

DEEPEMD: A TRANSFORMER-BASED FAST ESTIMATION OF THE EARTH MOVER’S DISTANCE

Anonymous authors

Paper under double-blind review

ABSTRACT

The Earth Mover’s Distance (EMD) is the measure of choice between point clouds. However the computational cost to compute it makes it prohibitive as a training loss, and the standard approach is to use a surrogate such as the Chamfer distance. We propose an attention-based model to compute an accurate approximation of the EMD that can be used as a training loss for generative models. To get the necessary accurate estimation of the gradients we train our model to explicitly compute the matching between point clouds instead of EMD itself. We cast this new objective as the estimation of an attention matrix that approximates the ground truth matching matrix. Experiments show that this model provides an accurate estimate of the EMD and its gradient with a wall clock speed-up of more than two orders of magnitude with respect to the exact Hungarian matching algorithm and one order of magnitude with respect to the standard approximate Sinkhorn algorithm, allowing in particular to train a point cloud VAE with the EMD itself. Extensive evaluation show the remarkable behaviour of this model when operating out-of-distribution, a key requirement for a distance surrogate. Finally, the model generalizes very well to point clouds during inference several times larger than during training.

1 INTRODUCTION

The *earth mover’s distance* (EMD), also known as *Wasserstein distance* is a distance between distributions that is defined as the minimum total of mass-time-distance displacement needed to transform one distribution to the other. In the case of uniform distributions over a finite number of points, it turns into a distance between point clouds that corresponds to finding the one-to-one matching that minimizes the sum of the distances between pairs of matched points. Since there is no inherent ordering in point cloud data, computing the EMD between two point clouds involves finding a matching based on the euclidean distance between points. The matching is constrained to be bipartite so that one point cloud is completely transformed to the other, without any fractional assignment, and the transport cost is minimal.

The EMD is the most commonly used distance metric on point clouds, and is extremely useful in many different contexts. In particular as we will see for both assessing the performance of, and for training variational autoencoders, since the generated point cloud should get as close as possible to the target in terms of displacement. It can also be interpreted as the distance between two distributions computed with a finite number of samples and reflects the notion of nearness properly, does not have quantization/binning and non-overlapping support problems of most other metrics, e.g., f -divergences, total variation distance, etc.

The EMD between point clouds can be computed exactly, but it is extremely expensive computationally. The standard method is the Hungarian matching Kuhn (1955) algorithm whose complexity is $O(N^3)$ where N is the number of points. Due to this computational cost, training deep generative models for point clouds is not done with this metric, even though it is a metric of choice for performance evaluation. The standard approach uses the *Chamfer distance* (CD) as surrogate. This metric can be computed in $O(N^2)$ time complexity but relaxes the one-to-one matching, which may create pathological situations.

We propose a deep architecture that takes as input two point clouds encoded as sets of geometric coordinate tuples, and computes an accurate estimate of the EMD. We show that the most efficient

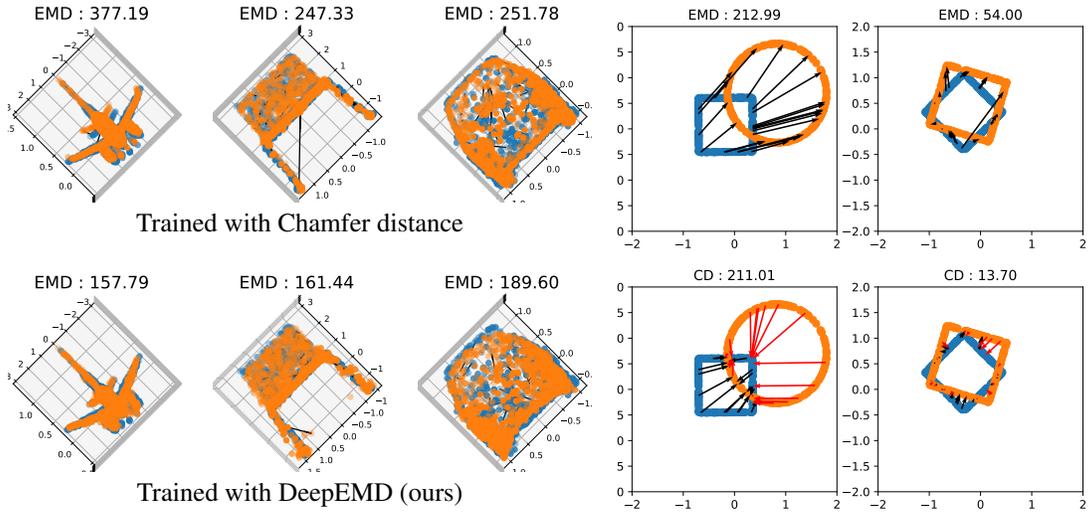


Figure 1: Example point clouds (blue) and their VAE reconstructions (orange) when trained with different reconstruction losses. Training with DeepEMD (bottom) consistently achieves lower reconstruction error (EMD, shown on top of each example) than with the standard Chamfer distance (top).

Figure 2: Example pairs of point clouds. The true earth mover’s distance (EMD, top) and Chamfer distance (CD, bottom) are shown above each example. Arrows indicate the matching between the two point clouds under their respective metrics.

approach, in particular if that estimate is used as a loss for a training process, is to estimate the matching matrix itself. Since the EMD is the sum of the distance between matched points, this approach provides a very accurate estimate of the gradient with respect to the point coordinates. Training a deep variational autoencoder with our model instead of an exact computation is up to $\times 100$ faster (see Fig. 8), and the resulting model performs far better than one trained with the usual Chamfer surrogate (see § 4.4 and Fig. 7).

The key contributions of this paper can be summarized as:

- We propose DeepEMD which approximates the EMD between point clouds in $O(N^2)$ time complexity vs $O(N^3)$ of the hungarian algorithm (§ 3).
- We propose to cast the prediction of a bipartite matching as an attention matrix from which we get an accurate estimate of the EMD and its gradient (§ 3.2).
- We show that DeepEMD generalizes well to unseen data distributions (§ 4.3), and can be used for evaluation of generative models. It provides accurate estimates of the gradients of the distance and demonstrate that DeepEMD can be used as a surrogate reconstruction loss for training deep generative models of point clouds (§ 4.4).
- We show that DeepEMD achieves about $40\times$ speed-up over Sinkhorn algorithm, achieving equal or better performance for various metrics (§ 4.3).

2 RELATED WORK

The two commonly used distance metrics for point clouds in literature are Earth Mover’s Distance (EMD) and Chamfer Distance (CD). Consider two point clouds $X = \{x_i\}_{i=1}^N$ and $Y = \{y_j\}_{j=1}^N$, where $x_i, y_j \in \mathbb{R}^d$. The EMD between the two point clouds can be computed as,

$$\text{EMD}(X, Y) = \min_{\phi \in \mathcal{M}(X, Y)} \sum_{x \in X} \|x - \phi(x)\|_2, \tag{1}$$

where $\mathcal{M}(X, Y)$ is the set of 1-to-1 (bipartite) mappings from X to Y . In addition to the distance, the optimal matching ϕ^* is also interesting for some applications. Since directly optimizing the EMD is computationally expensive, most methods in the literature rely on CD as a proxy similarity measure or reconstruction loss. The CD can be computed as,

$$\text{CD}(X, Y) = \sum_{x \in X} \min_{y \in Y} \|x - y\|_2^2 + \sum_{y \in Y} \min_{x \in X} \|x - y\|_2^2, \quad (2)$$

and in $O(N^2)$ time complexity. The CD solution leads to a non-bipartite one-to-many matching between $x \rightarrow y$ and vice versa. We can also use the L_2 measure with $d = \|x - y\|_2$ instead $d = \|x - y\|_2^2$ to make it comparable to EMD. Note that the above EMD for point clouds is related to the Wasserstein-2 metric (see appendix § A for details). The utility of EMD is limited by the $O(N^3)$ computational cost of evaluating it. There have been several research efforts to circumvent this issue in various application settings.

This is the case for application to point clouds where N is usually in the range of several thousands. Kim et al. (2021) trains a variational auto-encoder with CD as the reconstruction loss. EMD is still the metric of choice for evaluating point cloud generative models Huang et al. (2022); Luo and Hu (2021); Kim et al. (2021); Yang et al. (2019); Shu et al. (2019); Achlioptas et al. (2018). Another issue is disparity between performance measures (minimum matching distance, coverage, etc.) computed with EMD and CD, the comparisons are contradictory and often inconsistent across measures.

CD is usually insensitive to mismatched local density while EMD is dominated by global distribution and overlooks the fidelity of detailed structures Wu et al. (2021). Wu et al. (2021) proposes a new similarity metric called Density-aware Chamfer distance (DCD) to tackle these issues. DCD is derived from CD and can also be computed in $O(N^2)$ time complexity. Urbach et al. (2020) proposed Deep Point Cloud Distance (DPDist) which measures the distance between the points in one cloud and the estimated continuous surface from which the other point cloud is sampled. The surface is estimated locally by a network using the 3D modified Fisher vector representation.

Shirdhonkar and Jacobs (2008) proposed a linear time algorithm for approximating the EMD by exploiting the Hölder continuity constraint in its dual form to convert it into a simple optimization problem with an explicit solution in the wavelet domain and computed as the sum of absolute values of the weighted wavelet co-efficients of the difference histogram. However, their approach is limited to low dimensional histograms. In the optimal transport literature, several efforts have been taken towards improving the statistical and computational properties. Recently, Chuang et al. (2022) proposed Information Maximizing Optimal Transport (InfoOT) which is an information-theoretic extension of optimal transport based on kernel density estimation of the mutual information which introduces global structure into OT maps. The resulting solution maximizes the mutual information between domains while minimizing geometric distance and improves the capability for handling data clusters and outliers.

Other approaches focus on regularizing the OT problem for making it smooth and strictly convex Cuturi (2013); Flamary et al. (2016); Genevay et al. (2018); Blondel et al. (2018). *Sinkhorn distances* Cuturi (2013) smooth the classic OT problem with an entropic regularization term and can be computed through Sinkhorn’s matrix scaling algorithm at a speed that is several orders of magnitude faster than that of transport solvers. We provide more details in the appendix. Meta OT Amos et al. (2022) proposes a meta model to predict the solution to the optimal transport problem which is then used to initialize a standard Sinkhorn solver to further refine the predicted solution. The architecture of the meta model depends on the data domain, and DeepEMD can be utilized when working with point clouds. The choice of meta model architecture is contingent upon the specific data domain, and DeepEMD demonstrates its exceptional utility in point cloud processing.

In this paper, our goal is to approximate the EMD using a deep network in a learning based paradigm where each sample represents two distributions and the target for regression is either the true metric i.e. EMD or the optimal matching ϕ^* , or both. Existing point cloud datasets can serve as an interesting learning problem, where we can interpret each point cloud as a 2D or 3D distribution of points on a shape (manifold). It is posed as a supervised learning problem where the task is to estimate the true EMD, or the true bipartite matching, or both, between a pair of input point clouds.

3 METHOD

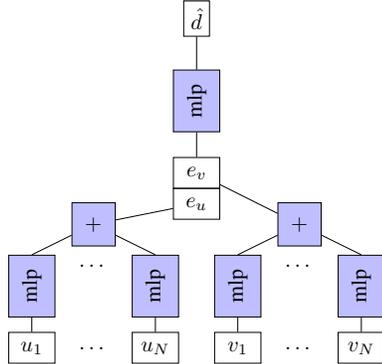


Figure 3: The MLP model (see § 3.1) predicts directly an estimate \hat{d} of the EMD.

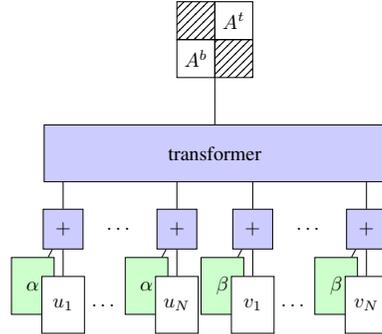


Figure 4: The transformer model we use for DeepEMD (see §3.2) predicts directly the bipartite graph as an attention matrix.

We are interested in building a model which operates on a pair of point clouds (U, V) as input, where $U, V \in \mathcal{D}^N$, $U = \{u_i\}_{i=1}^N$, $V = \{v_j\}_{j=1}^N$, $u_i, v_j \in \mathbb{R}^D$, and N is the cardinality of the point clouds. Note that the points are unordered and the indexing is arbitrary. We denote the earth mover’s distance between them as $d = \text{EMD}(U, V)$ where $d \in \mathbb{R}$. The goal of the model is to predict d and ∇d . Also, let $M \in \{0, 1\}^{N \times N}$ denote the ground truth bipartite matching from EMD, where $M_{i,j} = 1$ indicates that u_i is matched to v_j and vice-versa. Bipartite-ness implies $\forall j, \sum_i M_{i,j} = 1$ and $\forall i, \sum_j M_{i,j} = 1$.

Since point clouds are unordered and invariant to elementwise permutation, we seek mappings $f : \mathcal{D}^N \times \mathcal{D}^N \rightarrow \mathbb{R}$ which are permutation invariant for any permutations π and π' , i.e.,

$$f(U, V) = f(\pi(U), \pi'(V)), \quad (3)$$

In the following sections, we first introduce a simple MLP based baseline, followed by our transformer-based model, *DeepEMD*.

3.1 PREDICTING THE DISTANCE

We propose a simple MLP baseline composed of a point-wise MLP backbone, followed by a prediction head which is also a MLP (see figure 3). The backbone MLP takes a point cloud and returns an embedding $e \in \mathbb{R}^d$ as,

$$e_u = \sum_{i=1}^n g(u_i), \quad e_v = \sum_{j=1}^n g(v_j) \quad (4)$$

The prediction head then produces the final prediction as,

$$\hat{d} = h(e_u \oplus e_v) + h(e_v \oplus e_u), \quad (5)$$

where \oplus denotes vector concatenation. Both g and h are composed of sequential linear layers with ReLU non-linearity between layers. The embeddings are permutation equivariant because of the sum aggregation which does not depend on the ordering of points. Further, we concatenate the embeddings both ways as in Eq. (5), which makes the mapping symmetric. We train the model with mean-squared error loss, $l = (d - \hat{d})^2$. Since the model does not predict the matching, we can interpret it from the direction of the gradient of a point $\delta v_j = \left[\frac{\partial \hat{d}}{\partial V} \right]_j$, e.g., by taking cosine similarity between δv_j and $u_i - v_j$, where u_i is the point matching to v_j from EMD.

3.2 PREDICTING BIPARTITE MATCHING

The transformer Vaswani et al. (2017) intuitively seems to be a very good model for reasoning with point clouds and matching. Moreover, by considering that the output is the last layer’s attention

matrix, we can use it to directly predict the bipartite matching. Since the EMD—and consequently its gradient with respect to the point positions—is a function of the point positions and the matching array, predicting the latter leads to an accurate estimate of the former, that in particular is shielded from the issue of a possible decoupling between matching a functional point-wise (e.g. for MSE) and matching the gradients. While it is viable to directly predict the distance with a transformer-like architecture, we chose to only predict the matching since it is straightforward to estimate the distance using the predicted matching.

We propose *DeepEMD* composed of a sequence of multi-head full attention layers, followed by a prediction head which is also a full attention layer, but with a single head. Given two point clouds U and V , add a learned cloud-specific positional embedding to indicate if a point originates from U or V , we concatenate the points sequence and feed the resulting $I = U \cup V$ as input to the model. The group-id embedding helps the model in modulating attention locally within a point cloud as well as globally across both point clouds. We tried other variants with self-attention layers, cross attention layers, and an alternating mixture of both, but found full attention over both point clouds to be best performing.

For our problem, we get the input X for the transformer by adding positional embeddings to I . Let $\vec{0}^n$ and $\vec{1}^n$ denote a vector of n zeros and ones, respectively. X is then obtained as,

$$P = \vec{0}^n \cup \vec{1}^n, \quad X = I + W^P[P] \quad (\text{indexing}) \quad (6)$$

The intermediate feature $t(X)$ from the transformer encoder (see appendix for details) has the same number of elements as X with each element now being a contextualized representation for the corresponding point in the input. Further, these intermediate representation are fed into a single-head attention layer which outputs the attention matrix as,

$$K = t(X)W^K, \quad Q = t(X)W^Q, \quad A = \frac{QK^\top}{d_k} \quad (7)$$

$$A^t = A_{:,n,n}, \quad A^b = A_{n:,n} \quad (8)$$

Here, A is a $2N \times 2N$ matrix and we slice the top-right block (first N rows and last N columns) as A^t and bottom-left block (last N rows and first N columns) as A^b . $A^t_{i,j}$ can be interpreted as the relatedness of u_i with v_j . Similarly, $A^b_{i,j}$ can be interpreted as the relatedness of v_i with u_j .

Given M , the ground truth bipartite matching from EMD, we define the loss as the average of the cross-entropies as,

$$l(U, V) = \frac{1}{N} \sum_{i=1}^N \text{CE}(A^t_{i,\cdot}, M_{i,\cdot}) + \frac{1}{N} \sum_{i=1}^N \text{CE}(A^b_{\cdot,i}, M_{\cdot,i}) \quad (9)$$

The EMD is then estimated with the predicted matching as,

$$\phi^b(i) = \operatorname{argmax}_j A^b_{i,j}, \quad \phi^t(i) = \operatorname{argmax}_j A^t_{i,j} \quad (10)$$

$$\hat{d} = \frac{1}{2} \left(\sum_i \|u_i - v_{\phi^t(i)}\| + \sum_i \|v_i - u_{\phi^b(i)}\| \right) \quad (11)$$

4 EXPERIMENTS

In this section, we present the overall experimental setup, performance results and comparisons of DeepEMD across various tasks.

4.1 DATASETS

We consider different datasets for our experiments - *Syn2D*, ShapeNet Chang et al. (2015), ModelNet40 Wu et al. (2015) and ScanObjectNN Uy et al. (2019). *Syn2D* consists of 2D point clouds generated synthetically by sampling points on squares and circles (see Fig 2). ShapeNet and ModelNet40 are datasets of 3D point clouds derived from 3D CAD models for different real world objects

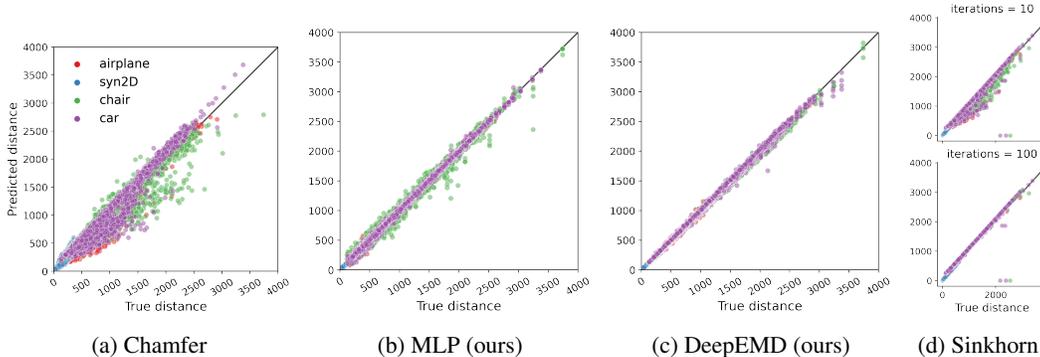


Figure 5: Scatter plot for true vs. approximate EMD from different models/metrics on validation splits for Syn2D and ShapeNet datasets. DeepEMD (ours) consistently performs better across different categories as it has less dispersion. Sinkhorn algorithm becomes more accurate with more iterations. Also note that it encounters numerical errors for some examples.

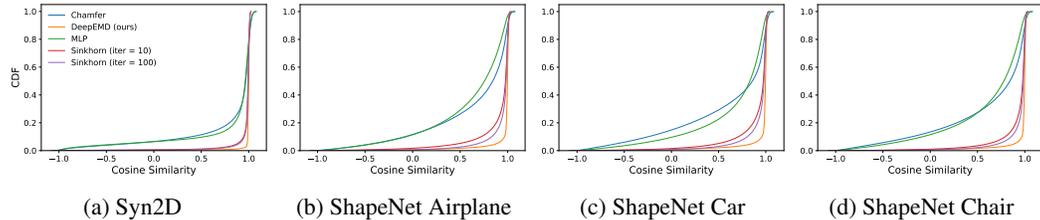


Figure 6: CDF of cosine similarity between true and estimated gradients for all points across all point clouds collected together on validation splits for Syn2D and ShapeNet datasets. The ideal cdf curve should have all the mass at cosine similarity 1. DeepEMD (ours) consistently outperforms all the other methods across different datasets.

like chairs, cars, airplanes, etc. ScanObjectNN is a relatively new real-world point cloud object dataset based on scanned indoor scene data. In order to improve and assess generalization, we augment the train and test splits with synthetic perturbations. We provide more details about the datasets and these augmentations in the Appendix § B.

4.2 PERFORMANCE MEASURES

We consider various measures to assess EMD approximation methods for distance as well as matching estimation. We compare accuracy and computation time to that of Sinkhorn and CD (see § 2).

Distance Estimation. We visualize the true vs. predicted distance through scatter plots (Fig. 5), we expect the data points to be close to $x = y$ line. We compare various correlation measures : linear correlation (r), Spearman correlation (ρ) and Kendall-Tau correlation (τ), to assess the quality of distance estimation. The Spearman and Kendall-Tau are rank-statistic based correlation measures, indicative of the correspondence between two rankings. Note that, correlation measures are useful metrics as they indicate appropriateness of the predicted metric as a distance measure, irrespective of their absolute values. Additionally we look at different quantiles (RE_n) of relative approximation error, which penalizes the difference between absolute values of the predicted and true distance.

Matching Estimation. In order to assess quality of the matching, we consider the cosine similarity between the true and predicted gradient. The true gradient of EMD is always along the matched point. We visualize the cumulative distribution function (cdf) of cosine similarities (Fig. 6), where we expect all the mass to be close to 1. We also look at different quantiles (CS_n) of the cosine similarity. We also consider accuracy which is computed as the average accuracy of matching source points to target points and vice-versa, bipartiteness (B) which is fraction of points with bipartite

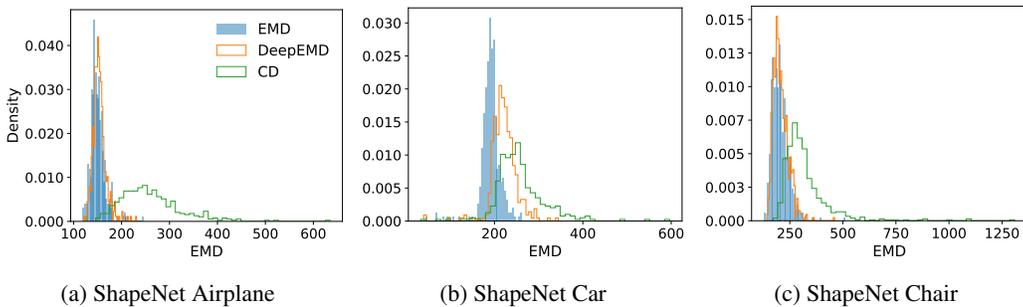


Figure 7: Comparison of EMD between input and reconstructed point clouds from SetVAE trained with different reconstruction losses. The better model should have smaller reconstruction loss and thus mass close to zero in the histograms. DeepEMD (ours) is consistently better as compared to Chamfer loss and very similar to EMD loss.

matching, and also bipartiteness-correctness (B_{corr}) which is fraction of points which are bipartite as well as matched correctly.

4.3 RESULTS

EMD Prediction. Fig. 5 shows the scatter plot of the true EMD vs. approximate EMD predicted from our trained models on the validation split for *Syn2D* and ShapeNet datasets. Note that the validation split also contains the augmentations as discussed in Sec. 4.1. We also validate on specific splits for which the results are shown in the appendix. The plots indicate that both DeepEMD (Fig. 5c) and MLP baseline (Fig. 5b) approximate the EMD faithfully. The MLP baseline seems to struggle a bit on ShapeNet Chair dataset. The higher dispersion in Chamfer (Fig. 5a) and Sinkhorn with 10 iterations (Fig. 5d, top) indicates poor EMD estimation. The approximation with Sinkhorn algorithm becomes more accurate with higher number of iterations (Fig. 5d, bottom), as expected.

We summarize various metrics in Tables 5, 6 and 7 the appendix. DeepEMD and MLP baseline both achieve linear correlation higher than 0.99 in each case. The models achieve Kendall-Tau correlation close to 0.99 and 0.96, respectively, and Spearman correlation close to 0.99 in each case, indicating that ordering of the samples based on approximate distances and true distances are very similar, and monotonicity of samples are preserved. It can be observed that DeepEMD is best except for relative error where it may do worse than our method with MLP. Also it is interesting to note that Sinkhorn performs worse than DeepEMD on correlation- or relative error-based measures.

Matching/Gradient Prediction. Estimating the matching and gradient of the distance is particularly important for training models with DeepEMD as a surrogate distance function. Note that the gradient of a point from true EMD is always along the matched point in the other point cloud. Fig. 6a shows the cdf of cosine similarity between the true and estimated gradient for all the points across all point clouds collected together for *Syn2D*, while Figs. 6b, 6c, and 6d for ShapeNet. The cdf has most mass at cosine similarity close to 1 with a very short tail and is never negative indicating that the estimated gradient is aligned with the true gradient for DeepEMD. This is particularly important when the model is used as a surrogate reconstruction loss. Ideally, the model should provide good estimate of the true gradient throughout training and more particularly in the very beginning when the reconstructions are very noisy, and also towards the end when reconstructions likely become very similar to the training distribution. We discuss more on this in the next sections and the appendix. The MLP baseline usually did not perform well in this regard and also in generalizing to unseen examples and thus was not useful for training generative models. The same issue can be observed with Chamfer distance as well. Table 5 (appendix) shows the cosine similarity quantiles (CS_n), as well as accuracy, bipartiteness (B) and bipartiteness-correctness (B_{corr}). DeepEMD performs better than other models and metrics in each of these measures and indicates better matching approximation.

Out-of-distribution generalization. The generalization of the prediction to a novel distribution is particularly important for a surrogate metric. We test the out-of-distribution behaviour of our models in two different settings : Table 1 shows the generalization performance of the model trained on

a single category of ShapeNet and tested on validation split of multi-category ModelNet40 dataset, while Tables 8, 9 and 10 in the appendix show the performance when tested on different ShapeNet categories. The results indicate that DeepEMD generalizes well when test and train data differ without any adaptation or fine-tuning. Further, the validation performance on a category of a model trained on another category (see Appendix for details) is very similar to the performance of the model trained on the same category. These quite remarkable behaviors point towards the network “meta-learning” in some way the matching algorithm. This is further strengthened by the results on scaling to different number of points during test time as shown in Table 2.

Table 1: Out-of-distribution (dataset) generalization for our models and comparison with other metrics (Chamfer and Sinkhorn), tested on full validation split for ModelNet40 (with 40 categories) and ScanObjectNN (with 15 categories). The models are trained on a single ShapeNet category. The reported numbers are averaged over these categories as well as four training seeds. The first six rows show distance estimation metrics (see § 4.2), while the last six rows correspond to matching estimation metrics. The arrows next to the metrics indicate whether higher (\uparrow) values are better or lower (\downarrow). Chamfer and Sinkhorn are deterministic, thus variances are not reported. Further, MLP does not provide accuracy and bipartiteness metrics.

DATASET	MODELNET40				SCANOBJECTNN			
	CHAMFER	SINKHORN	MLP	DEEPEMD	CHAMFER	SINKHORN	MLP	DEEPEMD
r (\uparrow)	0.951	0.971	0.959 ± 0.011	0.999 ± 0.0	0.971	0.929	0.965 ± 0.005	0.997 ± 0.001
ρ (\uparrow)	0.935	0.988	0.945 ± 0.017	0.999 ± 0.0	0.979	0.965	0.963 ± 0.007	0.999 ± 0.0
τ (\uparrow)	0.792	0.983	0.819 ± 0.024	0.974 ± 0.002	0.882	0.968	0.855 ± 0.011	0.973 ± 0.002
$RE_{0.1}$ (\downarrow)	0.03	0.057	0.009 ± 0.001	0.005 ± 0.002	0.025	0.038	0.013 ± 0.001	0.004 ± 0.001
$RE_{0.5}$ (\downarrow)	0.129	0.102	0.062 ± 0.005	0.019 ± 0.004	0.094	0.078	0.076 ± 0.005	0.019 ± 0.004
$RE_{0.9}$ (\downarrow)	0.321	0.2	0.257 ± 0.03	0.04 ± 0.004	0.282	0.244	0.299 ± 0.025	0.051 ± 0.005
$CS_{0.1}$ (\uparrow)	-0.067	0.824	-0.293 ± 0.047	0.927 ± 0.003	0.138	0.879	-0.208 ± 0.042	0.946 ± 0.002
$CS_{0.5}$ (\uparrow)	0.834	0.986	0.684 ± 0.023	1.0 ± 0.0	0.917	0.992	0.719 ± 0.02	0.999 ± 0.0
$CS_{0.9}$ (\uparrow)	0.997	0.999	0.96 ± 0.003	1.0 ± 0.0	0.998	1.0	0.965 ± 0.003	1.0 ± 0.0
ACCURACY (\uparrow)	12.651	31.91	-	56.38 ± 0.604	7.673	20.04	-	40.671 ± 0.62
B (\uparrow)	17.045	33.458	-	70.401 ± 0.672	9.474	19.43	-	56.269 ± 0.71
B_{corr} (\uparrow)	6.544	19.615	-	47.084 ± 0.741	3.38	9.961	-	30.055 ± 0.678

Scaling number of points. Remarkably, the size of point clouds during testing can differ greatly from those during training without degrading performance. Table 2 shows performance of the model for test point cloud sizes ranging from 256 to 8196, while training was done with only 1024 points. Prediction of the metric itself (top 6 rows) does not degrade for all practical purposes. Regarding the matching estimation, directional measure of performance related to the cosine similarity (rows CS_n) do not degrade neither. We can notice degradation in accuracy based measures (last 3 rows) which is natural since the problem becomes difficult with increasing number of points N because of its combinatorial nature. For training when memory requirement is much higher due to backprop, we can use smaller number of points, and scale it up during inference without any fine-tuning.

Table 2: Scaling number of points and out-of-distribution (scale) generalization for DeepEMD. The models are trained on a single ShapeNet category with 1024 points and tested on validation split of same category but with different number of points. Reported values are averaged over 4 training seeds. DeepEMD generalizes well to unseen number of points at test time without fine-tuning.

# points	← Less # points than training →			Trained	← More # points than training →		
	256	512	768		1024	2048	4096
r	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	0.999 ± 0.0	0.999 ± 0.001
ρ	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	0.999 ± 0.0	0.998 ± 0.0
τ	0.985 ± 0.0	0.987 ± 0.0	0.988 ± 0.0	0.988 ± 0.001	0.986 ± 0.001	0.981 ± 0.002	0.974 ± 0.004
$RE_{0.1}$	0.002 ± 0.001	0.002 ± 0.001	0.004 ± 0.002	0.007 ± 0.003	0.012 ± 0.005	0.013 ± 0.007	0.014 ± 0.008
$RE_{0.5}$	0.01 ± 0.002	0.011 ± 0.002	0.014 ± 0.003	0.017 ± 0.005	0.027 ± 0.009	0.034 ± 0.013	0.04 ± 0.016
$RE_{0.9}$	0.026 ± 0.003	0.026 ± 0.003	0.029 ± 0.004	0.032 ± 0.005	0.042 ± 0.009	0.054 ± 0.013	0.066 ± 0.018
$CS_{0.1}$	0.94 ± 0.002	0.955 ± 0.002	0.961 ± 0.001	0.964 ± 0.001	0.967 ± 0.001	0.967 ± 0.001	-
$CS_{0.5}$	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	-
$CS_{0.9}$	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	-
Accuracy	72.348 ± 0.44	69.384 ± 0.383	66.901 ± 0.379	64.648 ± 0.404	57.588 ± 0.464	47.78 ± 0.51	35.274 ± 0.483
B	81.857 ± 0.755	80.101 ± 0.547	78.013 ± 0.507	75.896 ± 0.521	68.658 ± 0.584	58.109 ± 0.597	44.603 ± 0.734
B_{corr}	64.838 ± 0.756	61.558 ± 0.587	58.545 ± 0.547	55.719 ± 0.568	46.831 ± 0.618	35.053 ± 0.6	21.606 ± 0.469

Computational Time and Complexity. Fig. 8 compares the evaluation time for different models and metrics. DeepEMD achieves a significant speedup of about $100\times$ as compared to EMD and $40\times$

as compared to Sinkhorn with 100 iterations. This speedup becomes more pronounced on bigger point clouds as hungarian algorithm takes $O(N^3)$ time vs. $O(N^2)$ for DeepEMD.

4.4 DEEPEMD USED AS A LOSS

Training a SetVAE, as for any auto-encoder, requires a reconstruction loss to assess the quality of the learned representation. While the eventual goal would be to minimize the EMD, standard approach uses Chamfer Distance due to the prohibitive computation cost of calculating the EMD. Instead of Chamfer Distance we propose to use DeepEMD and demonstrate its utility as a reconstruction loss as compared to Chamfer Distance.

DeepEMD was trained separately on each category of ShapeNet dataset and the trained model was then used as a surrogate reconstruction loss for training a variational auto-encoder. We use SetVAE Kim et al. (2021), a transformer based VAE adapted for point clouds and set-structured data. The parameters of DeepEMD module are frozen during training of the SetVAE. We follow exactly the same protocol as in SetVAE and train using ShapeNet categories of airplane, chair, and car and also the same hyper-parameters for training. Fig. 1 and Fig. 9 (appendix) shows the reconstruction on validation data achieved by SetVAE models trained with different reconstruction losses. DeepEMD consistently achieves lower reconstruction EMD as compared to CD. This is further verified from Fig. 7 which shows the distribution of true EMD between a point cloud and its reconstruction.

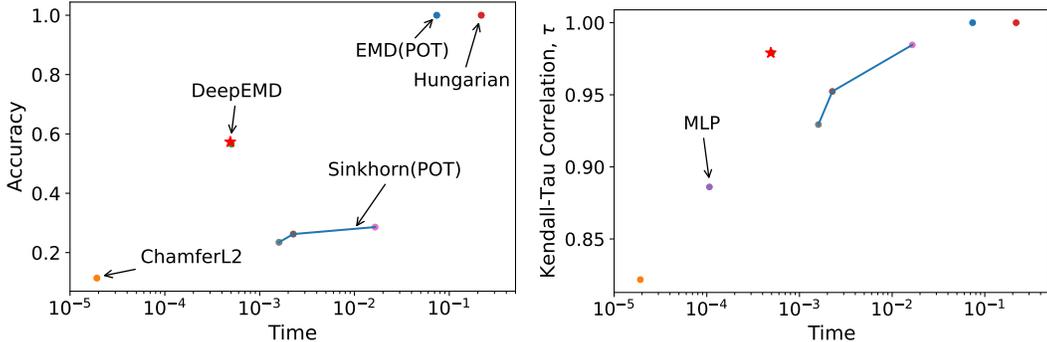


Figure 8: Comparison of empirical evaluation time and different performance measures : Accuracy (left) and Kendall-Tau correlation τ (right). We use Python Optimal Transport (POT) library for computing Sinkhorn distances, and show metrics at different iterations (5, 10 and 100). DeepEMD is $\sim 100\times$ and $\sim 40\times$ faster than Hungarian algorithm and Sinkhorn (100 iterations), respectively.

5 CONCLUSION AND FUTURE WORK

We propose DeepEMD, a method for fast approximation of EMD, improving time complexity from $O(N^3)$ to $O(N^2)$. It is composed of a multi-head multi-layer transformer, followed by a single-head full attention layer as the final output layer. It operates on two point clouds and outputs an attention matrix which is interpreted as the matching matrix and optimized to match the ground truth matching obtained from the hungarian algorithm. We demonstrated the effectiveness of DeepEMD in approximating the true EMD for synthetic 2D point clouds as well as real world datasets like ShapeNet, ModelNet40 and ScanObjectNN. It achieves a speed-up of $\times 100$ with 1024 points. Further, we show that it estimates the gradients well, generalizes well for unseen point clouds (or distributions), and can be used for end-to-end training of point cloud autoencoders achieving faster convergence than Chamfer distance surrogate.

It would be interesting to explore fast transformer variants to further improve from the quadratic time complexity for future work. In terms of architecture various pooling/un-pooling strategies can be explored which can help with both, better time complexity and improved feature learning. In this work, we estimate the Wasserstein -2 metric, and extension to other Wasserstein $-p$ metrics and other optimal transport problems could also be interesting for various applications. Lastly, extension to general probability distributions with fractional assignments (i.e. mass splitting) can also be very useful and valuable for some applications.

REFERENCES

- Panos Achlioptas, Olga Diamanti, Ioannis Mitliagkas, and Leonidas Guibas. Learning representations and generative models for 3d point clouds. In *International conference on machine learning*, pages 40–49. PMLR, 2018.
- Brandon Amos, Samuel Cohen, Giulia Luise, and Ievgen Redko. Meta optimal transport. *arXiv preprint arXiv:2206.05262*, 2022.
- Mathieu Blondel, Vivien Seguy, and Antoine Rolet. Smooth and sparse optimal transport. In *International conference on artificial intelligence and statistics*, pages 880–889. PMLR, 2018.
- Angel X Chang, Thomas Funkhouser, Leonidas Guibas, Pat Hanrahan, Qixing Huang, Zimo Li, Silvio Savarese, Manolis Savva, Shuran Song, Hao Su, et al. Shapenet: An information-rich 3d model repository. *arXiv preprint arXiv:1512.03012*, 2015.
- Ching-Yao Chuang, Stefanie Jegelka, and David Alvarez-Melis. Infoot: Information maximizing optimal transport. *arXiv preprint arXiv:2210.03164*, 2022.
- Marco Cuturi. Sinkhorn distances: Lightspeed computation of optimal transport. *Advances in neural information processing systems*, 26, 2013.
- R Flamary, N Courty, D Tuia, and A Rakotomamonjy. Optimal transport for domain adaptation. *IEEE Trans. Pattern Anal. Mach. Intell.*, 1, 2016.
- 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. *The Journal of Machine Learning Research*, 22(1):3571–3578, 2021.
- Aude Genevay, Gabriel Peyré, and Marco Cuturi. Learning generative models with sinkhorn divergences. In *International Conference on Artificial Intelligence and Statistics*, pages 1608–1617. PMLR, 2018.
- Tianxin Huang, Xuemeng Yang, Jiangning Zhang, Jinhao Cui, Hao Zou, Jun Chen, Xiangrui Zhao, and Yong Liu. Learning to train a point cloud reconstruction network without matching. In *Computer Vision—ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part I*, pages 179–194. Springer, 2022.
- Jinwoo Kim, Jaehoon Yoo, Juho Lee, and Seunghoon Hong. Setvae: Learning hierarchical composition for generative modeling of set-structured data. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 15059–15068, 2021.
- Harold W Kuhn. The hungarian method for the assignment problem. *Naval research logistics quarterly*, 2(1-2):83–97, 1955.
- Shitong Luo and Wei Hu. Diffusion probabilistic models for 3d point cloud generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2837–2845, 2021.
- Sameer Shirdhonkar and David W Jacobs. Approximate earth mover’s distance in linear time. In *2008 IEEE Conference on Computer Vision and Pattern Recognition*, pages 1–8. IEEE, 2008.
- Dong Wook Shu, Sung Woo Park, and Junseok Kwon. 3d point cloud generative adversarial network based on tree structured graph convolutions. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 3859–3868, 2019.
- Dahlia Urbach, Yizhak Ben-Shabat, and Michael Lindenbaum. Dpdist: Comparing point clouds using deep point cloud distance. In *European Conference on Computer Vision*, pages 545–560. Springer, 2020.
- Mikaela Angelina Uy, Quang-Hieu Pham, Binh-Son Hua, Thanh Nguyen, and Sai-Kit Yeung. Revisiting point cloud classification: A new benchmark dataset and classification model on real-world data. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 1588–1597, 2019.

Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017.

Tong Wu, Liang Pan, Junzhe Zhang, Tai Wang, Ziwei Liu, and Dahua Lin. Density-aware chamfer distance as a comprehensive metric for point cloud completion. *arXiv preprint arXiv:2111.12702*, 2021.

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*, pages 1912–1920, 2015.

Guandao Yang, Xun Huang, Zekun Hao, Ming-Yu Liu, Serge Belongie, and Bharath Hariharan. Pointflow: 3d point cloud generation with continuous normalizing flows. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 4541–4550, 2019.

A OPTIMAL TRANSPORT AND WASSERSTEIN DISTANCES

The Wasserstein- p metric between two probability distributions μ_X and ν_Y is defined as,

$$W_p(\mu, \nu) = \left(\inf_{\gamma \in \Gamma(\mu, \nu)} \mathbf{E}_{(x, y) \sim \gamma} \|x - y\|_p \right)^{1/p}, \quad (12)$$

where $\Gamma(\mu, \nu)$ are all possible joint distributions where $X, Y \in \mathcal{D}$, (\mathcal{D}, d) defines a metric space (here, $d = \|x - y\|_p$) and marginals satisfy $\int_{\mathcal{D}} \gamma(x, y) dy = \mu(x)$ and $\int_{\mathcal{D}} \gamma(x, y) dx = \nu(y)$. Note that, while the distance is useful in itself, the optimal transport plan γ^* is also interesting for some applications.

Given samples from μ and ν , $W_p(\mu, \nu)$ can be computed by solving the optimal transport problem,

$$\gamma^* = \operatorname{argmin}_{\gamma \in \mathbb{R}_+^{m \times n}} \sum_{i, j} \gamma_{i, j} M_{i, j} \quad (13)$$

$$\text{s.t. } \gamma \mathbf{1} = [\hat{\mu}]_m; \gamma^T \mathbf{1} = [\hat{\nu}]_n; \gamma \geq 0 \quad (14)$$

where $[\hat{\mu}]_m$ and $[\hat{\nu}]_n$ represent binned histograms derived from samples from μ and ν with m and n bins respectively. M is a $m \times n$ distance or cost matrix, $M_{i, j}$ represents the cost $d(x, y)$ to transport mass from bin $[\hat{\mu}]_m^i$ to bin $[\hat{\nu}]_n^j$.

It is also possible to avoid binning and compute the Wasserstein distance directly from samples. The problem can then be formulated as,

$$\gamma^* = \operatorname{argmin}_{\gamma \in \mathbb{R}_+^{n \times n}} \sum_{i, j} \gamma_{i, j} M_{i, j} \quad (15)$$

$$\text{s.t. } \gamma \mathbf{1} = 1; \gamma^T \mathbf{1} = 1; \gamma \in \{0, 1\} \quad (16)$$

where $M_{i, j}$ now represents the cost of transporting point x_i to y_j . Each point is considered to be sampled i.i.d. from their respective distributions. Unlike the previous case, we can no longer transport a quanta of a point or mass i.e. fractional assignment is not meaningful, naturally leading to a bipartiteness constraint. This is exactly same as Eq (1) and the optimal transport plan for the problem is a linear sum assignment problem and can be computed using the hungarian algorithm.

B DATASETS

We train and evaluate on *Syn2D*, ShapeNet and ModelNet40 datasets.

Syn2D. We generate 2D synthetic point clouds by uniformly sampling 200 points on simple 2D shapes of circles and squares. The circles are generated by uniformly sampling the center and radius from $(0, 1]$. For the squares, we sample its center, rotation and scale uniformly from $[-0.5, 0.5]$, $[0, \frac{\pi}{2}]$, and $[0.5, 1]$, respectively. We refer to the synthetic dataset as *Syn2D* (Fig. 2).

ShapeNet. ShapeNet Chang et al. (2015) is a richly-annotated, large-scale dataset of 3D shapes. We train and evaluate our models using point clouds from one of the three categories in the ShapeNet dataset : airplane, chair, and car. We sample 1024 points from the ShapeNet point clouds for training and evaluate on a range of point cloud sizes.

ModelNet40. Chang et al. (2015) introduced the ModelNet project to provide a comprehensive and clean collection of 3D CAD models for objects, compiled using a list of the most common object categories. In our experiments, we evaluate our models trained with a single ShapeNet category on the ModelNet40 dataset. ModelNet40 has 40 different categories.

ScanObjectNN. Uy et al. (2019) proposed the ScanObjectNN dataset, a real-world point cloud object dataset based on scanned indoor scene data, in order to provide a more realistic benchmark as compared to ModelNet40. In our experiments, we evaluate our models trained with a single ShapeNet category on pairs of point clouds derived from the training set of ScanObjectNN dataset.

Table 3 provides a summary and statistics about of the datasets.

Table 3: Summary of datasets. The top five rows show statistics about the original point cloud datasets, while the bottom four rows show statistics about the pair point cloud dataset that was used for our evaluation and/or training.

	SYN2D	MODELNET40	SCANOBJECTNN	SHAPENET	AIRPLANE	CAR	CHAIR
# CATEGORIES	2	40	15	55	1	1	1
FEATURE DIM	2	3	3	3	3	3	3
CARDINALITY	200	2048	2048	15000	15000	15000	15000
# TRAIN SAMPLES	8000	9840	2309	35708	2832	2458	4612
# VAL SAMPLES	2000	2468	581	5158	405	352	662
# TRAIN PAIR SAMPLES	20000	-	-	-	10000	10000	10000
TRAIN CARDINALITY	200	-	-	-	1024	1024	1024
# VAL PAIR SAMPLES	5000	2000	10000	-	2000	2000	2000
VAL CARDINALITY	200	1024	1024	-	256-8192	256-8192	256-8192

Further, we build pairs (U, V) of point clouds by randomly sampling pairs from the train or validation splits of point clouds datasets summarized in Table . We refer the first argument in the pair U as the *source* and the second argument V as *target*.

B.1 AUGMENTATIONS

In order to improve generalization, we augment the datasets with point cloud pairs. We randomly sample a point cloud pair (U, V) from the dataset and another noisy point cloud N by randomly sampling points from $\mathcal{N}(0, 1)$ and scaling the whole point cloud by $\sigma \sim \mathcal{U}(0.1, 1.1)$. Further, we augment the point cloud pairs according to the following schemes:

- (U, V) : the originally sampled pair.
- (U, N) : target is replaced by the noisy point cloud.
- $(U, U + N)$: target is a corrupted version of U with additive noise N .
- $(U, \tilde{U} + N)$: \tilde{U} denotes a point cloud which is similar to U . For *Syn2D*, we perturb the surface parameters (radius, scale, center, etc.) used for sampling U and sample points on the perturbed surface. For ShapeNet, we independently sample different set of points from the original surface.
- $(U, V + N)$: target is a corrupted version of a randomly sampled point cloud from the dataset.

The resulting dataset constitutes 20% samples from each of the above splits. Validation splits are generated randomly and independently for *Syn2D*, while for ShapeNet and ModelNet40, we use the validation split provided. We also augment the validation split with the same scheme as discussed above. The specifics of the pair point-cloud dataset is summarized in Table 3.

C MODELS

Transformers. Let $X = [x_1, \dots, x_m] \in \mathbb{R}^{m \times d_{\text{model}}}$ be the input sequence (or set) of m vectors. A transformer layer performs the following computation Vaswani et al. (2017) :

$$X' = \text{LN}(X + \text{Multihead}(X))$$

$$t_l(X) = \text{LN}(X' + \text{FFN}(X'))$$

where LN and FFN stand for layer norm and feed-forward network, respectively. Multihead denotes a multi-head attention layer which allows the model to jointly attend to information from different representation subspaces at different positions. It consists of a stack of H scaled dot-product attention layers and computes key, query and value matrices, followed by a softmax as follows :

$$K_h = XW_h^K, \quad Q_h = XW_h^Q, \quad V_h = XW_h^V$$

$$A_h = \text{soft-max} \left(\frac{Q_h K_h^\top}{d_k} \right) V_h$$

$$\text{MultiHead}(X) = \bar{A} = \text{concat}(A_1, \dots, A_H) W_O$$

The final output of the encoder can be written as a composition:

$$t(X) = t_N(t_{N-1}(\dots(t_1(X))))$$

MLP. The baseline MLP model predicts directly the distance and it does not use the matching information. It has about 110K parameters in total. The point-wise MLP backbone $g(\cdot)$ is composed of three hidden layer of sizes 4, 8 and 16, with ReLU non-linearity. It outputs a single embedding of dimension 128 for each point cloud after aggregating the point level features. The embedding of the point clouds are then concatenated and passed to the prediction head which is also an MLP with four hidden layers of sizes 256, 128, 64 and 16, and outputs a single scalar which is interpreted as the predicted distance.

DeepEMD. We use a transformer encoder backbone which transforms the raw input point clouds into contextualized point level features. The transformer encoder is followed by the output layer which computes the queries and keys and finally the attention matrix which is interpreted as the matching as explained in § 3.2. The model constitutes about 803K learnable parameters, with 8 transformer encoder layers, each with 6 heads. The latent dimensions (d_{model} , d_{keys} , etc.) for each layer were all set to 78.

We use the ADAM optimizer with a constant learning rate of 0.001 for DeepEMD and 0.0001 for the MLP.

D SINKHORN DISTANCE

The Sinkhorn distance Cuturi (2013) considers a regularized OT optimization problem :

$$\begin{aligned} \gamma^* &= \operatorname{argmin}_{\gamma \in \mathbb{R}_+^{m \times n}} \sum_{i,j} \gamma_{i,j} M_{i,j} + \lambda \Omega(\gamma) \\ \text{s.t. } \gamma \mathbf{1} &= [\hat{\mu}]_m; \gamma^T \mathbf{1} = [\hat{\nu}]_n; \gamma \geq 0 \end{aligned}$$

where λ is the regularization coefficient and $O(\gamma) = \sum_{i,j} \gamma_{i,j} \log(\gamma_{i,j})$ is an entropy regularization term which makes the optimization problem smooth and strictly convex allowing for optimization procedures such as the Sinkhorn-Knopp algorithm. In this paper, we use the Python Optimal Transport (POT) python library Flamary et al. (2021) for computing the Sinkhorn distances. It is an iterative algorithm and can be evaluated in $O(N^2)$ time complexity.

In our experiments, we used a regularization multiplier of 0.1. Reducing the multiplier typically resulted in numerical issues for a larger portion of the samples, whereas increasing it led to greater inaccuracies in estimation with the same computation time (for a fixed number of iterations). We show the performance metrics with different regularisation parameters in Table 4.

Table 4: Effect of regularization parameter on Sinkhorn’s algorithm (100 iterations). Results on ScanObjectNN dataset with 10000 samples. The second row shows number of instances which had numerical issues.

λ	0.08	0.1	0.12	0.14
# FAIL	686	80	6	0
r	0.616	0.929	0.993	0.999
ρ	0.708	0.965	0.997	0.999
τ	0.785	0.968	0.988	0.986
RE _{0.1}	0.033	0.038	0.045	0.052
RE _{0.5}	0.068	0.078	0.093	0.110
RE _{0.9}	0.288	0.244	0.166	0.342
CS _{0.1}	0.815	0.879	0.869	0.851
CS _{0.5}	0.991	0.992	0.991	0.989
CS _{0.9}	0.999	1.0	0.999	0.999
ACCURACY	21.29	20.049	18.49	17.22
B	20.72	19.439	18.09	17.03
B _{corr}	10.97	9.961	8.98	8.23

The difficulty of choosing the regularisation parameter also makes it difficult for training generative models with Sinkhorn as it often leads to numerical issues. We experimented with several variants

of Sinkhorn, including log-space Sinkhorn, log-stabilized Sinkhorn, Greenkhorn, etc. However, we discovered that they either exhibited significantly slower performance compared to the vanilla variant, frequently failed to converge to a satisfactory optimal transport matrix within a finite timeframe, or sometimes exhibited both issues.

E EXTENDED RESULTS

E.1 DISTANCE AND MATCHING ESTIMATION

We compare performance of different models and metrics in Table 5. The models were trained on a single ShapeNet category and evaluated on the validation split of the same category. The numbers are averaged over all training categories as well as four training seeds.

Table 5: Performance comparison of different metrics and models. The models are trained on a single ShapeNet category and evaluated on the test split of the same category. The reported numbers are averaged over all categories and four training seeds. The first six rows show distance estimation metrics (see § 4.2), while the last six rows correspond to matching estimation metrics. The arrows next to the metrics indicate whether higher (\uparrow) or lower (\downarrow) values are better. Chamfer and Sinkhorn are deterministic, thus variances are not reported. Further, our MLP model does not provide accuracy and bipartiteness metrics.

MODEL	CHAMFER	SINKHORN	MLP (OURS)	DEEPEMD (OURS)
r (\uparrow)	0.963	0.995	0.998 ± 0.0	1.0 ± 0.0
ρ (\uparrow)	0.953	0.997	0.998 ± 0.001	1.0 ± 0.0
τ (\uparrow)	0.827	0.987	0.966 ± 0.003	0.988 ± 0.001
$RE_{0.1}$ (\downarrow)	0.023	0.051	0.002 ± 0.0	0.007 ± 0.003
$RE_{0.5}$ (\downarrow)	0.109	0.106	0.015 ± 0.001	0.017 ± 0.005
$RE_{0.9}$ (\downarrow)	0.31	0.271	0.076 ± 0.006	0.032 ± 0.005
$CS_{0.1}$ (\uparrow)	-0.173	0.831	-0.034 ± 0.049	0.964 ± 0.001
$CS_{0.5}$ (\uparrow)	0.85	0.986	0.798 ± 0.018	1.0 ± 0.0
$CS_{0.9}$ (\uparrow)	0.998	0.999	0.974 ± 0.003	1.0 ± 0.0
ACCURACY (\uparrow)	11.677	28.407	-	64.648 ± 0.404
B (\uparrow)	17.784	31.889	-	75.896 ± 0.521
B_{corr} (\uparrow)	5.626	16.658	-	55.719 ± 0.568

Tables 6 and 7 show the per-category performance comparison, averaged over four training seeds.

Table 6: Per-category distance estimation performance measures of different models and metrics when train and test category are same. The reported number are averaged over four training seeds.

TRAIN_CATE	MODEL/METRIC	r	ρ	τ	$RE_{0.1}$	$RE_{0.5}$	$RE_{0.9}$
AIRPLANE	CHAMFER	0.9797	0.9647	0.8519	0.0141	0.0962	0.3018
	DEEPEMD	0.9998 ± 0.0	0.9997 ± 0.0	0.9879 ± 0.001	0.0039 ± 0.0014	0.0142 ± 0.0031	0.0306 ± 0.004
	MLP	0.9992 ± 0.0001	0.9986 ± 0.0001	0.9722 ± 0.0018	0.002 ± 0.0003	0.0125 ± 0.0018	0.0772 ± 0.0038
	SINKHORN	0.9998	0.9997	0.9881	0.0496	0.1104	0.2984
CAR	CHAMFER	0.9675	0.9564	0.8318	0.035	0.1167	0.324
	DEEPEMD	0.9997 ± 0.0001	0.9998 ± 0.0	0.9891 ± 0.0006	0.006 ± 0.0047	0.0156 ± 0.0068	0.0302 ± 0.0097
	MLP	0.9993 ± 0.0002	0.9988 ± 0.0004	0.9746 ± 0.0035	0.0018 ± 0.0002	0.0112 ± 0.001	0.0625 ± 0.012
	SINKHORN	0.991	0.9941	0.9854	0.0531	0.1124	0.2986
CHAIR	CHAMFER	0.9431	0.9382	0.7976	0.0192	0.1133	0.3037
	DEEPEMD	0.9997 ± 0.0	0.9997 ± 0.0001	0.9866 ± 0.0013	0.0103 ± 0.0088	0.0225 ± 0.0118	0.0356 ± 0.012
	MLP	0.9965 ± 0.0013	0.9959 ± 0.0015	0.9504 ± 0.0079	0.0037 ± 0.0006	0.0214 ± 0.0023	0.0881 ± 0.0124
	SINKHORN	0.9942	0.9969	0.9886	0.049	0.0955	0.2167

E.2 OUT-OF-DISTRIBUTION GENERALIZATION

Table 8 shows the out-of-distribution generalization performance for our models. The trained model on a particular ShapeNet category is evaluated on the validation split of other ShapeNet categories.

Table 7: Per-category matching estimation performance measures of different models and metrics when train and test category are same. The reported number are averaged over four training seeds.

TRAIN_CATE	MODEL/METRIC	CS _{0.1}	CS _{0.5}	CS _{0.9}	ACCURACY	B	B_CORR
AIRPLANE	CHAMFER	-0.0813	0.8446	0.9973	10.1768	16.7437	4.7461
	DEEPEMD	0.9643 \pm 0.0027	1.0 \pm 0.0	1.0 \pm 0.0	61.9119 \pm 0.9043	73.3128 \pm 1.1746	52.2082 \pm 1.246
	MLP	-0.0492 \pm 0.0533	0.7766 \pm 0.0263	0.9722 \pm 0.003	-	-	-
	SINKHORN	0.8314	0.9871	0.9994	25.2956	29.018	13.9732
CAR	CHAMFER	-0.246	0.8615	1.0	13.7079	20.7186	6.8617
	DEEPEMD	0.9585 \pm 0.0025	1.0 \pm 0.0	1.0 \pm 0.0	67.7243 \pm 0.6554	78.7286 \pm 0.797	59.6723 \pm 0.9245
	MLP	0.017 \pm 0.0652	0.8388 \pm 0.0187	0.9804 \pm 0.0029	-	-	-
	SINKHORN	0.8043	0.9845	0.9993	31.1621	35.5971	19.3105
CHAIR	CHAMFER	-0.1929	0.8442	0.9976	11.145	15.8883	5.2694
	DEEPEMD	0.9703 \pm 0.0007	1.0 \pm 0.0	1.0 \pm 0.0	64.3079 \pm 0.4712	75.6459 \pm 0.657	55.2757 \pm 0.7011
	MLP	-0.0695 \pm 0.1211	0.7793 \pm 0.0438	0.97 \pm 0.0072	-	-	-
	SINKHORN	0.8558	0.9876	0.9992	28.7644	31.0521	16.69

The numbers are averaged over these other categories as well as four training seeds. Tables 9 and 10 show the performance in the same setting but for each test category separately.

Table 8: Out-of-distribution (category) generalization for our models and comparison with other metrics (Chamfer and Sinkhorn). The models are trained on a single ShapeNet category and evaluated on other ShapeNet categories. The reported numbers are averaged over these categories as well as four training seeds. The first five rows show distance estimation metrics (see § 4.2), while the last five rows correspond to matching estimation metrics. The arrows next to the metrics indicate whether higher (\uparrow) or lower (\downarrow) values are better.

MODEL	MLP (OOD)	MLP	DEEPEMD (OOD)	DEEPEMD
r (\uparrow)	0.98 \pm 0.019	0.998 \pm 0.001	0.999 \pm 0.001	1.0 \pm 0.0
ρ (\uparrow)	0.976 \pm 0.02	0.998 \pm 0.001	0.999 \pm 0.0	1.0 \pm 0.0
τ (\uparrow)	0.886 \pm 0.04	0.966 \pm 0.004	0.977 \pm 0.003	0.988 \pm 0.001
RE _{0.1} (\downarrow)	0.014 \pm 0.003	0.002 \pm 0.0	0.009 \pm 0.009	0.007 \pm 0.005
RE _{0.5} (\downarrow)	0.065 \pm 0.018	0.015 \pm 0.002	0.024 \pm 0.012	0.017 \pm 0.007
RE _{0.9} (\downarrow)	0.319 \pm 0.132	0.076 \pm 0.009	0.05 \pm 0.011	0.032 \pm 0.008
CS _{0.1} (\uparrow)	-0.208 \pm 0.089	-0.034 \pm 0.074	0.933 \pm 0.004	0.964 \pm 0.002
CS _{0.5} (\uparrow)	0.714 \pm 0.047	0.798 \pm 0.027	1.0 \pm 0.0	1.0 \pm 0.0
CS _{0.9} (\uparrow)	0.963 \pm 0.006	0.974 \pm 0.004	1.0 \pm 0.0	1.0 \pm 0.0
ACCURACY (\uparrow)	-	-	54.35 \pm 1.16	64.648 \pm 0.606
B (\uparrow)	-	-	67.922 \pm 1.343	75.896 \pm 0.782
B _{corr} (\uparrow)	-	-	44.293 \pm 1.445	55.719 \pm 0.851

Table 9: Per-category distance estimation performance measures of different models and metrics when train and test category are different. The reported number are averaged over four training seeds.

TRAIN_CATE	TEST_CATE	r	ρ	τ	RE _{0.1}	RE _{0.5}	RE _{0.9}
AIRPLANE	AIRPLANE	0.9998 \pm 0.0	0.9997 \pm 0.0	0.9879 \pm 0.001	0.0039 \pm 0.0014	0.0142 \pm 0.0031	0.0306 \pm 0.004
	CAR	0.9989 \pm 0.0006	0.9997 \pm 0.0001	0.9864 \pm 0.001	0.0035 \pm 0.0005	0.016 \pm 0.0036	0.0308 \pm 0.0039
	CHAIR	0.9981 \pm 0.0003	0.9983 \pm 0.0003	0.9689 \pm 0.0016	0.0044 \pm 0.0011	0.0212 \pm 0.0033	0.0491 \pm 0.0048
CAR	AIRPLANE	0.9993 \pm 0.0002	0.9989 \pm 0.0003	0.9766 \pm 0.0017	0.0092 \pm 0.0065	0.0256 \pm 0.011	0.0634 \pm 0.0114
	CAR	0.9997 \pm 0.0001	0.9998 \pm 0.0	0.9891 \pm 0.0006	0.006 \pm 0.0047	0.0156 \pm 0.0068	0.0302 \pm 0.0097
	CHAIR	0.9978 \pm 0.0001	0.9981 \pm 0.0002	0.9671 \pm 0.0017	0.0045 \pm 0.0019	0.0206 \pm 0.006	0.0525 \pm 0.0022
CHAIR	AIRPLANE	0.999 \pm 0.0004	0.9983 \pm 0.0008	0.9751 \pm 0.0047	0.0164 \pm 0.0114	0.034 \pm 0.0127	0.0631 \pm 0.0092
	CAR	0.9987 \pm 0.0007	0.9996 \pm 0.0001	0.9867 \pm 0.0007	0.0141 \pm 0.0116	0.0275 \pm 0.014	0.0422 \pm 0.0139
	CHAIR	0.9997 \pm 0.0	0.9997 \pm 0.0001	0.9866 \pm 0.0013	0.0103 \pm 0.0088	0.0225 \pm 0.0118	0.0356 \pm 0.012

E.3 DEEPEMD AS A LOSS

Fig. 9 shows more samples with the input point cloud and the reconstructed output from SetVAE when trained with EMD, Chamfer or DeepEMD as the reconstruction loss.

Table 10: Per-category matching estimation performance measures of different models and metrics when train and test category are different. The reported number are averaged over four training seeds.

TRAIN_CATE	TEST_CATE	CS _{0.1}	CS _{0.5}	CS _{0.9}	ACCURACY	B	B_CORR
AIRPLANE	AIRPLANE	0.9643 ± 0.0027	1.0 ± 0.0	1.0 ± 0.0	61.9119 ± 0.9043	73.3128 ± 1.1746	52.2082 ± 1.246
	CAR	0.9347 ± 0.004	1.0 ± 0.0	1.0 ± 0.0	61.4086 ± 0.9383	72.6142 ± 1.3357	51.6545 ± 1.3383
	CHAIR	0.926 ± 0.0045	0.9994 ± 0.0001	1.0 ± 0.0	48.9237 ± 0.9275	63.1441 ± 1.0583	38.1275 ± 1.1302
CAR	AIRPLANE	0.9279 ± 0.0027	0.9997 ± 0.0	1.0 ± 0.0	50.7472 ± 1.1176	66.2 ± 1.0652	40.6049 ± 1.3128
	CAR	0.9585 ± 0.0025	1.0 ± 0.0	1.0 ± 0.0	67.7243 ± 0.6554	78.7286 ± 0.797	59.6723 ± 0.9245
	CHAIR	0.9218 ± 0.0035	0.9994 ± 0.0001	1.0 ± 0.0	47.7712 ± 1.1373	63.8324 ± 1.1691	37.6786 ± 1.2958
CHAIR	AIRPLANE	0.9401 ± 0.0031	0.9998 ± 0.0	1.0 ± 0.0	53.7969 ± 0.8777	67.1421 ± 1.0913	43.3759 ± 1.0926
	CAR	0.945 ± 0.0006	1.0 ± 0.0	1.0 ± 0.0	63.4508 ± 0.5761	74.5965 ± 0.7858	54.3188 ± 0.8313
	CHAIR	0.9703 ± 0.0007	1.0 ± 0.0	1.0 ± 0.0	64.3079 ± 0.4712	75.6459 ± 0.657	55.2757 ± 0.7011

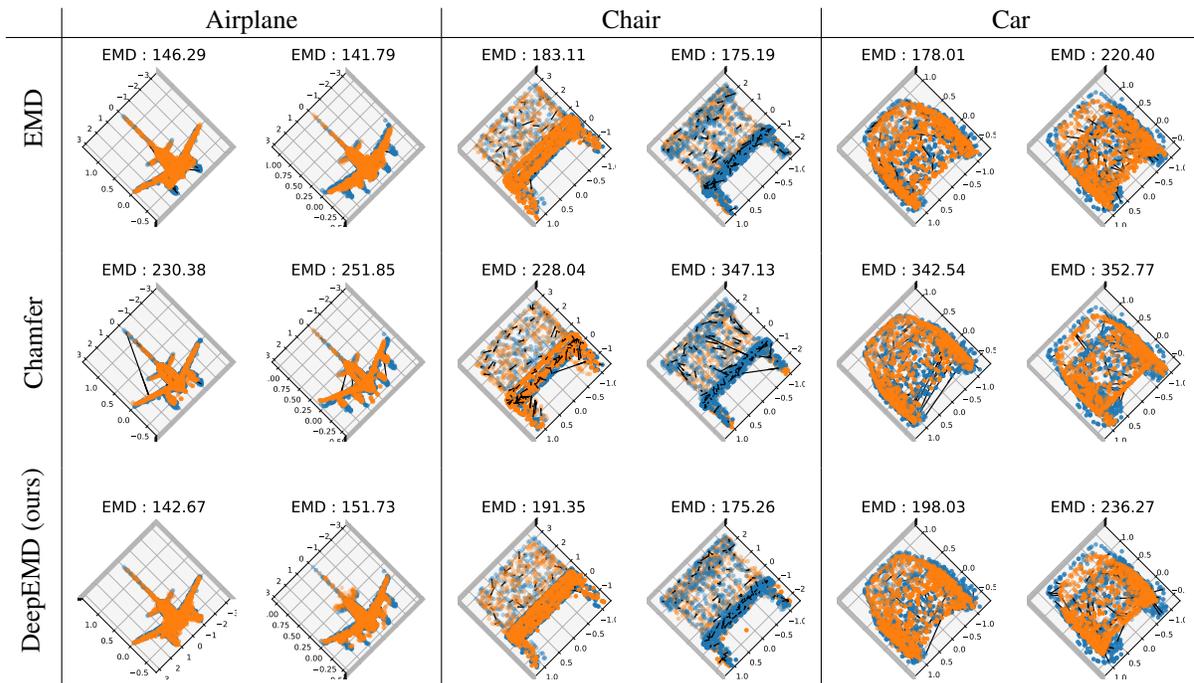


Figure 9: Reconstruction on validation data with SetVAE trained with different reconstruction losses : EMD (top), Chamfer (middle) and DeepEMD surrogate (bottom). Training with DeepEMD as a loss consistently achieves lower reconstruction EMD as compared to Chamfer loss.