leq n}, \quad \phi_i^Y \leftarrow (\|\boldsymbol{y}_j - \boldsymbol{m}_i^Y\|)_{1 \leq j \leq m} 
278
279
281
282
283
                                             9: Concatenate representations: \Phi^{\mathbf{X}} \leftarrow [\phi_0^{\mathbf{X}}, \phi_1^{\mathbf{X}}, \dots, \phi_k^{\mathbf{X}}], \quad \Phi^{\mathbf{Y}} \leftarrow [\phi_0^{\mathbf{Y}}, \phi_1^{\mathbf{Y}}, \dots, \phi_k^{\mathbf{Y}}]

10: Calculate the \mathcal{W}_p distance: \mathcal{QD}_p \leftarrow \mathcal{W}_p \left( (\Phi^{\mathbf{X}})_{\#} \mu_X, (\Phi^{\mathbf{Y}})_{\#} \mu_Y \right) \quad \triangleright \mathcal{O}(n^2 \log n + m^2 \log m)
284
                                              11: return QD_p
287
```

**Computational Cost.** For simplicity, we assume  $m \le n$  and  $q \le d$ . The initial norm computation requires  $\mathcal{O}(nd)$  operations. Computing the quantiles takes  $\mathcal{O}(n\log n)$  time. The main loop for computing the QWs, QDMs, and QDMDs has a total cost of  $\mathcal{O}(knd)$ . It is noteworthy that the final representations are generated in nearly linear time with respect to n. The final step of calculating the Wasserstein distance has a complexity of  $\mathcal{O}(n^2\log n)$  by Sinkhorn Algorithm(Cuturi, 2013) or Earth Moving Distance(Rubner et al., 2000).

To mitigate the high complexity of the final step, we also propose an Integral-QDOT approach.

#### 3.4 INTEGRAL-QDOT

To address the high computational cost of the standard Wasserstein distance, methods based on slicing(Bonneel et al., 2015; Deshpande et al., 2019) and the closed-form solution of one-dimensional Optimal Transport have become increasingly popular. We observe that the QDMD representation,  $\phi^X(x,q)$ , defined in the previous section, is a scalar value for any given quantile level q. This structure naturally inspires an alternative approach: instead of comparing the multi-dimensional representations in  $\mathbb{R}^{k+1}$ , we can compare the one-dimensional distributions of the QDMD scalars for each q and then aggregate the results. This leads to the Integral-QDOT(IQDOT) approach.

**Definition 4 (Integral-QDOT)** Given two mm-spaces  $\mathcal{X} = (X, d_X, \mu_X)$  and  $\mathcal{Y} = (Y, d_Y, \mu_Y)$ , where  $X \subseteq \mathbb{R}^d$ ,  $Y \subseteq \mathbb{R}^q$ , and  $d_X, d_Y$  are the standard Euclidean distances. For any quantile level  $q \in (0,1)$ , we can obtain their corresponding scalar QDMD representations,  $(\phi^X(\cdot,q))_{\#\mu_X}$  and  $(\phi^Y(\cdot,q))_{\#\mu_Y}$ . The Integral-QDOT distance is then defined as the  $L_p$ -norm of the 1-D Wasserstein distances between these push-forward measures, integrated over all  $q \in (0,1)$ :

$$\mathcal{IQD}_p(\mathcal{X}, \mathcal{Y}) := \inf_{\pi \in \Pi(\mu_X, \mu_Y)} \left( \int_{(0,1)} \int_{X \times Y} |\phi^X(x, q) - \phi^Y(y, q)|^p d\pi(x, y) dq \right)^{1/p}. \tag{8}$$

We further establish that IQDOT also constitutes a well-defined metric.

**Theorem 3 (Metric Property of IQDOT)** Let  $\mathcal{X} = (X, d_X, \mu_X)$  or  $\mathcal{Y} = (Y, d_Y, \mu_Y)$  be two mmspaces embedded in Euclidean spaces. If the set of QDMs satisfies the dimensionality condition  $\dim(\{\boldsymbol{m}^X(q)\}_{q\in(0,1)}) = \dim(X)$  and  $\dim(\{\boldsymbol{m}^Y(q)\}_{q\in(0,1)}) = \dim(Y)$ , then  $\mathcal{IQD}_p$  defines a metric on the isometry classes of these mm-spaces.

As for the numerical implementation of IQDOT, we consider computing the mean of the one-dimensional Wasserstein distances for the QDMD representations corresponding to k quantiles. Since the 1d Wasserstein distance has a closed-form solution, the computational cost of the final step in Algorithm 1 is reduced to  $\mathcal{O}(kn\log n)$ , thereby achieving highly efficient computation.

## 4 EXPERIMENTS

We next present a series of experiments designed to empirically validate several key properties of our proposed framework. The evaluation aims to demonstrate QDOT's: (1) effectiveness in cross-space tasks; (2) fast computational efficiency; (3) transferability and versatility in comparing diverse distributions; and (4) strong performance in complex models. All CPU-based experiments were conducted on an Intel(R) Xeon(R) Platinum 8280 CPU @ 2.70GHz with 256GB RAM. All GPU-based experiments were performed on a single NVIDIA RTX 4090 GPU with 24GB VRAM.

#### 4.1 CROSS SPACE TASKS

To evaluate the cross-space alignment capability and the metric accuracy of our proposed method, we conduct experiments on the camel-gallop data(Sumner & Popović, 2004). This data consists of a reference 3D point cloud model and a corresponding 48-frame sequence of a galloping camel. For our tests, we subsample the point clouds to 10,000 points. The experimental task is to match each frame of the 3D sequence against projections of the static reference model onto three distinct 2D subspaces. We assess alignment quality using the Transformed Mean Squared Error (TMSE) and the Inlier Ratio (IR), formally defined in Appendix D.2. TMSE measures the alignment cost in the original 3D space, while IR quantifies the percentage of correctly matched points within a given tolerance. Our comparisons are structured as follows: First, we evaluate our QDOT by Sinkhorn algorithm, against Entropic Gromov-Wasserstein (EGW) (Peyré et al., 2016), testing both methods with regularization parameters  $\lambda=0.1$  and  $\lambda=0.01$ . Second, to assess performance with sparse couplings, we compare QDOT by EMD against the classical GW method.

Table 1: Cross Space Results on the camel-gallop dataset

Methods	Transformed MSE $\downarrow$				Inlier Ratio (%) ↑						
	3D	$2D_{1st}$	2D <sub>2nd</sub>	$2D_{3rd}$	Avg.	3D	$2D_{1st}$	2D <sub>2nd</sub>	$2D_{3rd}$	Avg.	$\mathit{Time}_{(\times 10^2 s)}$
$EGW_{(\lambda=0.1)}$	0.37	0.37	0.39	0.40	0.38	41.71	41.90	39.66	38.69	40.49	3.03
QDOT <sub>(Sink-0.1)</sub>	0.32	0.32	0.32	0.33	0.32	48.75	48.25	48.84	47.03	48.22	0.13
$EGW_{(\lambda=0.01)}$	0.21	0.23	0.22	0.33	0.25	72.12	64.92	69.51	50.89	64.36	21.51
QDOT <sub>(Sink-0.01)</sub>	0.22	0.24	0.23	0.25	0.24	71.15	63.75	68.48	63.20	66.65	0.74
GW	0.26	0.25	0.30	0.35	0.29	61.29	63.04	56.78	49.16	58.06	14.63
QDOT <sub>(EMD)</sub>	0.25	0.27	0.25	0.27	0.26	63.35	58.60	63.87	60.02	61.46	0.38

As shown in Table 1, the couplings produced by QDOT achieve comparable or superior results to those from GW and EGW. Notably, this performance is attained with a computational cost that is up to 30 times lower, underscoring the effectiveness and efficiency of our algorithm for cross-space alignment. Furthermore, to evaluate the ability of cross-space metrics, we plotted the dissimilarity trends over the point-cloud sequence, and included IQDOT, SGW, and RISGW for comparison.

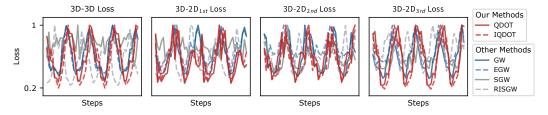


Figure 4: Loss trends between the camel-gallop sequence and the four reference in different spaces. It is evident that QDOT, IQDOT, GW, and EGW successfully capture the four periodic cycles of the camel's gallop, exhibiting consistent fluctuation amplitudes and trends across all reference spaces. In contrast, SGW and RISGW fail to reveal a clear, meaningful pattern.

The results, visualized in Figure 4, demonstrate that methods possessing well-defined metric properties consistently produce effective and meaningful dissimilarity curves. In stark contrast, SGW and RISGW, which lack these theoretical guarantees, fail to reveal this underlying geometric pattern, underscoring the importance of sound metric properties.

#### 4.2 TIME COST

378

379 380

381

382

383

384

385

386

387

388

389

390

391

392

393

394

395 396

397 398

399

400

401

402

403

404

405

406 407

408 409

410

411

412

413

414

415

416

417

418

419 420

421

422

423

424

425 426

427 428

429

430

431

To benchmark the practical runtime performance of our methods, we conduct an experiment using randomly generated 2D point clouds, with the number of support points ranging from  $10^2$  to  $10^6$ . We compare our methods against GW, EGW, SGW, and RISGW. For this test, all methods are configured with 50 projections (or quantiles); IQDOT is additionally implemented using 5 quantiles, and QDOT uses an EMD solver. The results are visualized in Figure 5. The plot clearly shows that for smaller sample sizes, the QDOT family of methods is faster than all baselines. In large-scale scenarios involving tens of thousands of points, IQDOT achieves a remarkable speedup, running up to thousands of times faster than the classic GW method. QDOT also maintains a significant advantage. The overall trends in the log-log plot confirm our complexity analysis: IQDOT exhibits a quasi-linear time complexity, whereas QDOT scales quadratically.

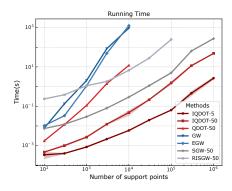


Figure 5: Runtime comparison of various methods. All experiments were conducted in a CPU environment.

#### 4.3 Transfer Learning

Leveraging the robust cross-domain behavior of QDOT, we propose a parameter-free transfer learning approach for point cloud classification. We evaluate it on ModelNet40 (Wu et al., 2015) and ShapeNetPart (Fan et al., 2017), focusing on the seven shared classes (3,072 samples from ModelNet and 15,402 from ShapeNet). This setting entails tens of millions of pairwise comparisons, making computationally intensive approaches infeasible; we therefore restrict our baselines to linear-time methods. To test rotational robustness, random vertical-axis rotations are applied, and experiments are conducted with point clouds subsampled to n=1024 and n=2048. The task treats one dataset as the source and the other as the target: each target cloud is classified via 1-Nearest Neighbor using distances to all source clouds computed by our metric. Results are reported below.

Methods (Mo→Sh) chair Time(h) airplane car guitar lamp laptop table Avg. SGW-1024 96.69 99.10 60.67 71.29 98.66 94.25 79.77 13.11 0.00SGW-2048 96.98 99.10 60.24 0.00 71.62 98.66 94.47 79.83 21.84 IQDOT-1024 1.59 94.86 84.18 83.82 96.56 51.13 96.45 84.95 83.89 **IODOT-2048** 95.35 87.86 85.01 96.44 52.55 98.00 87.15 85.42 5.98 guitar Methods (Sh $\rightarrow$ Mo) airplane car chair lamp laptop table Avg. Time(h) SGW-1024 89.66 37.03 43.68 0.00 80.55 20.11 99.59 59.66 13.11 SGW-2048 99.79 91.18 37.03 44.08 0.00 83.33 20.11 60.31 21.84 **IQDOT-1024** 93.97 98.89 87.20 91.10 98.43 81.94 97.04 96.74 1.59 **IQDOT-2048** 98.89 91.91 93.22 98.03 77.77 97.63 97.56 95.05 5.98

Table 2: Point Cloud Transfer Learning Classification Accuracy (%)

As shown in Table 2, IQDOT demonstrates more consistent performance compared to SGW. In the ModelNet-to-ShapeNet transfer task, IQDOT achieves superior overall results. Furthermore, its performance in the ShapeNet-to-ModelNet direction is exceptionally strong, improving the average accuracy by 35% over the SGW method. Notably, IQDOT completed this large-scale task in just a few hours, highlighting its remarkable efficiency.

#### 4.4 MOLECULE GENERATION

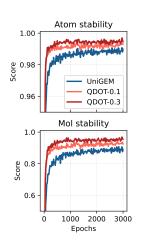
Molecular generation is a central challenge in drug discovery and molecular science. Recent advances show that diffusion models with equivariant neural architectures have become the dominant paradigm (Hoogeboom et al., 2022; Xu et al., 2023; Song et al., 2024; Feng et al., 2025). Since molecular properties are invariant to absolute position and orientation, prior works attempted to enforce stability through explicit structural alignment during training (Song et al., 2023; Hassan et al.,

2024), but such procedures incur high computational cost. We instead integrate the QDOT distance directly into the diffusion loss, providing an alignment-free learning signal that is inherently invariant and emphasizes geometric fidelity. This encourages physically stable and generalizable molecule generation without the overhead of explicit alignment.

We evaluate this approach on QM9 (Ramakrishnan et al., 2014) by adapting the classic model EDM (Hoogeboom et al., 2022) and a recently model UniGEM (Feng et al., 2025), with QDOT-augmented loss functions. Both are trained for 3000 epochs, and the results are reported below.

Table 3: Molecule Generation Results on QM9 Dataset

Methods	Atom sta(%)	Mol sta(%)	Valid(%)	V * U(%)
Data	99.0	95.2	97.7	97.7
EDM	$98.70_{\pm 0.01}$	86.56 <sub>±0.27</sub>	$93.73_{\pm 0.12}$	92.04 <sub>±0.03</sub>
$\rightarrow$ MSE <sub>0.1</sub> $\rightarrow$ MSE <sub>0.3</sub>	$98.81_{\pm 0.04} \ 98.68_{\pm 0.07}$	$87.63_{\pm 0.21}$ $86.04_{\pm 0.74}$	$94.54_{\pm 0.15}$ $93.83_{\pm 0.43}$	$92.81_{\pm 0.20}$ $92.39_{\pm 0.55}$
$\rightarrow$ QDOT <sub>0.1</sub> $\rightarrow$ QDOT <sub>0.3</sub>	$99.15_{\pm 0.01}$ $99.35_{\pm 0.06}$	$90.91_{\pm 0.21}$ $93.29_{\pm 0.81}$	$96.75_{\pm 0.33}$ $97.80_{\pm 0.22}$	$89.06_{\pm 0.19} \\ 81.21_{\pm 2.03}$
UniGEM	$98.90_{\pm 0.03}$	$89.40_{\pm0.02}$	$94.58_{\pm 0.07}$	$92.75_{\pm0.11}$
$\rightarrow$ MSE <sub>0.1</sub>	$98.91_{\pm 0.04}$	$89.00_{\pm 0.22}$	$95.00_{\pm 0.16}$	$93.13_{\pm 0.10}$
$\rightarrow$ MSE <sub>0.3</sub>	$99.00_{\pm 0.07}$	$89.48_{\pm 0.70}$	$95.16_{\pm0.16}$	$93.34_{\pm0.18}$
$\rightarrow$ QDOT <sub>0.1</sub> $\rightarrow$ QDOT <sub>0.3</sub>	$99.24_{\pm 0.03}$ $99.44_{\pm 0.01}$	$92.73_{\pm 0.09}$ $95.23_{\pm 0.18}$	$96.85_{\pm 0.11}$ $97.94_{\pm 0.06}$	$93.42_{\pm 0.23}$ $83.88_{\pm 0.30}$



<sup>&</sup>lt;sup>1</sup> The notation "QDOT<sub>0.1</sub>" indicates that the error function  $\mathcal{L} \leftarrow 0.1\mathcal{L}_{QDOT} + 0.9\mathcal{L}_{MSE}$ ; "MSE<sub>0.1</sub>" indicates that  $\mathcal{L} \leftarrow 0.9\mathcal{L}_{MSE}$ .

Figure 6: Stability trends across epochs.

As shown in Table 3, our QDOT loss significantly enhances both EDM and UniGEM baselines. For UniGEM, a mere 10% QDOT weight (QDOT<sub>0.1</sub>) improves all four metrics, while a 30% weight (QDOT<sub>0.3</sub>) establishes a new state-of-the-art, boosting Atom and Molecule Stability to 99.44% and 95.23%, respectively. An ablation study confirms these gains stem from QDOT's geometric guidance, not merely from re-weighting the original loss. Furthermore, as illustrated in Figure 6, QDOT accelerates training, halving the convergence time for UniGEM to just 1000 epochs and thereby improving training efficiency and stability.

To assess the fine-tuning capability of the QDOT loss, we conducted an additional experiment on GEOM-Drug (Axelrod & Gomez-Bombarelli, 2022). Starting from a UniGEM model pre-trained for 13 epochs (Feng et al., 2025), we introduced QDOT into the loss and fine-tuned for 3 more epochs with a learning rate of  $10^{-4}$ . The results in Table 4 show that while continued training with the original MSE loss yields no further gains, incorporating QDOT leads to substantial improvements in

Table 4: Fine-tuning Results on GEOM-Drugs

Methods	Atom sta(%)	Mol sta(%)	Valid(%)
UniGEM	84.84	1.20	98.29
$\begin{array}{l} \rightarrow \text{MSE} \\ \rightarrow \text{QDOT}_{0.1} \\ \rightarrow \text{QDOT}_{0.2} \\ \rightarrow \text{QDOT}_{0.3} \end{array}$	84.39 87.94 <b>91.87</b> 91.73	1.13 6.04 <u>17.14</u> <b>22.05</b>	99.01 98.46 97.86 94.62

Atom Stability and Molecule Stability, with only a negligible drop in validity.

#### 5 CONCLUSION

This work introduces QDOT, a novel geometric metric. By constructing isometry-invariant anchors from distance quantiles and deriving a lossless distance representation through trilateration theory, QDOT establishes a rigorous metric on isometry classes of mm-spaces. Experimental results demonstrate its effectiveness across cross-space alignment and comparison, computational efficiency, transfer learning, and molecular generation tasks. Future directions include: (1) extending QDOT to hyperbolic and spherical spaces; (2) developing more general approaches for constructing isometry-invariant anchors and representations; and (3) applying QDOT-based alignment techniques and QDOT as a loss function in complex models.

#### ETHICS STATEMENT

All datasets utilized in this study are publicly available, and their use raises no ethical concerns.

#### REPRODUCIBILITY STATEMENT

We are committed to ensuring the reproducibility of our research. The source code for all experiments is provided in the supplementary materials, which can be used to reproduce the results presented in this paper. For the molecule generation experiment, we have included the evaluate checkpoint of QDOT-0.3. This checkpoint can also be reproduced using the provided training code. Furthermore, we have explicitly stated all key assumptions for the theorems presented. The sources and preprocessing scripts for all datasets used in this work are also provided.

#### REFERENCES

- Sayali Anil Alatkar and Daifeng Wang. CMOT: cross-modality optimal transport for multimodal inference. *Genome Biology*, 24(1):163, 2023.
- Jason Altschuler, Francis Bach, Alessandro Rudi, and Jonathan Niles-Weed. Massively scalable sinkhorn distances via the nyström method. *Advances in neural information processing systems*, 32, 2019.

Martin Arjovsky, Soumith Chintala, and Léon Bottou. Wasserstein generative adversarial networks. In *International conference on machine learning*, pp. 214–223. PMLR, 2017.

Arip Asadulaev, Rostislav Korst, Aleksandr Korotin, Vage Egiazarian, Andrey Filchenkov, and Evgeny Burnaev. Rethinking optimal transport in offline reinforcement learning. *Advances in Neural Information Processing Systems*, 37:123592–123607, 2024.

Simon Axelrod and Rafael Gomez-Bombarelli. Geom, energy-annotated molecular conformations for property prediction and molecular generation. *Scientific Data*, 9(1):185, 2022.

Serge Belongie, Jitendra Malik, and Jan Puzicha. Shape matching and object recognition using shape contexts. *IEEE transactions on pattern analysis and machine intelligence*, 24(4):509–522, 2002.

Nicolas Bonneel, Michiel Van De Panne, Sylvain Paris, and Wolfgang Heidrich. Displacement interpolation using lagrangian mass transport. In *Proceedings of the 2011 SIGGRAPH Asia conference*, pp. 1–12, 2011.

Nicolas Bonneel, Julien Rabin, Gabriel Peyré, and Hanspeter Pfister. Sliced and Radon Wasserstein barycenters of measures. *Journal of Mathematical Imaging and Vision*, 51:22–45, 2015.

Chao Chen, Guanbin Li, Ruijia Xu, Tianshui Chen, Meng Wang, and Liang Lin. Clusternet: Deep hierarchical cluster network with rigorously rotation-invariant representation for point cloud analysis. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 4994–5002, 2019.

Scott Cohen and L Guibasm. The earth mover's distance under transformation sets. In *Proceedings of the Seventh IEEE International Conference on Computer Vision*, volume 2, pp. 1076–1083. IEEE, 1999.

Marco Cuturi. Sinkhorn distances: Lightspeed computation of optimal transport. *Advances in neural information processing systems*, 26, 2013.

Valentin De Bortoli, James Thornton, Jeremy Heng, and Arnaud Doucet. Diffusion schrödinger bridge with applications to score-based generative modeling. *Advances in neural information processing systems*, 34:17695–17709, 2021.

- Ishan Deshpande, Yuan-Ting Hu, Ruoyu Sun, Ayis Pyrros, Nasir Siddiqui, Sanmi Koyejo, Zhizhen
   Zhao, David Forsyth, and Alexander G Schwing. 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, 2019.
  - Haoqiang Fan, Hao Su, and Leonidas J Guibas. A point set generation network for 3d object reconstruction from a single image. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 605–613, 2017.
  - Shikun Feng, Yuyan Ni, Lu yan, Zhi-Ming Ma, Wei-Ying Ma, and Yanyan Lan. UniGEM: A unified approach to generation and property prediction for molecules. In *The Thirteenth International Conference on Learning Representations*, 2025. URL https://openreview.net/forum?id=Lb91pXwZMR.
  - Rémi Flamary, Nicolas Courty, Alexandre Gramfort, Mokhtar Z Alaya, Aurélie Boisbunon, Stanislas Chambon, Laetitia Chapel, Adrien Corenflos, Kilian Fatras, Nemo Fournier, et al. Pot: Python optimal transport. *Journal of Machine Learning Research*, 22(78):1–8, 2021.
  - Nicolas Fournier and Arnaud Guillin. On the rate of convergence in wasserstein distance of the empirical measure. *Probability theory and related fields*, 162(3):707–738, 2015.
  - Fabian Fuchs, Daniel Worrall, Volker Fischer, and Max Welling. SE (3)-transformers: 3d roto-translation equivariant attention networks. *Advances in neural information processing systems*, 33:1970–1981, 2020.
  - Michael Gromov. Groups of polynomial growth and expanding maps (with an appendix by Jacques Tits). *Publications Mathématiques de l'IHÉS*, 53:53–78, 1981.
  - Majdi Hassan, Nikhil Shenoy, Jungyoon Lee, Hannes Stärk, Stephan Thaler, and Dominique Beaini. Et-flow: Equivariant flow-matching for molecular conformer generation. *Advances in Neural Information Processing Systems*, 37:128798–128824, 2024.
  - Emiel Hoogeboom, Victor Garcia Satorras, Clement Vignac, and Max Welling. Equivariant diffusion for molecule generation in 3d. In *International conference on machine learning*, pp. 8867–8887. PMLR, 2022.
  - Seohyun Kim, Jaeyoo Park, and Bohyung Han. Rotation-invariant local-to-global representation learning for 3d point cloud. *Advances in Neural Information Processing Systems*, 33:8174–8185, 2020.
  - Pascal Klink, Haoyi Yang, Carlo D'Eramo, Jan Peters, and Joni Pajarinen. Curriculum reinforcement learning via constrained optimal transport. In *International Conference on Machine Learning*, pp. 11341–11358. PMLR, 2022.
  - Matt Kusner, Yu Sun, Nicholas Kolkin, and Kilian Weinberger. From word embeddings to document distances. In *International conference on machine learning*, pp. 957–966. PMLR, 2015.
  - Tam Le, Makoto Yamada, Kenji Fukumizu, and Marco Cuturi. Tree-sliced variants of wasserstein distances. *Advances in neural information processing systems*, 32, 2019.
  - Mengyu Li, Jun Yu, Tao Li, and Cheng Meng. Importance sparsification for sinkhorn algorithm. *Journal of Machine Learning Research*, 24(247):1–44, 2023a.
  - Mengyu Li, Jun Yu, Hongteng Xu, and Cheng Meng. Efficient approximation of Gromov-Wasserstein distance using importance sparsification. *Journal of Computational and Graphical Statistics*, 32(4):1512–1523, 2023b.
  - Tao Li, Cheng Meng, Hongteng Xu, and Jun Yu. Hilbert curve projection distance for distribution comparison. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 46(7):4993–5007, 2024.
  - Tao Li, Cheng Meng, Hongteng Xu, and Jun Zhu. Efficient variants of wasserstein distance in hyperbolic space via space-filling curve projection. *IEEE Transactions on Neural Networks and Learning Systems*, 2025.

- Igor Melnyk, Youssef Mroueh, Brian Belgodere, Mattia Rigotti, Apoorva Nitsure, Mikhail
   Yurochkin, Kristjan Greenewald, Jiri Navratil, and Jarret Ross. Distributional preference alignment of Ilms via optimal transport. Advances in Neural Information Processing Systems, 37:
   104412–104442, 2024.
  - Facundo Mémoli. Gromov–Wasserstein distances and the metric approach to object matching. *Foundations of computational mathematics*, 11:417–487, 2011.
  - Lawrence Narici and Edward Beckenstein. *Topological vector spaces*. Chapman and Hall/CRC, 2010.
  - Gabriel Peyré, Marco Cuturi, and Justin Solomon. Gromov-Wasserstein averaging of kernel and distance matrices. In *International conference on machine learning*, pp. 2664–2672. PMLR, 2016.
  - Gabriel Peyré, Marco Cuturi, et al. Computational optimal transport: With applications to data science. *Foundations and Trends*® *in Machine Learning*, 11(5-6):355–607, 2019.
  - Raghunathan Ramakrishnan, Pavlo O Dral, Matthias Rupp, and O Anatole Von Lilienfeld. Quantum chemistry structures and properties of 134 kilo molecules. *Scientific data*, 1(1):1–7, 2014.
  - Yossi Rubner, Carlo Tomasi, and Leonidas J Guibas. The earth mover's distance as a metric for image retrieval. *International journal of computer vision*, 40(2):99–121, 2000.
  - Meyer Scetbon, Gabriel Peyré, and Marco Cuturi. Linear-time Gromov Wasserstein distances using low rank couplings and costs. In *International Conference on Machine Learning*, pp. 19347–19365. PMLR, 2022.
  - RJ Serfling. Approximation theorems of mathematical statistics. Wiley series in probability and mathematical statistics Show all parts in this series, 1980.
  - Liangliang Shi, Jack Fan, and Junchi Yan. Ot-clip: Understanding and generalizing clip via optimal transport. In *Forty-first International Conference on Machine Learning*, 2024.
  - Yuxuan Song, Jingjing Gong, Minkai Xu, Ziyao Cao, Yanyan Lan, Stefano Ermon, Hao Zhou, and Wei-Ying Ma. Equivariant flow matching with hybrid probability transport for 3d molecule generation. *Advances in Neural Information Processing Systems*, 36:549–568, 2023.
  - Yuxuan Song, Jingjing Gong, Hao Zhou, Mingyue Zheng, Jingjing Liu, and Wei-Ying Ma. Unified generative modeling of 3d molecules with bayesian flow networks. In *The Twelfth International Conference on Learning Representations*, 2024.
  - Karl-Theodor Sturm. On the geometry of metric measure spaces. *Acta Mathematica*, 196(1):65–131, 2006.
  - Karl-Theodor Sturm. The space of spaces: curvature bounds and gradient flows on the space of metric measure spaces, volume 290. American Mathematical Society, 2023.
  - Robert W Sumner and Jovan Popović. Deformation transfer for triangle meshes. *ACM Transactions on graphics (TOG)*, 23(3):399–405, 2004.
  - 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, 2019.
  - Federico Thomas and Lluis Ros. Revisiting trilateration for robot localization. *IEEE Transactions on robotics*, 21(1):93–101, 2005.
  - Vayer Titouan, Rémi Flamary, Nicolas Courty, Romain Tavenard, and Laetitia Chapel. Sliced Gromov-Wasserstein. *Advances in Neural Information Processing Systems*, 32, 2019.
  - Alexander Tong, Kilian FATRAS, Nikolay Malkin, Guillaume Huguet, Yanlei Zhang, Jarrid Rector-Brooks, Guy Wolf, and Yoshua Bengio. Improving and generalizing flow-based generative models with minibatch optimal transport. *Transactions on Machine Learning Research*, 2024. ISSN 2835-8856. URL https://openreview.net/forum?id=CD9Snc73AW. Expert Certification.

- Cédric Villani et al. Optimal transport: old and new, volume 338. Springer, 2008.
  - Zhirong Wu, Shuran Song, Aditya Khosla, Fisher Yu, Linguang Zhang, Xiaoou Tang, and Jianxiong Xiao. 3d shapenets: A deep representation for volumetric shapes. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 1912–1920, 2015.
  - Hongteng Xu, Dixin Luo, Hongyuan Zha, and Lawrence Carin Duke. Gromov-Wasserstein learning for graph matching and node embedding. In *International conference on machine learning*, pp. 6932–6941. PMLR, 2019.
  - Minkai Xu, Alexander S Powers, Ron O Dror, Stefano Ermon, and Jure Leskovec. Geometric latent diffusion models for 3d molecule generation. In *International Conference on Machine Learning*, pp. 38592–38610. PMLR, 2023.
  - Yingxue Xu and Hao Chen. Multimodal optimal transport-based co-attention transformer with global structure consistency for survival prediction. In *Proceedings of the IEEE/CVF international conference on computer vision*, pp. 21241–21251, 2023.
  - Jianyu Yang, Hongxing Wang, Junsong Yuan, Youfu Li, and Jianyang Liu. Invariant multi-scale descriptor for shape representation, matching and retrieval. *Computer Vision and Image Under-standing*, 145:43–58, 2016.
  - Hao Yu, Zheng Qin, Ji Hou, Mahdi Saleh, Dongsheng Li, Benjamin Busam, and Slobodan Ilic. Rotation-invariant transformer for point cloud matching. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 5384–5393, 2023.

#### A ADDITIONAL BACKGROUND

 Here we introduce additional background concepts that are related to our work.

One-Dimensional Wasserstein Distance. A key special case of the Wasserstein distance defined in equation 1 occurs in one dimension, where it admits a convenient closed-form solution. For two one-dimensional distributions  $\mu_X$  and  $\mu_Y$ , let  $F_X$  and  $F_Y$  be their respective cumulative distribution functions , and let  $F_X^{-1}$  and  $F_Y^{-1}$  be the corresponding quantile functions. The p-Wasserstein distance is then given by:

$$W_p(\mu_X, \mu_Y) = \left( \int_0^1 |F_X^{-1}(u) - F_Y^{-1}(u)|^p \, du \right)^{\frac{1}{p}}. \tag{9}$$

For discrete distributions, the empirical quantile functions are readily obtained by sorting, making the computation of the 1D Wasserstein distance explicit and efficient.

**Gromov-Wasserstein Distance**(Mémoli, 2011). While the original formulation of the Gromov-Wasserstein distance in equation 2 is theoretically elegant, its computation is NP-hard. Consequently, an alternative and more commonly used formulation was introduced. Given two mm-spaces  $\mathcal{X}=(X,d_X,\mu_X)$  and  $\mathcal{Y}=(Y,d_Y,\mu_Y)$ , for  $p\geq 1$ , the p-Gromov-Wasserstein distance is defined as:

$$\mathcal{GW}_p(\mathcal{X}, \mathcal{Y}) := \left(\inf_{\pi \in \Pi(\mu_X, \mu_Y)} \int_{X \times Y} \int_{X \times Y} |d_X(x, x') - d_Y(y, y')|^p d\pi(x, y) d\pi(x', y')\right)^{\frac{1}{p}}.$$
(10)

Mémoli (2011) proved that the formulation in equation 10 is bi-Hölder equivalent to the one in equation 2. Moreover, equation 2 constitutes a metric in the sense of Definition 1.

**Computational Costs.** We now briefly review the computational complexity of these OT-based methods for discrete distributions supported on n points. Computing the standard Wasserstein distance in equation 1 using a classic EMD solver has a worst-case complexity of  $\mathcal{O}(n^3)$  and an average-case complexity of  $\mathcal{O}(n^2)$  (Bonneel et al., 2011). Alternatively, solving the entropyregularized OT problem with the Sinkhorn algorithm (Cuturi, 2013) has a complexity of up to  $\mathcal{O}(n^2 \log n)$ . While some targeted algorithms may achieve faster runtimes (Altschuler et al., 2019; Li et al., 2023a), this often comes at the cost of sacrificing metric properties. Computing the Gromov-Wasserstein distance in equation 10 involves solving a non-convex quadratic program, which has a complexity of at least  $\mathcal{O}(n^3 \log n)$  (Xu et al., 2019). Faster algorithms, such as those employing entropic regularization (Peyré et al., 2016; Scetbon et al., 2022; Li et al., 2023b), typically represent a trade-off between speed and accuracy. As for the one-dimensional Wasserstein distance in equation 9, the computation is significantly faster, requiring only  $\mathcal{O}(n \log n)$  time via sorting. This efficiency has motivated many methods that project higher-dimensional distributions onto one-dimensional lines for comparison (Bonneel et al., 2015; Deshpande et al., 2019; Le et al., 2019; Li et al., 2024; 2025). However, for the Gromov-Wasserstein distance, the development of sliced variants is less mature, with Sliced Gromov-Wasserstein (SGW) (Titouan et al., 2019) being the most prominent example that has achieved empirical success.

#### **B** Proofs

#### B.1 Proof of Theorem 1

The proof for the symmetry and triangle inequality properties is straightforward, as they are directly inherited from the Wasserstein distance  $W_p$ . We focus on proving the **Identity of Indiscernibles**, which is the most involved part.

To do so, we first recall an equivalent characterization of isometry from (Sturm, 2023, Lemma 1.10):

(i)  $\mathcal{X}, \mathcal{Y}$  are isometric.

 (ii) There exists a coupling  $\pi \in \Pi(\mu_X, \mu_Y)$  such that  $d_X(\boldsymbol{x}_0, \boldsymbol{x}_1) = d_Y(\boldsymbol{y}_0, \boldsymbol{y}_1)$  holds for  $\pi \otimes \pi$ -almost every pair of points  $((\boldsymbol{x}_0, \boldsymbol{y}_0), (\boldsymbol{x}_1, \boldsymbol{y}_1))$ .

The "if" part is straightforward: if  $\mathcal{X}$  and  $\mathcal{Y}$  are isometric, their QDMD representations will be identically distributed by construction, making their QDOT distance zero. We now prove the "only if" part: if  $\mathcal{QD}_p(\mathcal{X}, \mathcal{Y}) = 0$ , then  $\mathcal{X}$  and  $\mathcal{Y}$  are isometric by satisfying condition (ii).

**Proof.** By definition, the condition  $\mathcal{QD}_p(\mathcal{X}, \mathcal{Y}) = 0$  implies that there exists an optimal coupling  $\pi$  between the pushforward measures on the representation space, such that  $\phi^X(x) = \phi^Y(y)$  holds for all (x, y) in the support of the corresponding coupling on  $X \times Y$ .

For a given quantile level vector  $\mathbf{q} \in (0,1)^k$ , let us denote the corresponding QDMs as  $\mathbf{m}_i^X = \mathbf{m}^X(q_i)$  and  $\mathbf{m}_i^Y = \mathbf{m}^Y(q_i)$  for  $i = 1, 2, \dots, k$ . The equality of the representations implies:

$$\left\{ egin{aligned} \|oldsymbol{x} - oldsymbol{m}_i^X\|_2 &= \|oldsymbol{y} - oldsymbol{m}_i^Y\|_2, \ \|oldsymbol{x}\|_2 &= \|oldsymbol{y}\|_2, \end{aligned} 
ight. \quad orall_i (oldsymbol{x}, oldsymbol{y}) \in \operatorname{supp}(\pi).$$

From  $\|x\|_2 = \|y\|_2$  and the fact that isometric spaces have identical norm distributions (and thus identical quantile functions), it directly follows that the weights are equal for coupled points:

$$w^X(\boldsymbol{x}, q_i) = w^Y(\boldsymbol{y}, q_i), \quad \forall (\boldsymbol{x}, \boldsymbol{y}) \in \operatorname{supp}(\pi), i = 1, 2, \dots k.$$

For any  $(x, y) \in \text{supp}(\pi)$ , expanding the squared distance equality yields:

$$\begin{aligned} &\|\boldsymbol{x} - \boldsymbol{m}_i^X\|_2^2 = \|\boldsymbol{y} - \boldsymbol{m}_i^Y\|_2^2 \\ \Rightarrow & \boldsymbol{x}^\top \boldsymbol{x} + \boldsymbol{m}_i^{X^\top} \boldsymbol{m}_i^X - 2\boldsymbol{x}^\top \boldsymbol{m}_i^X = \boldsymbol{y}^\top \boldsymbol{y} + \boldsymbol{m}_i^{Y^\top} \boldsymbol{m}_i^Y - 2\boldsymbol{y}^\top \boldsymbol{m}_i^Y \\ \Rightarrow & \boldsymbol{m}_i^{X^\top} \boldsymbol{m}_i^X - 2\boldsymbol{x}^\top \boldsymbol{m}_i^X = \boldsymbol{m}_i^{Y^\top} \boldsymbol{m}_i^Y - 2\boldsymbol{y}^\top \boldsymbol{m}_i^Y, \end{aligned}$$

where the last step uses  $\|\boldsymbol{x}\|_2^2 = \|\boldsymbol{y}\|_2^2$ . Integrating this equality against the weight function  $w(\cdot, q_j)$  over the respective spaces gives:

$$\int_{X} w^{X}(q_{j}, \boldsymbol{x}) (\boldsymbol{m}_{i}^{X^{\top}} \boldsymbol{m}_{i}^{X} - 2\boldsymbol{x}^{\top} \boldsymbol{m}_{i}^{X}) d\mu_{X}(\boldsymbol{x})$$

$$= \int_{Y} w^{Y}(q_{j}, \boldsymbol{y}) (\boldsymbol{m}_{i}^{Y^{\top}} \boldsymbol{m}_{i}^{Y} - 2\boldsymbol{y}^{\top} \boldsymbol{m}_{i}^{Y}) d\mu_{Y}(y)$$

$$= \mathbb{E}_{X}(w^{X}(q_{j}, X)) \boldsymbol{m}_{i}^{X^{\top}} \boldsymbol{m}_{i}^{X} - 2\mathbb{E}_{X}(w^{X}(q_{j}, X)X)$$

$$= \mathbb{E}_{Y}(w^{X}(q_{j}, Y)) \boldsymbol{m}_{i}^{Y^{\top}} \boldsymbol{m}_{i}^{Y} - 2\mathbb{E}_{Y}(w^{Y}(q_{j}, Y)Y).$$

By the definition of QDM in Eq. equation 5, this simplifies to:

$$\boldsymbol{m}_{i}^{X^{\top}} \boldsymbol{m}_{i}^{X} - 2 \boldsymbol{m}_{i}^{X^{\top}} \boldsymbol{m}_{j}^{X} = \boldsymbol{m}_{i}^{Y^{\top}} \boldsymbol{m}_{i}^{Y} - 2 \boldsymbol{m}_{i}^{Y^{\top}} \boldsymbol{m}_{j}^{Y}.$$

By setting i=j, we find  $\|\boldsymbol{m}_i^X\|_2^2 = \|\boldsymbol{m}_i^Y\|_2^2$ , which implies  $\|\boldsymbol{m}_i^X\|_2 = \|\boldsymbol{m}_i^Y\|_2$ . Substituting this back, we obtain the key result that the Gram matrices of the QDM sets are identical:  $\boldsymbol{m}_i^{X^\top} \boldsymbol{m}_j^X = \boldsymbol{m}_i^{Y^\top} \boldsymbol{m}_i^Y$ .

Now, for any  $(x, y) \in \text{supp}(\pi)$ , the relation  $\|x - m_i^X\|_2^2 = \|y - m_i^Y\|_2^2$  simplifies further:

$$egin{aligned} oldsymbol{x}^{ op} oldsymbol{x} + oldsymbol{m}_i^{X^{ op}} oldsymbol{m}_i^{X} - 2 oldsymbol{x}^{ op} oldsymbol{m}_i^{X} = oldsymbol{y}^{ op} oldsymbol{m}_i^{Y} - 2 oldsymbol{y}^{ op} oldsymbol{m}_i^{Y} \ \Rightarrow oldsymbol{x}^{ op} oldsymbol{m}_i^{X} = oldsymbol{y}^{ op} oldsymbol{m}_i^{Y}. \end{aligned}$$

Due to the dimensionality condition  $\dim(\{\boldsymbol{m}^X(q_i)\}_{1\leq i\leq k})=\dim(X)$ , any point  $\boldsymbol{x}\in \operatorname{supp}(\mu_X)$  can be written as a linear combination  $\boldsymbol{x}=\sum_{i=1}^k\alpha_i\boldsymbol{m}_i^X$  for some coefficients  $\alpha_1,\ldots,\alpha_k$ . We can show that its coupled counterpart  $\boldsymbol{y}$  must be the same linear combination of the corresponding

QDMs in Y:

$$\begin{aligned} &\|\boldsymbol{y} - \sum_{i=1}^{k} \alpha_{i} \boldsymbol{m}_{i}^{Y} \|_{2}^{2} \\ &= \|\boldsymbol{y}\|_{2}^{2} + \|\sum_{i=1}^{k} \alpha_{i} \boldsymbol{m}_{i}^{Y} \|_{2}^{2} - 2 \sum_{i=1}^{k} \alpha_{i} \boldsymbol{y}^{\top} \boldsymbol{m}_{i}^{Y} \\ &= \|\boldsymbol{y}\|_{2}^{2} + \sum_{i=1}^{k} \sum_{j=1}^{k} \alpha_{i} \alpha_{j} \boldsymbol{m}_{i}^{Y^{\top}} \boldsymbol{m}_{j}^{Y} - 2 \sum_{i=1}^{k} \alpha_{i} \boldsymbol{x}^{\top} \boldsymbol{m}_{i}^{X} \\ &= \|\boldsymbol{y}\|_{2}^{2} + \sum_{i=1}^{k} \sum_{j=1}^{k} \alpha_{i} \alpha_{j} \boldsymbol{m}_{i}^{X^{\top}} \boldsymbol{m}_{j}^{X} - 2 \boldsymbol{x}^{\top} \sum_{i=1}^{k} \alpha_{i} \boldsymbol{m}_{i}^{X} \\ &= \|\boldsymbol{y}\|_{2}^{2} + \boldsymbol{x}^{\top} \boldsymbol{x} - 2 \boldsymbol{x}^{\top} \boldsymbol{x} = 0, \end{aligned}$$

where the last equality uses  $\|x\|_2^2 = \|y\|_2^2$ . This implies  $y = \sum_{i=1}^k \alpha_i m_i^Y$ .

Finally, for any two pairs  $(\boldsymbol{x}_1, \boldsymbol{y}_1), (\boldsymbol{x}_2, \boldsymbol{y}_2) \in \operatorname{supp}(\pi)$ , let  $\boldsymbol{x}_1 = \sum_{i=1}^k \alpha_i \boldsymbol{m}_i^X$  and  $\boldsymbol{x}_2 = \sum_{j=1}^k \beta_j \boldsymbol{m}_j^X$ . It follows that their distances are preserved:

$$||\mathbf{x}_{1} - \mathbf{x}_{2}||_{2}^{2} - ||\mathbf{y}_{1} - \mathbf{y}_{2}||_{2}^{2}$$

$$= (||\mathbf{x}_{1}||_{2}^{2} - ||\mathbf{y}_{1}||_{2}^{2}) + (||\mathbf{x}_{2}||_{2}^{2} - ||\mathbf{y}_{2}||_{2}^{2}) - (2\mathbf{x}_{1}^{\top}\mathbf{x}_{2} - 2\mathbf{y}_{1}^{\top}\mathbf{y}_{2})$$

$$= -2\sum_{i=1}^{k} \sum_{j=1}^{k} \alpha_{i}\beta_{j} (\mathbf{m}_{i}^{X^{\top}}\mathbf{m}_{j}^{X} - \mathbf{m}_{i}^{Y^{\top}}\mathbf{m}_{j}^{Y}) = 0.$$

Thus,  $d_X(\boldsymbol{x}_1, \boldsymbol{x}_2) = d_Y(\boldsymbol{y}_1, \boldsymbol{y}_2)$  holds for  $\pi \otimes \pi$ -a.e. pairs, fulfilling condition (ii) and completing the proof.

#### B.2 PROOF OF THEOREM 2

This section provides the proof for Theorem 2. We begin by establishing the necessary notation and preliminary lemmas. Let  $r_q = F_{\|X\|_2}^{-1}(q)$  be the true q-th quantile of the norm distribution, and let  $\hat{r}_q = \hat{F}_{\|X_n\|_2}^{-1}(q)$  be its empirical counterpart estimated from the sample set  $\mathbb{X}_n = \{x_1, \dots, x_n\}$ . For brevity, we denote the weight functions as  $w_i^X(x) = w^X(x, q_i)$  and its empirical version as  $\hat{w}_i^{X_n}(x) = e^{-\sigma(\|x\|_2 - \hat{r}_q)^2}$ .

Our proof relies on the well-established convergence rates for the Wasserstein distance, summarized in the following lemma.

Lemma 1 (Convergence Rate of Wasserstein Distance (Fournier & Guillin, 2015)) For a distribution  $\mu$  defined on a vector measure space  $(X, \mathcal{B}(X))$  and its empirical version  $\mu_n$ , if the q-th moment  $\mathcal{M}_q(\mu)$  is finite, i.e.,  $\mathcal{M}_q(\mu) = \int_X \|\mathbf{x}\|^q d\mu(x) < \infty$ , then the following holds:

$$\begin{split} &\mathbb{E}\left(\mathcal{W}_{p}\left(\mu_{n},\mu\right)\right) \leq CM_{q}^{p/q}(\mu) \\ &\times \begin{cases} n^{-1/2} + n^{-(q-p)/q} & \text{if } p > d/2 \text{ and } q \neq 2p \\ n^{-1/2}\log(1+n) + n^{-(q-p)/q} & \text{if } p = d/2 \text{ and } q \neq 2p \\ n^{-p/d} + n^{-(q-p)/q} & \text{if } p \in (0,d/2) \text{ and } q \neq d/(d-p) \end{cases} \end{split}$$

As discussed in (Fournier & Guillin, 2015), for a sufficiently large moment order q, the term  $n^{-(q-p)/q}$  becomes negligible compared to the leading term. For the common case of p=2, this holds for  $q\geq 4$ . We will assume this condition holds in our subsequent analysis.

Next, we state the standard convergence rate for empirical quantiles.

**Lemma 2 (Convergence Rate of Quantiles)** Let  $F_X$  be the CDF of a distribution  $\mu_X$ , and assume it is continuous on  $r_q$  and has a strictly positive density. Then, the mean squared error of the empirical quantile converges as follows:

$$\mathbb{E}|r_q - \hat{r}_q|^2 = O(n^{-1}).$$

This is a classical result in asymptotic theory, which can be found in (Serfling, 1980, Chap 2.3).

Furthermore, our proof requires the following bounds. The first lemma ensures that the denominator in the QDM definition is well-behaved.

**Lemma 3** If  $\mathbb{E}||X||_2 < \infty$ , then for any given quantile level  $q_i$ , the expected weight is bounded away from zero, i.e.,  $\frac{1}{\mathbb{E}_{\mu_X}(w_i^X(X))} < \infty$ .

**Proof.** Since  $\mathbb{E}_{\mu_X}(\|X\|_2) < \infty$ , for any given  $p_0 < 1$ , there exists an  $M_0$  such that  $\Pr(\|X\|_2 \le M_0) \ge p_0$ . Therefore, we have:

$$\mathbb{E}_{\mu_X}(w_i^X(X)) \ge \int_{\|\mathbf{x}\|_2 < M_0} e^{-\sigma(x - r_{q_i})^2} d\mu_X(x) \ge \int_{\|\mathbf{x}\|_2 < M_0} e^{-\sigma \max\{x, r_{q_i}\}^2} d\mu_X(x)$$

$$\ge \int_{\|\mathbf{x}\|_2 < M_0} e^{-\sigma \max\{M_0, r_{q_i}\}^2} d\mu_X(x) \ge p_0 e^{-\sigma \max\{M_0, r_{q_i}\}^2}.$$

This implies that the expectation is strictly positive, and thus  $\frac{1}{\mathbb{E}_{\mu_X}(w_i^X(X))} < \infty$ .

**Lemma 4** For any quantile level  $q_i$ , let the normalized empirical weights be  $\hat{p}_j = \frac{\hat{w}_i^{\chi_n}(\mathbf{x}_j)}{\sum_{j'=1}^n \hat{w}_i^{\chi_n}(\mathbf{x}'_j)}$ . Then for any power  $1 \leq \alpha < \infty$ , we have the following bound:

$$\sum_{j=1}^{n} \mathbb{E}(\hat{p}_{j}^{\alpha})^{1/\alpha} < \infty.$$

**Proof.** Let  $M_{\alpha}(\hat{w}_i) = \mathbb{E}(\hat{w}_i^{\mathbb{X}_n}(\boldsymbol{x}_j)^{\alpha})$ . Since  $\hat{w}_i^{\mathbb{X}_n}(\boldsymbol{x}_j) < 1$ , it is clear that  $M_{\alpha}(\hat{w}_i) < 1$ . Define the sample mean of the weights as  $\bar{w}(r) = \frac{1}{n} \sum_{j=1}^n e^{-\sigma(\|\boldsymbol{x}_j\|_2 - r)^2}$ . We then have  $\mathbb{E}(\bar{w}(\hat{r}_{q_i})) = M_1(\hat{w}_i)$ . Now, consider the following probability:

$$\mathbb{P}\big(\bar{w}_i(\hat{r}_{q_i}) < \tfrac{1}{2}M_1(\hat{w}_i)\big) \leq \underbrace{\mathbb{P}\big(\bar{w}_i(r_{q_i}) < \tfrac{3}{4}M_1(\hat{w}_i)\big)}_{\mathbb{P}_1} + \underbrace{\mathbb{P}\big(|\bar{w}_i(r_{q_i}) - \bar{w}_i(r_{q_i})| > \tfrac{1}{4}M_1(\hat{w}_i)\big)}_{\mathbb{P}_2}.$$

For the first term,  $\mathbb{P}_1$ , we note that  $\mathbb{E}(\bar{w}_i(r_{q_i})) = M_1(\hat{w}_i)$  and the terms in the sum are i.i.d. Applying Hoeffding's inequality yields:

$$\begin{split} \mathbb{P}_{1} = & \mathbb{P} \big( \mathbb{E} (\bar{w}_{i}(r_{q_{i}})) - \bar{w}_{i}(r_{q_{i}}) > \frac{n}{4} M_{1}(\hat{w}_{i}) \big) \\ = & \mathbb{P} \Big( \sum_{j=1}^{n} \mathbb{E} (w_{i}^{X}(\boldsymbol{x}_{j})) - \sum_{j=1}^{n} w_{i}^{X}(\boldsymbol{x}_{j}) > \frac{n}{4} M_{1}(\hat{w}_{i}) \Big) \\ \leq & \exp(-\frac{n}{8} M_{1}^{2}(\hat{w}_{i})). \end{split}$$

For the second term,  $\mathbb{P}_2$ , we first bound the difference by applying the Mean Value Theorem:

$$\left| \bar{w}_{i}(r_{q_{i}}) - \bar{w}_{i}(r_{q_{i}}) \right| \leq \frac{1}{n} \sum_{j=1}^{n} \left| e^{-\sigma(\|\mathbf{x}_{j}\|_{2} - r_{q_{i}})} - e^{-\sigma(\|\mathbf{x}_{j}\|_{2} - \hat{r}_{q_{i}})} \right|$$

$$\leq \frac{1}{n} \sum_{j=1}^{n} \left| 2\sigma(\|\mathbf{x}_{j}\|_{2} - r_{j0}) e^{-\sigma(\|\mathbf{x}_{j}\|_{2} - r_{j0}0)^{2}} \right| \left| r_{q_{i}} - \hat{r}_{q_{i}} \right|$$

$$\leq \sqrt{2\sigma} e^{-1/2} \left| r_{q_{i}} - \hat{r}_{q_{i}} \right| \quad \text{(since } \sup_{t \geq 0} 2\sigma t e^{-\sigma t^{2}} \leq \sqrt{2\sigma} e^{-1/2} \right).$$
(11)

Consequently, we have:

$$\mathbb{P}_2 \leq \mathbb{P}\Big(|r_{q_i} - \hat{r}_{q_i}| > \sqrt{\frac{e}{32\sigma}} M_1(\hat{w}_i)\Big) \leq 2\exp(-2n\delta),$$

where  $\delta = \min\{F_X(r_{p_i} + \sqrt{\frac{e}{32\sigma}}M_1(\hat{w}_i)) - p_i, p_i - F_X(r_{p_i} - \sqrt{\frac{e}{32\sigma}}M_1(\hat{w}_i))\} > 0$ . The final inequality follows from a standard result on the concentration of empirical quantiles, see, e.g.,

 (Serfling, 1980, Theorem 2.3.2). By the symmetry of the i.i.d. samples,  $\mathbb{E}(\hat{p}_j^{\alpha}) = \mathbb{E}(\hat{p}_l^{\alpha})$  for all  $j, l \in \{1, \dots, n\}$ . Thus we can write:

$$\begin{split} \sum_{j=1}^{n} \mathbb{E}(\hat{p}_{j}^{\alpha})^{1/\alpha} &= n \mathbb{E}(\hat{p}_{1}^{\alpha})^{1/\alpha} \\ &= \mathbb{E}\left(\frac{\hat{w}_{i}^{\mathbb{X}_{n}}(\boldsymbol{x}_{1})^{\alpha}}{\bar{w}_{i}(\hat{r}_{q_{i}})^{\alpha}}\right)^{1/\alpha} \\ &\leq \mathbb{E}_{\left\{\bar{w}_{i}(\hat{r}_{q_{i}}) \geq \frac{1}{2}M_{1}(\hat{w}_{i})\right\}} \left(\frac{\hat{w}_{i}^{\mathbb{X}_{n}}(\boldsymbol{x}_{1})^{\alpha}}{\bar{w}_{i}(\hat{r}_{q_{i}})^{\alpha}}\right)^{1/\alpha} + n \mathbb{P}(\bar{w}_{i}(\hat{r}_{q_{i}}) < \frac{1}{2}M_{1}(\hat{w}_{i})) \\ &\leq \frac{2\mathbb{E}(\hat{w}_{i}^{\mathbb{X}_{n}}(\boldsymbol{x}_{1})^{\alpha})^{1/\alpha}}{M_{1}(\hat{w}_{i}))} + n \exp(-\frac{n}{8}M_{1}^{2}(\hat{w}_{i})) + 2n \exp(-2n\delta) \\ &\leq \frac{2M_{\alpha}(\hat{w}_{i})^{1/\alpha}}{M_{1}(\hat{w}_{i}))} + \frac{8e^{-1}}{M_{1}^{2}(\hat{w}_{i})} + \frac{e^{-1}}{\delta} < \infty. \end{split}$$

The final line is finite, leveraging the result from Lemma 3 which ensures  $M_1(\cdot)$  is bounded away from zero. This completes the proof.

**Lemma 5** If the fourth moment of the norm is finite, i.e.,  $\mathbb{E}||X||_2^4 < \infty$ , then the expected squared norm of the empirical QDM is also finite:  $\mathbb{E}||\hat{m}_i^{\mathbb{X}_n}||_2^2 < \infty$ .

**Proof.** The proof follows from applying the Cauchy-Schwarz or Hölder's inequality and leveraging the result from Lemma 4.

$$\begin{split} \mathbb{E} \| \hat{m}_{i}^{\mathbb{X}_{n}} \|_{2}^{2} &= \sum_{j=1}^{n} \sum_{l=1}^{n} \mathbb{E} (\hat{p}_{j} \hat{p}_{l} \boldsymbol{x}_{j}^{\top} \boldsymbol{x}_{l}) \\ &\leq \sum_{j=1}^{n} \sum_{l=1}^{n} \mathbb{E} (\hat{p}_{j}^{4})^{1/4} \mathbb{E} (\hat{p}_{l}^{4})^{1/4} \mathbb{E} ((\boldsymbol{x}_{j}^{\top} \boldsymbol{x}_{l})^{2})^{1/2} \\ &\leq \sum_{j=1}^{n} \mathbb{E} (\hat{p}_{j}^{4})^{1/2} \mathbb{E} (\| \boldsymbol{x}_{j} \|_{2}^{4})^{1/2} + \sum_{j \neq l} \mathbb{E} (\hat{p}_{j}^{4})^{1/4} \mathbb{E} (\hat{p}_{l}^{4})^{1/4} \mathbb{E} (\| \boldsymbol{x}_{j} \|_{2}^{2} \| \boldsymbol{x}_{l} \|_{2}^{2})^{1/2} \\ &= \mathbb{E} (\| \boldsymbol{X} \|_{2}^{4})^{1/2} \sum_{j=1}^{n} \mathbb{E} (\hat{p}_{j}^{4})^{1/2} + \mathbb{E} (\| \boldsymbol{X} \|_{2}^{2}) \left( \sum_{j=1}^{n} \mathbb{E} (\hat{p}_{j}^{4})^{1/2} \right)^{2} < \infty. \end{split}$$

**Lemma 6** If  $\mathbb{E}||X||_2^4 < \infty$ , then the expectation of the product of the empirical QDM norm, a sample norm, and the quantile error converges at the rate of  $O(n^{-1/2})$ :  $\mathbb{E}(\|\hat{m}_i^{\mathbb{X}_n}\|_2 \|\mathbf{x}_1\| \|r_{q_i} - \hat{r}_{q_i}\|) = O(n^{-1/2})$ .

**Proof.** This proof also relies on the Cauchy-Schwarz inequality and the results from Lemma 2 and Lemma 4.

$$\mathbb{E}(\|\hat{m}_{i}^{\mathbb{X}_{n}}\|_{2}\|\boldsymbol{x}_{1}\|_{2}|r_{q_{i}}-\hat{r}_{q_{i}}|) \leq \mathbb{E}(\hat{q}_{1}\|\boldsymbol{x}_{1}\|_{2}^{2}|r_{q_{i}}-\hat{r}_{q_{i}}|) + \sum_{j=2}^{n} \mathbb{E}(\hat{q}_{j}\|\boldsymbol{x}_{1}\|_{2}\|\boldsymbol{x}_{j}\|_{2}|r_{q_{i}}-\hat{r}_{q_{i}}|) \\
\leq \mathbb{E}(\|X\|_{2}^{4})^{1/2}\mathbb{E}(|r_{q_{i}}-\hat{r}_{q_{i}}|^{2})^{1/2} \\
+ \sum_{j=2}^{n} \mathbb{E}(\hat{q}_{j}^{4})^{1/4}\mathbb{E}(\|\boldsymbol{x}_{1}\|_{2}^{4}\|\boldsymbol{x}_{j}\|_{2}^{4})^{1/4}\mathbb{E}(|r_{q_{i}}-\hat{r}_{q_{i}}|^{2})^{1/2} \\
= \mathbb{E}(|r_{q_{i}}-\hat{r}_{q_{i}}|^{2})^{1/2}\mathbb{E}(\|X\|_{2}^{4})^{1/2}\left(1+\sum_{j=2}^{n} \mathbb{E}(\hat{q}_{j}^{4})^{1/4}\right).$$

From Lemma 2, we know that  $\mathbb{E}(|r_{q_i}-\hat{r}_{q_i}|^2)^{1/2}=O(n^{-1/2})$ . Since  $\mathbb{E}(\|X\|_2^4)<\infty$  is assumed, and Lemma 4 ensures that  $\sum_{j=2}^n\mathbb{E}(\hat{q}_j^4)^{1/4}<\infty$ , the result follows.

**Proof of Theorem 2.** First, we use a property derived from the triangle inequality:

$$|d(\boldsymbol{x}, \boldsymbol{x}') - d(\boldsymbol{y}, \boldsymbol{y}')| \le |d(\boldsymbol{x}, \boldsymbol{x}') - d(\boldsymbol{x}, \boldsymbol{y}')| + |d(\boldsymbol{x}, \boldsymbol{y}') - d(\boldsymbol{y}, \boldsymbol{y}')|$$

$$\le d(\boldsymbol{x}, \boldsymbol{y}) + d(\boldsymbol{x}', \boldsymbol{y}').$$

This allows us to bound the QDOT distance.

where  $m_0^X = 0$ ,  $m_0^Y = 0$ , and the constants  $C_0 = (k+1)2^{p-1}$  and  $C_1 = 2^{p-1}$  are finite. From (Fournier & Guillin, 2015, Theorem 1), we have the bound for the Wasserstein Discrepancy between the true measure and its empirical version:

$$\mathbb{E}(\mathcal{W}_p(\mu_X, \mu_n)^p) \le C n^{-p/d}.$$

Next, we analyze the QDM Discrepancy term,  $d(\boldsymbol{m}_i^X, \hat{\boldsymbol{m}}_i^{\mathbb{X}_n})$ .

$$\begin{split} d(\boldsymbol{m}_{i}^{X}, \hat{\boldsymbol{m}}_{i}^{\mathbb{X}_{n}}) &= \left\| \frac{\frac{1}{n} \sum_{j=1}^{n} \hat{w}_{i}^{\mathbb{X}_{n}}(\boldsymbol{x}_{j}) \boldsymbol{x}_{j}}{\frac{1}{n} \sum_{j=1}^{n} \hat{w}_{i}^{\mathbb{X}_{n}}(\boldsymbol{x}_{j})} - \frac{\mathbb{E}_{\mu_{X}}(w_{i}^{X}(X)X)}{\mathbb{E}_{\mu_{X}}(w_{i}^{X}(X))} \right\|_{2} \\ &= \left\| \frac{\frac{1}{n} \sum_{j=1}^{n} \hat{w}_{i}^{\mathbb{X}_{n}}(\boldsymbol{x}_{j}) \boldsymbol{x}_{j} \mathbb{E}_{\mu_{X}}(w_{i}^{X}(X)) - \mathbb{E}_{\mu_{X}}(w_{i}^{X}(X)X) \frac{1}{n} \sum_{j=1}^{n} \hat{w}_{i}^{\mathbb{X}_{n}}(\boldsymbol{x}_{j})}{\frac{1}{n} \sum_{j=1}^{n} \hat{w}_{i}^{\mathbb{X}_{n}}(\boldsymbol{x}_{j}) \mathbb{E}_{\mu_{X}}(w_{i}^{X}(X))} \right\|_{2} \\ &\leq C_{2} \left\| \mathbb{E}_{\mu_{X}}\left(w_{i}^{X}(X)\right) - \frac{1}{n} \sum_{j=1}^{n} w_{i}^{X}(x_{j}) \right\|_{2} \\ &+ C_{2} \left\| \hat{m}_{i}^{\mathbb{X}_{n}} \right\|_{2} \left\| \frac{1}{n} \sum_{j=1}^{n} \hat{w}_{i}^{\mathbb{X}_{n}}(x_{j}) x_{j} - \mathbb{E}_{\mu_{X}}\left(w_{i}^{X}(X)X\right) \right\|_{2}. \end{split}$$

Here,  $C_2 = \frac{1}{\mathbb{E}_{\mu_X}(w_i^X(X))}$ . By Lemma 3, under the condition that  $q \ge 4$ ,  $C_2$  is a finite constant. We now focus on the two main terms. Term (I) can be decomposed as:

Term (I) 
$$\leq \left| \underbrace{\mathbb{E}_{\mu_X} \left( w_i^X(X) \right) - \frac{1}{n} \sum_{j=1}^n w_i^X(\boldsymbol{x}_j)}_{\text{Term (I.1)}} + \underbrace{\left| \frac{1}{n} \sum_{j=1}^n w_i^X(\boldsymbol{x}_j) - \frac{1}{n} \sum_{j=1}^n \hat{w}_i^{\mathbb{X}_n}(\boldsymbol{x}_j) \right|}_{\text{Term (I.2)}}.$$
 (14)

For Term (I.1), standard results for the mean of i.i.d. variables give:

$$\mathbb{E}(\text{Term (I.1)}) \leq \sqrt{\mathbb{E}\Big(\mathbb{E}_{\mu_X}\left(w_i^X(X)\right) - \frac{1}{n}\sum_{j=1}^n w_i^X(\boldsymbol{x}_j)\Big)^2} = \sqrt{\frac{1}{n}\operatorname{Var}\left(w_i^X(\boldsymbol{x}_1)\right)} = O(n^{-1/2}).$$

Using the bound from equation 11, we have:

Term (I.2) 
$$\leq \sqrt{2\sigma}e^{-1/2} |r_{q_i} - \hat{r}_{q_i}|$$
.

Since  $\mathbb{E}|r_{q_i} - \hat{r}_{q_i}| = O(n^{-1/2})$ , it follows that  $\mathbb{E}(\text{Term (I)}) = O(n^{-1/2})$ . For Term (II), we use a similar decomposition:

$$\operatorname{Term}\left(\Pi\right) \leq \underbrace{\left\|\hat{m}_{i}^{\mathbb{X}_{n}}\right\|_{2} \left\|\frac{1}{n} \sum_{j=1}^{n} \hat{w}_{i}^{\mathbb{X}_{n}}(x_{j}) \, x_{j} - \frac{1}{n} \sum_{j=1}^{n} w_{i}^{X}(x_{j}) \, x_{j} \, \right\|_{2}}_{\operatorname{Term}\left(\Pi, 1\right)} + \underbrace{\left\|\hat{m}_{i}^{\mathbb{X}_{n}}\right\|_{2} \left\|\sum_{j=1}^{n} w_{i}^{X}(x_{j}) \, x_{j} - \mathbb{E}_{\mu_{X}}\!\!\left(w_{i}^{X}(X)X\right) \right\|_{2}}_{\operatorname{Term}\left(\Pi, 2\right)}.$$

For Term (II.1), we have:

$$\begin{split} \text{Term (II.1)} = & \left\| \hat{m}_i^{\mathbb{X}_n} \right\|_2 \cdot \left\| \frac{1}{n} \sum_{j=1}^n \boldsymbol{x}_j \left( e^{-\sigma(\|\boldsymbol{x}_j\|_2 - r_{q_i})} - e^{-\sigma(\|\boldsymbol{x}_j\|_2 - \hat{r}_{q_i})} \right) \right\|_2 \\ \leq & \frac{1}{n} \left\| \hat{m}_i^{\mathbb{X}_n} \right\|_2 \sum_{j=1}^n \left( \|\boldsymbol{x}_j\|_2 \left| e^{-\sigma(\|\boldsymbol{x}_j\|_2 - r_{q_i})} - e^{-\sigma(\|\boldsymbol{x}_j\|_2 - \hat{r}_{q_i})} \right| \right) \\ \leq & \frac{\sqrt{2\sigma}e^{-1/2}}{n} \sum_{j=1}^n \left\| \hat{m}_i^{\mathbb{X}_n} \right\|_2 \|\boldsymbol{x}_j\|_2 |r_{q_i} - \hat{r}_{q_i}|. \end{split}$$

From Lemma 6, we know  $\mathbb{E}(\|\hat{m}_i^{\mathbb{X}_n}\|_2\|\boldsymbol{x}_j\|_2|r_{q_i}-\hat{r}_{q_i}|)=O(n^{-1/2})$ , which implies  $\mathbb{E}(\text{Term (II.1)})=O(n^{-1/2})$ . For Term (II.2), since  $w_i^X(x)\leq 1$ , we have  $\mathbb{E}\|w(X)X\|_2^2<\infty$ . Using the Cauchy-Schwarz inequality, we get:

$$\mathbb{E}(\text{Term (II.2)}) \leq \sqrt{\mathbb{E} \|\hat{m}_i^{\mathbb{X}_n}\|_2^2 \cdot \mathbb{E} \Big(\mathbb{E}_{\mu_X} \|w_i^X(X)X\Big) - \frac{1}{n} \sum_{j=1}^n w_i^X(\boldsymbol{x}_j) \boldsymbol{x}_j \|_2^2 \Big)}$$
$$= \sqrt{\mathbb{E} \|\hat{m}_i^{\mathbb{X}_n}\|_2^2} \sqrt{\frac{1}{n} \operatorname{tr} \operatorname{Cov} \left(w_i^X(\boldsymbol{x}_1) \boldsymbol{x}_1\right)} = O(n^{-1/2}).$$

Thus, we also have  $\mathbb{E}(\operatorname{Term}(\Pi)) = O(n^{-1/2})$ . In summary, we have shown that  $\mathbb{E}(d(\boldsymbol{m}_i^X, \hat{\boldsymbol{m}}_i^{\mathbb{X}_n})) = O(n^{-1/2})$ , which means the convergence rate for the QDM Discrepancy is  $O(n^{-p/2})$ . For typical cases where  $d \geq 2$ , this rate is faster than the Wasserstein rate, so the overall convergence is dominated by the Wasserstein Discrepancy. Therefore, we obtain the final rate:

$$\mathbb{E}(\mathcal{QD}_p(\mathcal{X}, \mathcal{X}_n)) = O(n^{-1/d}).$$

#### B.3 Proof of Theorem 3

The proof for the Identity of Indiscernibles is analogous to that of Theorem 1, and symmetry holds trivially by definition. Therefore, we focus on proving the Triangle Inequality. We begin by introducing the following well-known Gluing Lemma.

**Lemma 7 (Gluing Lemma)** For three mm-spaces  $\mathcal{X}, \mathcal{Y}, \mathcal{Z}$  with corresponding measures  $\mu_X, \mu_Y, \mu_Z$ , and given couplings  $\pi_{X,Z} \in \Pi(\mu_X, \mu_Z)$  and  $\pi_{Y,Z} \in \Pi(\mu_Y, \mu_Z)$ , there exists a joint coupling  $\pi_{X,Y,Z} \in \Pi(\mu_X, \mu_Y, \mu_Z)$  such that its marginals satisfy

$$(\text{proj}_{(X,Z)})_{\#}\pi_{X,Y,Z} = \pi_{X,Z}, \quad (\text{proj}_{(Y,Z)})_{\#}\pi_{X,Y,Z} = \pi_{Y,Z},$$

and its (X,Y)-marginal,  $(\operatorname{proj}_{(X,Y)})_{\#}\pi_{X,Y,Z}$ , is a coupling in  $\Pi(\mu_X,\mu_Y)$ .

A proof of this lemma can be found in (Villani et al., 2008, Chap 1).

**Proof of Triangle Inequality.** Consider three mm-spaces  $\mathcal{X}=(X,d_X,\mu_X), \ \mathcal{Y}=(Y,d_Y,\mu_Y),$  and  $\mathcal{Z}=(Z,d_Z,\mu_Z).$  Let  $\pi_{X,Y}^\star\in\Pi(\mu_X,\mu_Y), \pi_{X,Z}^\star\in\Pi(\mu_X,\mu_Z),$  and  $\pi_{Y,Z}^\star\in\Pi(\mu_Y,\mu_Z)$  denote the respective pairwise optimal couplings for the IQDOT distance. According to Lemma 7, we can find a joint coupling  $\pi_{X,Y,Z}\in\Pi(\mu_X,\mu_Y,\mu_Z)$  such that

$$(\text{proj}_{(X,Z)})_{\#}\pi_{X,Y,Z} = \pi_{X,Z}^{\star}, \quad (\text{proj}_{(Y,Z)})_{\#}\pi_{X,Y,Z} = \pi_{Y,Z}^{\star}.$$

Let  $\pi_{X,Y} := (\operatorname{proj}_{(X,Y)})_{\#} \pi_{X,Y,Z}$ . By construction,  $\pi_{X,Y} \in \Pi(\mu_X, \mu_Y)$ . We can now bound the IQDOT distance as follows:

$$\begin{split} \mathcal{IQD}_{p}(\mathcal{X},\mathcal{Y}) = & \mathbb{E}_{(0,1) \times \pi_{X,Y}^{\star}} \left( |\phi^{X}(x,q) - \phi^{Y}(y,q)|^{p} \right)^{1/p} \\ \leq & \mathbb{E}_{(0,1) \times \pi_{X,Y}} \left( |\phi^{X}(x,q) - \phi^{Y}(y,q)|^{p} \right)^{1/p} \\ \leq & \mathbb{E}_{(0,1) \times \pi_{X,Y,Z}} \left( (|\phi^{X}(x,q) - \phi^{Z}(z,q)| + |\phi^{Z}(z,q) - \phi^{Y}(y,q|)^{p} \right)^{1/p} \\ \leq & \mathbb{E}_{(0,1) \times \pi_{X,Z}^{\star}} \left( |\phi^{X}(x,q) - \phi^{Z}(z,q)|^{p} \right)^{1/p} + \mathbb{E}_{(0,1) \times \pi_{Y,Z}^{\star}} \left( |\phi^{Z}(z,q) - \phi^{Y}(y,q)|^{p} \right)^{1/p} \\ = & \mathcal{IQD}_{p}(\mathcal{X}, \mathcal{Z}) + \mathcal{IQD}_{p}(\mathcal{Z}, \mathcal{Y}), \end{split}$$

where the last inequality follows from Minkowski's inequality. This completes the proof.  $\Box$ 

#### C IMPLEMENTATION DETAILS

#### C.1 Details of Algorithm 1

The following provides further details on some of the operations and parameter choices in Algorithm 1.

**Data Initialization.** While optional, we recommend centering the data as a pre-processing step. The intuition for this is rooted in both theory and practice. Theoretically, it is known that a surjective isometry between real normed spaces that maps the origin to the origin also preserves the mean of a distribution (Narici & Beckenstein, 2010, Theorem 9.1.2). This implies that the distributions of norms for two centered, isometric spaces must be identical, which aligns with the foundational principles of our method. From a practical standpoint, centering is a common alignment technique used in numerical implementations of Gromov-Wasserstein (Flamary et al., 2021). However, this initialization step is not necessary for intra-space comparisons or in cases where a correspondence between the origins is already known.

Quantile Level Vector q. The quantile levels in the vector q are theoretically chosen from the open interval (0,1). In practice, we select these levels by taking equispaced points within a truncated interval  $[\delta,1-\delta]$ . The default value for  $\delta$  is 0.1, which avoids numerical instability at the extreme tails of the distribution. For larger datasets where the empirical quantiles are more stable, a smaller  $\delta$  may be used to satisfy the theoretical conditions more closely.

Choice of the Number of Quantiles k. A large number of quantiles, k, is often not required to achieve strong performance. For low-dimensional data (e.g., 3D point clouds), the trilateration condition is readily satisfied, and a small value such as k=5 can be sufficient. For high-dimensional data, which often exhibits low-rank structure, a value of k smaller than the ambient dimension can also be effective. While a larger k can lead to improved numerical stability, we find that k=50 provides a robust default choice across most applications.

**Bandwidth Parameter**  $\sigma$ . The choice of the bandwidth parameter  $\sigma$  in the Gaussian kernel is related to the number of support points, n. A very small  $\sigma$  will cause the weights to concentrate on a single point, while a very large  $\sigma$  will lead to nearly uniform weights across all samples. Empirically, we recommend choosing  $\sigma$  such that the maximum weight is on the order of the average weight, i.e.,  $\max_i \{w_i\} \sim n^{-1/2}$ . In practice, a default value of  $\sigma = 10$  provides robust and effective performance across most of our experiments.

We analyze the robustness of our method to these parameters through a toy example detailed in Appendix D.1.

#### C.2 ALGORITHM OF IQDOT

1134

1135 1136

1137

1138 1139

1157 1158 1159

1160

1161

1162

1163

1164 1165

1166 1167

1168 1169

1177

1178

1179

1180

1181

1182

1183

1184

1185

1186

1187

For numerical inputs, given by discrete sample matrices  $\mathbf{X} \in \mathbb{R}^{n \times d}$  and  $\mathbf{Y} \in \mathbb{R}^{m \times q}$  with corresponding probability vectors  $p^X \in \Delta^{n-1}$  and  $p^Y \in \Delta^{m-1}$ , we now detail the numerical implementation of the IQDOT, as introduced in Definition 4. The procedure is summarized in Algorithm 2.

#### Algorithm 2 Intergal-QDOT

```
1140
1141
                                  Require: X, Y, p^X, p^Y, q
1142
                                      1: Initialize the data matrices X and Y.
1143
                                     2: Compute the sample norms of \mathbf{X} and \mathbf{Y}: \phi_0^{\mathbf{X}} \leftarrow (\|\mathbf{x}_i\|)_{1 \leq i \leq n}, \phi_0^{\mathbf{Y}} \leftarrow (\|\mathbf{y}_j\|)_{1 \leq j \leq m}.

3: Compute the \mathbf{q}-quantiles of the sample norms: \qquad \qquad \qquad \triangleright \mathcal{O}(n \log n + m \log m)
\mathbf{r}^{\mathbf{X}} \leftarrow (F_{\|\mathbf{X}\|}^{-1}(q_1), \dots, F_{\|\mathbf{X}\|}^{-1}(q_k)), \quad \mathbf{r}^{\mathbf{Y}} \leftarrow (F_{\|\mathbf{Y}\|}^{-1}(q_1), \dots, F_{\|\mathbf{Y}\|}^{-1}(q_k)).
1144
1145
1146
                                                                                                                                                                                                                                                                                                                                                             \triangleright \mathcal{O}(knd + kmq)
                                      4: for i \leftarrow 1 to k do
1147
                                                                Compute the quantile weights:  \boldsymbol{w}_{ij}^{\mathbf{X}} \leftarrow \frac{p_{j}^{Y} \exp\{-\sigma(\boldsymbol{d}_{j}^{\mathbf{X}} - \boldsymbol{r}_{i}^{\mathbf{X}})^{2}\}}{\sum_{j'=1}^{n} p_{j'}^{Y} \exp\{-\sigma(\boldsymbol{d}_{j'}^{\mathbf{X}} - \boldsymbol{r}_{i}^{\mathbf{X}})^{2}\}}, \quad \boldsymbol{w}_{ij}^{\mathbf{Y}} \leftarrow \frac{p_{j}^{Y} \exp\{-\sigma(\boldsymbol{d}_{j}^{\mathbf{Y}} - \boldsymbol{r}_{i}^{\mathbf{Y}})^{2}\}}{\sum_{j'=1}^{m} p_{j'}^{Y} \exp\{-\sigma(\boldsymbol{d}_{j'}^{\mathbf{Y}} - \boldsymbol{r}_{i}^{\mathbf{Y}})^{2}\}}  Compute the q_{i}-quantile means: \boldsymbol{m}_{i}^{X} \leftarrow \mathbf{X}^{\top} \boldsymbol{w}_{i}^{X}, \quad \boldsymbol{m}_{i}^{Y} \leftarrow \mathbf{Y}^{\top} \boldsymbol{w}_{i}^{Y}
1148
1149
1150
                                      6:
1151
                                                                Compute the distances to quantity means: \phi_i^{\mathbf{X}} \leftarrow (\|\boldsymbol{x}_j - \boldsymbol{m}_i^X\|)_{1 \leq j \leq n}, \quad \phi_i^{\mathbf{Y}} \leftarrow (\|\boldsymbol{y}_j - \boldsymbol{m}_i^Y\|)_{1 \leq j \leq m}
Calculate the 1D-\mathcal{W}_p loss: \mathcal{L}_i \leftarrow \mathcal{W}_p\left((\phi_i^{\mathbf{X}})_{\#}\mu_{\mathbf{X}}, (\phi_i^Y)_{\#}\mu_{\mathbf{Y}}\right)
1152
1153
                                                                                                                                                                                                                                                                                                                                        \triangleright \mathcal{O}(n \log n + m \log m)
1154
                                      9: end for
                                  10: Set the final loss : \mathcal{IQD}_p \leftarrow \frac{1}{k} \sum_{i=1}^k \mathcal{L}_i
1155
1156
                                  11: return \mathcal{IQD}_p
```

**Computational Cost.** Similar to the analysis for QDOT, the pre-computation of QDMs and QD-MDs is highly efficient. The final one-dimensional Wasserstein distances can also be computed in  $O(n \log n)$  time by North-West corner rule(Peyré et al., 2019). Consequently, the overall time complexity of the IQDOT algorithm is  $\mathcal{O}(n \log n)$ , assuming k and the data dimensions are small constants.

#### D ADDITIONAL EXPERIMENT RESULTS

#### TOY EXAMPLE AND PARAMETER ANALYSIS

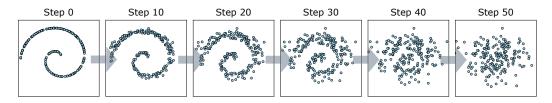


Figure 7: Illustration of the Wasserstein flow from a spiral point cloud to a random Gaussian noise distribution over 50 steps.

To provide an intuitive visualization of QDOT's behavior and to analyze its sensitivity to key parameters, we present a toy example. As illustrated in Figure 7, the experiment consists of a Wasserstein flow sequence that interpolates between a 2D spiral distribution and a random Gaussian noise distribution. Each point cloud in the 50-step sequence comprises n=200 points. We conduct a series of tests to: (1) compare the dissimilarity trends produced by different geometric metrics as the spiral deforms into noise; (2) analyze the effect of varying the bandwidth parameter  $\sigma$ ; (3) analyze the effect of varying the quantile interval parameter  $\delta$ ; and (4) assess the impact of different initialization and scaling strategies.

The results, presented in Figure 8, reveal several key insights. First, when compared to baselines such as GW, EGW, SGW, and RISGW, both QDOT and IQDOT produce significantly smoother

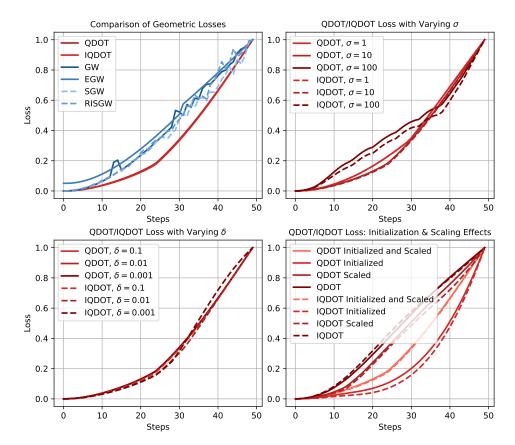


Figure 8: Results from the four experimental settings. The x-axis represents the step in the Wasserstein flow sequence, and the y-axis represents the computed dissimilarity (loss).

and more stable dissimilarity curves. Second, the analysis of the bandwidth parameter shows that for  $\sigma=1$  and  $\sigma=10$ , the resulting curves are stable and smooth. A larger value of  $\sigma=100$  introduces some volatility, which can be attributed to the interaction between a high bandwidth and the relatively small sample size. Third, the dissimilarity curve is largely insensitive to the choice of the quantile interval parameter  $\delta$  in this experiment, indicating good robustness. Finally, while different initialization and scaling strategies affect the absolute values of the dissimilarity, they all produce smooth, monotonically increasing curves that preserve the overall trend, demonstrating the robustness of the underlying geometric representation.

#### D.2 CROSS SPACE TASKS

**Evaluation Metrics.** For the numerical experiments, the inputs are a 3D point cloud sequence, represented by sample matrices  $\mathbf{X} \in \mathbb{R}^{n \times 3}$ , and a reference shape. The reference is provided both in its original 3D form,  $\mathbf{Y} \in \mathbb{R}^{m \times 3}$ , and as a 2D projection,  $\hat{\mathbf{Y}} \in \mathbb{R}^{m \times 2}$ . Both distributions have corresponding probability vectors,  $\mathbf{p}^{\mathbf{X}}$  and  $\mathbf{p}^{\mathbf{Y}}$ , we assume access to a ground-truth alignment between the sequence frame  $\mathbf{X}$  and the reference shape  $\mathbf{Y}$ , denoted by a map  $\mathcal{T}$ , i.e.,  $\mathbf{X} = \mathcal{T}(\mathbf{Y})$ . Algorithm 1 computes a coupling matrix  $\mathbf{\Pi} \in \mathbb{R}^{n \times m}$  between the input distributions, which satisfies the marginal constraints  $\mathbf{\Pi} \mathbf{1}_m = \mathbf{p}^{\mathbf{X}}$  and  $\mathbf{\Pi}^{\top} \mathbf{1}_n = \mathbf{p}^{\mathbf{Y}}$ . To evaluate the quality of the resulting match, we define the following metrics.

The **Transformed Mean Squared Error** (**TMSE**) evaluates the alignment cost in the original 3D ambient space. Even when the coupling  $\Pi$  is computed between a 3D shape X and a 2D projection  $\hat{Y}$ , the error is measured by using the original 3D coordinates of the reference shape, Y:

$$TMSE = \sum_{i=1}^{n} \sum_{j=1}^{m} \Pi_{ij} d_X(\boldsymbol{x}_i, \mathcal{T}(\boldsymbol{y}_j));$$

The **Inlier Ratio** (**IR**) then measures the total probability mass of the coupling  $\Pi$  that is placed on geometrically correct matches. A match between  $x_i$  and  $y_j$  is considered correct if  $x_i$  is close to the ground-truth location corresponding to  $y_j$ :

$$IR = \sum_{i=1}^{n} \sum_{j=1}^{m} \Pi_{ij} \mathbf{1}(d_X(\boldsymbol{x}_i, \mathcal{T}(\boldsymbol{y}_j)) < \tau) \times 100\%;$$

where the threshold  $\tau$  is set to the 0.2-quantile of the intra-point distance distribution within  $\mathbf{X}$ , i.e.,  $\tau = F_{d_X}^{-1}(0.2)$ .

#### D.3 MOLECULE GENERATION

For the molecular generation experiment, we present the complete results in Table 5.

Table 5: Molecule Generation Results on QM9 Dataset.

Methods	Atom sta. (%)	Mol sta. (%)	Validity (%)	Uniqueness (%)	V * U (%)
Data	99.0	95.2	97.7	100.0	97.7
EDM	$98.70_{\pm 0.01}$	$86.56_{\pm0.27}$	$93.73_{\pm 0.12}$	$98.19_{\pm 0.14}$	$92.04_{\pm 0.03}$
$\rightarrow$ MSE <sub>0.1</sub>	$98.81_{\pm 0.04}$	$87.63_{\pm 0.21}$	$94.54_{\pm 0.15}$	$98.17_{\pm 0.08}$	$92.81_{\pm 0.20}$
$\rightarrow$ MSE <sub>0.2</sub>	$98.78_{\pm 0.01}$	$87.60_{\pm0.30}$	$94.18_{\pm 0.15}$	$98.25_{\pm 0.10}$	$92.54_{\pm 0.19}$
$\rightarrow$ MSE <sub>0.3</sub>	$98.68_{\pm 0.07}$	$86.04_{\pm 0.74}$	$93.83_{\pm 0.43}$	$98.47_{\pm 0.22}$	$92.39_{\pm 0.55}$
$\rightarrow$ MSE <sub>0.4</sub>	$98.80_{\pm 0.02}$	$87.44_{\pm 0.35}$	$94.59_{\pm 0.25}$	$98.38_{\pm 0.07}$	$93.05_{\pm 0.31}$
$\rightarrow$ QDOT <sub>0.1</sub>	$99.15_{\pm 0.01}$	$90.91_{\pm 0.21}$	$96.75_{\pm 0.33}$	$92.06_{\pm 0.15}$	$89.06_{\pm0.19}$
$\rightarrow$ QDOT <sub>0.2</sub>	$99.30_{\pm 0.03}$	$92.73_{\pm 0.21}$	$97.63_{\pm 0.11}$	$84.48_{\pm 0.59}$	$82.48_{\pm 0.49}$
$\rightarrow$ QDOT <sub>0.3</sub>	$99.35_{\pm 0.06}$	$93.29_{\pm 0.81}$	$97.80_{\pm 0.22}$	$83.05_{\pm 2.27}$	$81.21_{\pm 2.03}$
$\rightarrow$ QDOT <sub>0.4</sub>	$99.48_{\pm 0.01}$	$94.37_{\pm 0.18}$	$98.16_{\pm 0.13}$	$79.42_{\pm 0.27}$	$77.96_{\pm0.17}$
UniGEM	$98.90_{\pm 0.03}$	$89.40_{\pm 0.02}$	$94.58_{\pm 0.07}$	$98.07_{\pm 0.05}$	$92.75_{\pm 0.11}$
$\rightarrow$ MSE <sub>0.1</sub>	$98.91_{\pm 0.04}$	$89.00_{\pm 0.22}$	$95.00_{\pm 0.16}$	$98.04_{\pm 0.07}$	$93.13_{\pm 0.10}$
$\rightarrow$ MSE <sub>0.2</sub>	$98.65_{\pm 0.06}$	$85.27_{\pm 0.35}$	$93.93_{\pm 0.50}$	$98.05_{\pm 0.13}$	$92.15_{\pm 0.42}$
$\rightarrow$ MSE <sub>0.3</sub>	$99.00_{\pm 0.07}$	$89.48_{\pm 0.70}$	$95.16_{\pm0.16}$	$98.09_{\pm 0.23}$	$93.34_{\pm0.18}$
$\rightarrow$ MSE <sub>0.4</sub>	$99.03_{\pm 0.05}$	$90.09_{\pm 0.53}$	$95.17_{\pm 0.30}$	$98.06_{\pm0.13}$	$93.32_{\pm0.18}$
$\rightarrow$ QDOT <sub>0.1</sub>	$99.24_{\pm 0.03}$	$92.73_{\pm 0.09}$	$96.85_{\pm0.11}$	$96.45_{\pm 0.21}$	$93.42 \pm 0.23$
$\rightarrow$ QDOT <sub>0.2</sub>	$99.35_{\pm 0.07}$	$94.06_{\pm 0.43}$	$97.42_{\pm 0.28}$	$92.59_{\pm 0.13}$	$90.20_{\pm 0.36}$
$\rightarrow$ QDOT <sub>0.3</sub>	$99.44_{\pm0.01}$	$95.23_{\pm 0.18}$	$97.94_{\pm0.06}$	$85.64_{\pm0.28}$	$83.88_{\pm0.30}$
$\rightarrow$ QDOT <sub>0.4</sub>	$99.38_{\pm 0.02}$	$94.81_{\pm 0.25}$	$97.66_{\pm 0.24}$	$78.67_{\pm 0.38}$	$76.83_{\pm 0.39}$

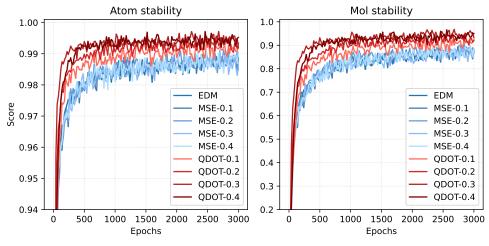
<sup>&</sup>lt;sup>1</sup> The notation  $\rightarrow$ QDOT<sub>0.2</sub> indicates that the objective function  $\mathcal{L} = 0.2\mathcal{L}_{QDOT} + 0.8\mathcal{L}_{EDM}$ .

As shown in Table 5, simply reducing the weight of the MSE loss does not yield significant improvements for the baseline models. We also observe that the benefits of the QDOT loss do not necessarily increase monotonically with its proportion in the final objective. Overall, these results suggest that the standard MSE loss primarily captures information about absolute atomic positions, while the QDOT loss focuses on the intrinsic molecular structure. The strongest performance is achieved when these two complementary objectives are fused, indicating that they provide a more comprehensive learning signal.

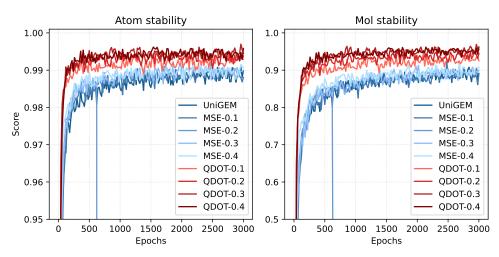
As illustrated in Figure 9, models trained with the QDOT loss achieve stable convergence on both Atom and Molecule Stability metrics in just 1000 epochs, for both the EDM and UniGEM architectures. In contrast, the baseline models and those fine-tuned by simply adjusting the MSE loss weight require significantly more training: approximately 2500 epochs on EDM and over 2000 epochs on UniGEM to reach a similar level of stability. This ablation confirms that simply reducing the MSE loss weight does not accelerate convergence, highlighting the efficiency gains are a direct result of the structural guidance provided by the QDOT loss.

#### THE USE OF LARGE LANGUAGE MODELS

We acknowledge the use of Large Language Models in the preparation of this manuscript. Their role was strictly limited to functioning as a tool for language polishing and improving readability.



(a) Stability trends across epochs on the EDM model



(b) Stability trends across epochs on the UniGEM model

Figure 9: Convergence trends of Atom Stability and Molecule Stability for various methods on the EDM and UniGEM models. The horizontal axis represents the number of training epochs, and the vertical axis represents the stability percentage. The comparison includes the baseline models and variants with different weights for the MSE and QDOT loss terms (MSE-0.1 to MSE-0.4 and QDOT-0.1) to QDOT-0.4).

We affirm that all definitions and theorems presented in this paper are the original intellectual contributions of the authors. The core code is intricately linked to the algorithm we propose, and the experiments were designed specifically to investigate the properties of our method. We take full responsibility for the integrity of this work, which is free from any form of academic misconduct.