

Learning with Batch-wise Optimal Transport Loss for 3D Shape Recognition

Lin Xu^{1,2*}Han Sun^{1,2}Yuai Liu^{1,2}¹Nanjing Institute of Advanced Artificial Intelligence ²Horizon Robotics

{lin01.xu, han.sun, yuai.liu}@horizon.ai

Abstract

Deep metric learning is essential for visual recognition. The widely used pair-wise (or triplet) based loss objectives cannot make full use of semantical information in training samples or give enough attention to those hard samples during optimization. Thus, they often suffer from a slow convergence rate and inferior performance. In this paper, we show how to learn an importance-driven distance metric via optimal transport programming from batches of samples. It can automatically emphasize hard examples and lead to significant improvements in convergence. We propose a new batch-wise optimal transport loss and combine it in an end-to-end deep metric learning manner. We use it to learn the distance metric and deep feature representation jointly for recognition. Empirical results on visual retrieval and classification tasks with six benchmark datasets, i.e., MNIST, CIFAR10, SHREC13, SHREC14, ModelNet10, and ModelNet40, demonstrate the superiority of the proposed method. It can accelerate the convergence rate significantly while achieving a state-of-the-art recognition performance. For example, in 3D shape recognition experiments, we show that our method can achieve better recognition performance within only 5 epochs than what can be obtained by mainstream 3D shape recognition approaches after 200 epochs.

1. Introduction

Learning a semantical embedding metric to make similar positive samples cluster together, while dissimilar negative ones widen apart is an essential part for modern recognition tasks [24, 12]. With the flourish of deep learning technologies [31, 47, 53], deep metric learning has gained more attention in recent years [26, 5, 44, 15, 50]. By training deep neural networks discriminatively end-to-end, a more complex and highly-nonlinear deep feature representation (from the input space to a lower dimensional semantical embedding metric space) can be learned. The jointly learned deep feature representation and embedding metric yield significant improvement for recognition applications, such as 2D image retrieval [59, 5, 37] or classification [60, 42], sig-

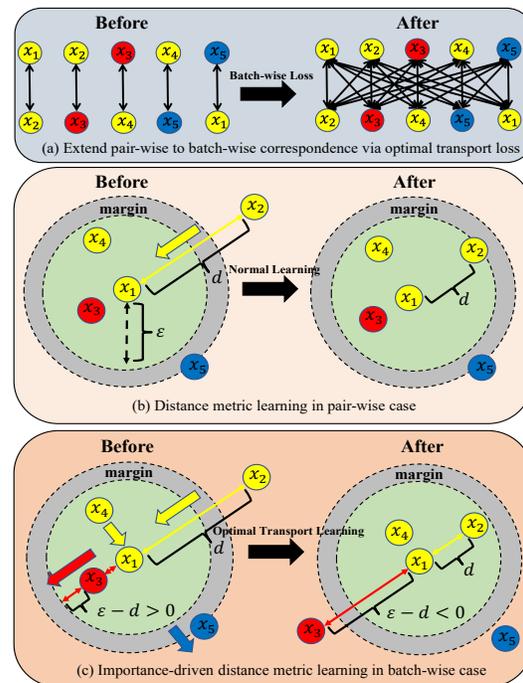


Figure 1. Schematic illustration of learning with the proposed batch-wise loss objective as compared to the pair-wise loss objective. Colors of circles represent semantical (or category) information. **(a)**: The relationships among batches of samples of these two loss objectives. **(b)**: Only the semantical information of a pair of examples is considered at each update in the pair-wise case. **(c)**: The importance-driven distance metric is optimized using all available information within training batches so that similar positive examples with large ground distances and dissimilar negative examples with small ground distances are emphasized automatically. Arrows indicate the weights (or importance) on distances arising from the proposed batch-wise optimal transport loss.

nature verification [6], face recognition [12, 60, 44], and sketch-based 3D shape cross-modality retrieval [33, 58, 63].

Despite the progress made, most of the pre-existing loss objectives [6, 12, 26, 44, 5] do have some limitations for metric learning. Commonly used contrastive loss [24, 12] or triplet loss [60, 10] only considers the semantical information within individual pairs or triplets of examples at each

update, while the interactions with the rest ones within the batch are ignored. It would bias the learned embedding metric and feature representation. Moreover, they do not give enough attention to hard positive or negative examples, by cause of the fact that these samples are often sparsely distributed and expensive to seek out. These *hard samples* can strongly influence parameters during the network is learned to correct them. As a consequence, methods which neglect them often suffer from slow convergence rate and suboptimal performance. Occasionally, such methods require expensive sampling techniques to accelerate the training process and boost the learning performance [10, 44, 36, 15].

In this paper, we propose a novel batch-wise optimal transport loss objective for deep metric learning. It can learn an importance-driven distance metric via optimal transport programming from batches of samples simultaneously. As we know, the fundamental idea behind metric learning is minimizing the intra-category variations (or distances) while maximizing the inter-category variations (or distances). Thus, those semantically similar positive samples with large ground distances and dissimilar negative examples with small ground distances should be regarded as *hard samples*. Such samples should be emphasized correctly to accelerate the metric learning process. Figure 1 illustrates our main idea of proposing the new batch-wise optimal transport loss objective. As illustrated, learning with the proposed loss can utilize all available semantical information of training batches simultaneously. The introduced importance-driven distance metric is partly obtained as a solution to the optimal transport program [56, 16]. It can mine and emphasize those *hard samples* automatically. Thus, the convergence rate of distance metric learning process can be significantly improved. We further develop the new loss objective in a deep metric learning manner. The whole network can be trained discriminatively in an end-to-end fashion. The jointly learned semantical embedding metric and deep feature representation would be more robust to intra-class and inter-class variations. We finally verify the performance of our proposed method applying to various visual recognition tasks, including 2D image recognition, sketch-based 3D shape cross-modality retrieval, and 3D shape recognition. Experiment results on six widely used benchmark datasets, i.e., *MNIST*, *CIFAR10*, *SHREC13*, *SHREC14*, *ModelNet10* and *ModelNet40*, demonstrate the superiority of the proposed method. Our method can achieve state-of-the-art recognition performances with a notably fast convergence rate.

In a nutshell, our main contributions in the present work can be summarized as follows:

(1) We propose a novel batch-wise optimal transport loss objective for learning an importance-driven distance metric to improve the existing pair-wise based loss objectives.

(2) We develop a deep metric learning method based on

the proposed loss objective, which learns the importance-driven metric and deep feature representation jointly.

(3) We verify the superiority of our proposed method on visual recognition tasks, including 2D image recognition, sketch-based 3D shape retrieval, and 3D shape recognition.

2. Related Work

Recognition of 3D shapes is becoming prevalent with the advancement of modeling, digitizing, and visualizing techniques for 3D objects. The increasing availability of 3D CAD models, both on the Internet, e.g., *Google 3D Warehouse* [1] and *Turbosquid* [2], and in the domain-specific field, e.g., *ModelNet* [3] and *SHREC* [34], has led to the development of several scalable and efficient methods to study and analyze them, as well as to facilitate practical applications. For 3D shape recognition, one fundamental issue is how to construct a determinative yet robust 3D shape descriptor and feature representation. Compared to 2D images, 3D shapes have more complex geometric structures. Their appearance can be affected significantly by innumerable variations such as viewpoint, scale, and deformation. These have brought great challenges into the recognition task. A natural method is to construct a shape descriptor based on the native 3D structures, e.g., point clouds, polygon meshes, and volumetric grid. Then, shapes can be represented with distances, angles, triangle areas, tetrahedra volumes, local shape diameters [38, 9], heat kernel signatures [7, 29], extensions of handcrafted SIFT, SURF [28], and learned 3D CNNs [62, 35] on 3D volumetric grids. An alternative way is describing a 3D shape by a collection of 2D view-based projections. It can make use of CNN models, which have been pre-trained on large 2D image datasets such as ImageNet [31] and gained a decent ability of generalization [20, 47, 23]. In this context, DeepPano [46] and PANORAMA-NN [45] are developed to convert 3D shapes into panoramic views, e.g., a cylinder projection around its principal axis. Multi-view CNN (*MVCNN*) [52] groups multiple CNNs with a view-pooling structure to process and learn from all available 2D projections of a 3D shape jointly.

3. Background

Loss objective for metric learning: Metric learning aims to learn a semantical metric from input samples. Let $\mathbf{x} \in \mathbf{X}$ be an input sample. The kernel function $f(\cdot; \theta) : \mathbf{X} \rightarrow \mathbb{R}^d$ takes input \mathbf{x} and generates an feature representation or embedding $f(\mathbf{x})$. In deep metric learning [24, 12, 50], kernel $f(\cdot; \theta)$ is usually defined by a deep neural network, parameterized by a series of weights and bias θ . Metric learning optimizes a discriminative loss objective to minimize intra-class distances while maximizing inter-class distances. For example, the contrastive loss in seminal Siamese network [24, 12] takes pairs of samples as input and trains two iden-

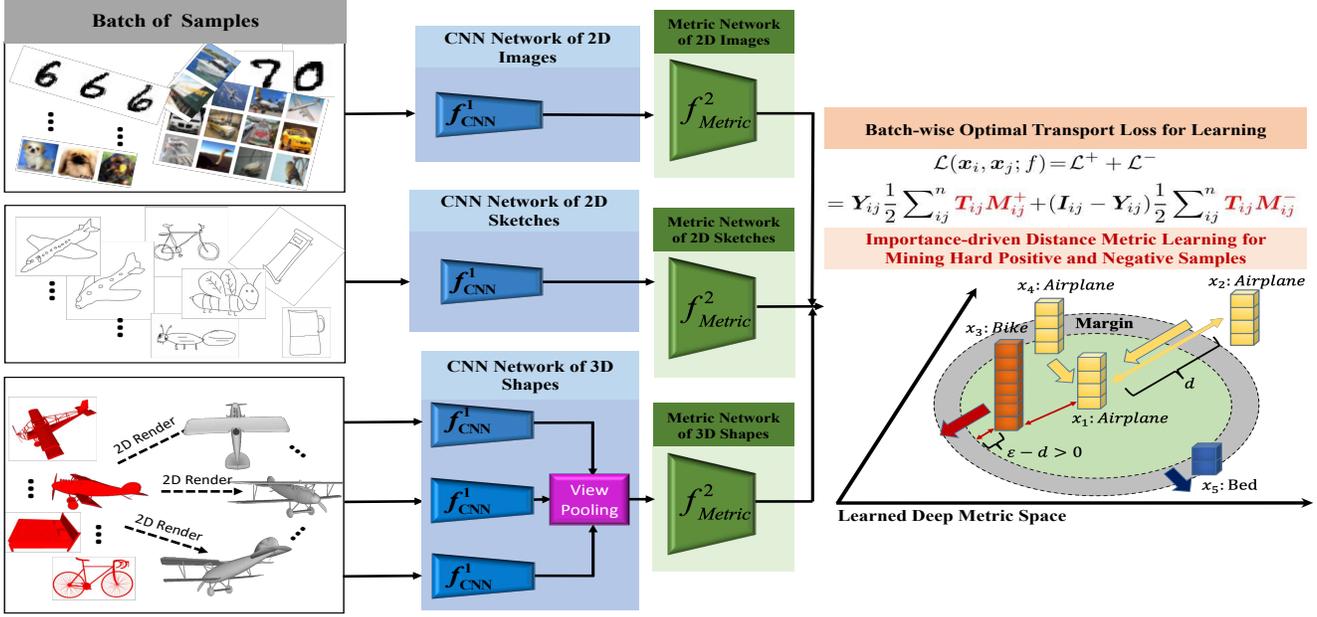


Figure 2. We formulate the proposed loss into a deep metric learning framework. Given batches of each modality samples, we use *LeNet-5* [32], *ResNet-50* [25], and *MVCNN* [52] as f_{CNN}^1 to extract deep CNN features for 2D images, 2D sketches, and 3D shapes, respectively. The metric network f_{Metric}^2 consisting of four fully connected (FC) layers, i.e., 4096-2048-512-128 (two FC layers 512-256 for *LeNet-5*) is used to perform dimensionality reduction of the CNN features. We add three sigmoid functions as activation among these FC layers to generate normalized and dense feature vectors. The whole framework can be end-to-end trained discriminatively with the new batch-wise optimal transport loss. The highlighted importance-driven distance metrics $T_{ij}M_{ij}^+$ and $T_{ij}M_{ij}^-$ are used for emphasizing hard positive and negative samples. It jointly learns the semantic embedding metric and deep feature representation for retrieval and classification.

tical networks to learn a deep distance metric M_{ij} as

$$\mathcal{L}(\mathbf{x}_i, \mathbf{x}_j; f) = \mathbf{y}_{ij} M_{ij} + (1 - \mathbf{y}_{ij}) \max\{0, \varepsilon - M_{ij}\}, \quad (1)$$

where the label $\mathbf{y}_{ij} \in \{0, 1\}$ indicates whether a pair of $(\mathbf{x}_i, \mathbf{x}_j)$ is from the same class or not. The margin parameter ε imposes a threshold of the distances among dissimilar samples. The metric M_{ij} can be an arbitrary distance measurement in the feature embedding space. Conventionally, the Euclidian metric $M_{ij} = \|(f(\mathbf{x}_i) - f(\mathbf{x}_j))\|_2^2$ is used to denote the distance between a pair of samples. Triplet loss [60, 10] shares a similar idea with contrastive loss, but extends a pair of samples to a triplet. For a given query \mathbf{x}_i , a similar sample \mathbf{x}_j to the query, and a dissimilar one \mathbf{x}_k , the triplet loss can be formulated as

$$\mathcal{L}(\mathbf{x}_i, \mathbf{x}_j, \mathbf{x}_k; f) = \max\{0, M_{ij} - M_{ik} + \varepsilon\}. \quad (2)$$

Intuitively, it encourages the distance between the dissimilar pair $M_{ik} = \|(f(\mathbf{x}_i) - f(\mathbf{x}_k))\|_2^2$ to be larger than the distance between the similar pair $M_{ij} = \|(f(\mathbf{x}_i) - f(\mathbf{x}_j))\|_2^2$ by at least a margin ε .

Optimal transport distances: Optimal transport distances [16], also known as Wasserstein distances [56] or Earth Mover’s distances [43], define a distance between two probability distributions according to principles from optimal transport theory [57, 61]. Formally, let \mathbf{r} and \mathbf{c} be n -dimensional probability measures. The set of transportation

plans between probability distributions \mathbf{r} and \mathbf{c} is defined as $U(\mathbf{r}, \mathbf{c}) := \{\mathbf{T} \in \mathbb{R}_+^{n \times n} | \mathbf{T}\mathbf{1} = \mathbf{r}, \mathbf{T}^T\mathbf{1} = \mathbf{c}\}$, where $\mathbf{1}$ is an all-ones vector. The set of transportation plans $U(\mathbf{r}, \mathbf{c})$ contains all nonnegative $n \times n$ elements with row and column sums \mathbf{r} and \mathbf{c} , respectively.

Give an $n \times n$ ground distance matrix M , the cost of mapping \mathbf{r} to \mathbf{c} using a transport matrix \mathbf{T} can be quantified as $\langle \mathbf{T}, M \rangle$, where $\langle \cdot, \cdot \rangle$ stands for the Frobenius dot-product. Then the problem defined in Equation (3)

$$D_M(\mathbf{r}, \mathbf{c}) := \min_{\mathbf{T} \in U(\mathbf{r}, \mathbf{c})} \langle \mathbf{T}, M \rangle, \quad (3)$$

is called an optimal transport problem between \mathbf{r} and \mathbf{c} given ground cost M . The optimal transport distance $D_M(\mathbf{r}, \mathbf{c})$ measures the cheapest way to transport the mass in probability measure \mathbf{r} to match that in \mathbf{c} .

Optimal transport distances define a more powerful cross-bin metric to measure probabilities compared with some commonly used bin-to-bin metrics, e.g., Euclidean, Hellinger, and Kullback-Leibler divergences. However, the cost of computing D_M is at least $\mathcal{O}(n^3 \log(n))$ when comparing two n -dimensional probability distributions in a general metric space [39]. To alleviate it, Cuturi [16] formulated a regularized transport problem by adding an entropy regularizer to Equation (3). This makes the objective function strictly convex and allows it to be solved efficiently. Particularly, given a transport matrix \mathbf{T} , let

$h(\mathbf{T}) = -\sum_{ij} \mathbf{T}_{ij} \log \mathbf{T}_{ij}$ be the entropy of \mathbf{T} . For any $\lambda > 0$, the regularized transport problem can be defined as

$$D_M^\lambda(\mathbf{r}, \mathbf{c}) := \min_{\mathbf{T} \in U(\mathbf{r}, \mathbf{c})} \langle \mathbf{T}, \mathbf{M} \rangle - \frac{1}{\lambda} h(\mathbf{T}), \quad (4)$$

where the larger λ is, the closer this relaxation $D_M^\lambda(\mathbf{r}, \mathbf{c})$ is to original $D_M(\mathbf{r}, \mathbf{c})$. Cuturi [16] also proposed the Sinkhorn’s algorithm [49] to solve Equation (4) for the optimal transport \mathbf{T}^* . Specifically, let the matrix $\mathbf{K} = \exp(-\lambda \mathbf{M})$ and solve it for the scaling vectors \mathbf{u} and \mathbf{v} to a fixed-point by computing $\mathbf{u} = \mathbf{r} / \mathbf{K} \mathbf{v}$, $\mathbf{v} = \mathbf{c} / \mathbf{K}^T \mathbf{u}$ in an alternating way. These yield the optimal transportation plan $\mathbf{T}^* = \text{diag}(\mathbf{u}) \mathbf{K} \text{diag}(\mathbf{v})$. This algorithm can be solved with complexity $\mathcal{O}(n^2)$ [16], which is significantly faster than exactly solving the original optimal transport problem.

4. Our Method

In this section, we propose a deep metric learning scheme by using principles of the optimal transport theory [57]. Currently, research works with optimal transport distance [16, 18, 17] mainly focus on theoretical analysis and simulation verification. Thus, it is hard to apply them into a large-scale 3D shape recognition contest directly. To this end, we have done the following three works to construct a trainable batch-wise optimal transport loss objective.

4.1. Importance-driven Distance Metric Learning

Assuming we are given two batches of samples, each batch has n examples $\mathbf{X} \in \mathbb{R}^{d \times n}$. Let $\mathbf{x}_i \in \mathbb{R}^d$ be the representation of the i^{th} shape. Additionally, let \mathbf{r} and \mathbf{c} be the n -dimensional probability vectors for two batches, where r_i and c_i denote the number of times shape i occurs in \mathbf{r} and \mathbf{c} (normalized overall samples in \mathbf{r} and \mathbf{c}). The optimal transport introduces a transportation plan $\mathbf{T} \in \mathbb{R}^{n \times n}$ such that \mathbf{T}_{ij} describes how much of r_i should be transported to c_j . As described in Equation (4), the optimal transport distance between batches \mathbf{r} and \mathbf{c} can be re-formulated as

$$D_{\text{OT}}^\lambda(\mathbf{r}, \mathbf{c}) = \min_{\mathbf{T} \geq 0} \sum_{i,j=1}^n \mathbf{T}_{ij} \mathbf{M}_{ij} - \frac{1}{\lambda} h(\mathbf{T}_{ij}) \quad (5)$$

s.t. $\sum_{j=1}^n \mathbf{T}_{ij} = \mathbf{r}$ and $\sum_{i=1}^n \mathbf{T}_{ij} = \mathbf{c} \quad \forall i, j.$

The learned optimal transportation plan \mathbf{T}^* is a probability distribution [16], which aims to find the least amount of cost needed to transport the mass from batch \mathbf{r} to batch \mathbf{c} . The unit of cost corresponds to transporting a sample by the unit of ground distance. Thus, \mathbf{T}^* solved by Equation (5) prefers to assign higher importance values to samples with small ground distances while leaving fewer for others.

Utilizing such property, we define the importance-driven distance metric via imposing semantical information of samples. Specifically, we first define the ground distances

for a pair of similar positive samples as

$$\mathbf{G}_{ij}^+(\mathbf{x}_i, \mathbf{x}_j; f) = e^{-\gamma \|f(\mathbf{x}_i) - f(\mathbf{x}_j)\|_2^2}, \quad (6)$$

where γ is a hype-parameter controlling the extent of re-scaling. This re-scaling operator shrinks large Euclidian distance between similar samples. After re-scaling \mathbf{G}^+ , the learned \mathbf{T}^* solved by Equation (5) tends to put higher importance values on those similar samples which have far Euclidian distances among each other (a.k.a., *hard positive samples*), while putting lower on the others accordingly. Thus, it would accelerate the process that similar samples are getting close to each other. For dissimilar negative samples, we define the ground distances correspondingly as

$$\mathbf{G}_{ij}^-(\mathbf{x}_i, \mathbf{x}_j; f) = e^{-\gamma \max\{0, \varepsilon - \|f(\mathbf{x}_i) - f(\mathbf{x}_j)\|_2^2\}}. \quad (7)$$

The hinge loss $\max\{0, \varepsilon - \|f(\mathbf{x}_i) - f(\mathbf{x}_j)\|_2^2\}$ penalizes the dissimilar samples within the margin ε and ignores the others. Thus, contrary to the above similar samples case, here the learned \mathbf{T}^* in Equation (5) will pay higher importance values on those dissimilar samples with small Euclidian distances (a.k.a., *hard negative samples*), while assigning fewer on the others. Thus, it could accelerate the process that dissimilar samples are getting apart to each other.

4.2. Batch-wise Optimal Transport Loss Learning

Based on the Euclidian distances of positive and negative samples \mathbf{M}_{ij}^+ , \mathbf{M}_{ij}^- , and optimal transportation plan \mathbf{T}_{ij}^* , now we can formulate a batch-wise optimal transport loss for metric learning. It can be viewed as an n -pairs extension version of the contrastive loss or triplet loss. We define the loss objective as

$$\begin{aligned} \mathcal{L}(\mathbf{x}_i, \mathbf{x}_j; f) &= \mathcal{L}^+ + \mathcal{L}^- \\ &= \mathbf{Y}_{ij} \frac{1}{2} \sum_{ij} \mathbf{T}_{ij} \mathbf{M}_{ij}^+ + (\mathbf{I}_{ij} - \mathbf{Y}_{ij}) \frac{1}{2} \sum_{ij} \mathbf{T}_{ij} \mathbf{M}_{ij}^-, \end{aligned} \quad (8)$$

where \mathbf{Y}_{ij} is a binary label assigned to a pair of training batches. Let $\mathbf{Y}_{ij} = 1$ if sample \mathbf{x}_i and \mathbf{x}_j are deemed similar, and $\mathbf{Y}_{ij} = 0$ otherwise. An all-ones matrix is denoted as $\mathbf{I} \in \mathbb{R}^{n \times n}$ and n is the size of each training batch. In practice, $\mathbf{T}_{ij} \mathbf{M}_{ij}^+$ and $\mathbf{T}_{ij} \mathbf{M}_{ij}^-$ can be regarded as the importance-driven distance metric for positive and negative samples, respectively. Substituting our pre-defined ground distances \mathbf{G}_{ij} for Euclidian distances \mathbf{M}_{ij} in Equation (5), the optimal transportation plan \mathbf{T}_{ij}^* can be obtained. It is a probability distribution of weights for emphasizing hard positive and negative samples during the loss objective optimization. We just write the loss regarding only one pair of batches here for simplicity. The overall data loss objective based on all training batches can be easily derived as $\sum \mathcal{L}$.

4.3. Batch Gradient Descent Optimization

We further derive the back-propagation form of the batch-wise optimal transport loss objective. The proposed

loss objective can be embedded into a deep metric learning framework, so that the whole network can be trained discriminatively end-to-end via batch gradient descent.

Since the batch-wise optimal transport distance (i.e., $\mathbf{T}_{ij}^* \mathbf{M}_{ij}^+$ and $\mathbf{T}_{ij}^* \mathbf{M}_{ij}^-$) is a fully connected dense matrix of pairs-wise ground distances, its gradient can be deduced as network flow manner. Specifically, we compute the gradient of corresponding loss $\mathcal{L}(\mathbf{x}_i, \mathbf{x}_j; f)$ with respect to embedding representations $f(\mathbf{x}_i)$ and $f(\mathbf{x}_j)$ as follows,

$$\begin{aligned} \frac{\partial \mathcal{L}}{\partial f(\mathbf{x}_i)} &= \sum_{j=1}^n \mathbf{T}_{ij}^* (f(\mathbf{x}_i) - f(\mathbf{x}_j)) (\mathbf{Y}_{ij} - (\mathbf{I}_{ij} - \mathbf{Y}_{ij}) \delta_{ij}) \\ \frac{\partial \mathcal{L}}{\partial f(\mathbf{x}_j)} &= - \sum_{i=1}^n \mathbf{T}_{ij}^* (f(\mathbf{x}_i) - f(\mathbf{x}_j)) (\mathbf{Y}_{ij} - (\mathbf{I}_{ij} - \mathbf{Y}_{ij}) \delta_{ij}), \end{aligned} \quad (9)$$

where \mathbf{T}^* is the optimizer obtained from Equation (4). Motivated by the fast optimal distance computation [16, 18, 21], we relax the linear program in Equation (5) using the regularized entropy as in Equation (4). It allows us to approximately solve Equation (4) in $\mathcal{O}(n^2)$ time via $\mathbf{T}^* = \text{diag}(\mathbf{u}) \mathbf{K} \text{diag}(\mathbf{v})$, where n is the size of batch.

The δ here is also a binary indicator assigned to the pairs. Let $\delta_{ij} = 1$ when the Euclidean distance between shape \mathbf{x}_i and \mathbf{x}_j is within the margin (i.e., $\varepsilon - \|f(\mathbf{x}_i) - f(\mathbf{x}_j)\|_2 > 0$), and $\delta_{ij} = 0$ otherwise. The $f(\mathbf{x}_i)$ and $f(\mathbf{x}_j)$ are feature representations obtained through deep neural networks. Therefore, the gradient with respect to the network can be computed easily with the chain-rule in a back-propagation fashion, as far as $\frac{\partial \mathcal{L}}{\partial f(\mathbf{x}_i)}$ and $\frac{\partial \mathcal{L}}{\partial f(\mathbf{x}_j)}$ are derived. We also note that the defined ground distances \mathbf{G}_{ij}^+ and \mathbf{G}_{ij}^- are just used to determine the optimal transportation plan \mathbf{T}^* for re-weighting the importance of similar positive and dissimilar negative samples. We do not consider them as variables to compute gradient in Equation (9) for gradient updating.

5. Experiments

In this section, we evaluated the performance of the proposed method with applications to 2D image recognition (i.e., retrieval and classification), sketch-based 3D shape retrieval, and 3D shape recognition tasks. Six widely used benchmark datasets were employed in our experiments, including *MNIST* [32], *CIFAR10* [30], *SHREC13* [33], *SHREC14* [41], *ModelNet10*, and *ModelNet40* [62].

5.1. Experimental settings

Architecture: Figure 2 illustrates network architecture of deep metric learning with our batch-wise loss objective.

Datasets: The *MNIST* [32] is a large handwritten digits dataset, which has 60,000 28×28 black-and-white training images and 10,000 testing images. The *CIFAR10* [30] dataset consists of 60,000 32×32 RGB images in 10 different categories, with 6,000 images per category. There are 50,000 training images and 10,000 test images. *SHREC13*

[33] and *SHREC14* [41] are two large-scale datasets for sketch-based 3D shape retrieval. *SHREC13* contains 7,200 human-drawn sketches and 1,258 3D shapes from 90 different categories. For each category, 50 sketches are used for training and remaining 30 sketches are used for the test. There are 14 3D shapes per category generally. *SHREC14* is larger than *SHREC13*, which has 13,680 sketches and 8,987 3D shapes from 171 categories. Each of the categories has 53 3D shapes on average. There are 8,550 sketches for training and 5,130 for test. *ModelNet* [3] is a large-scale 3D shape dataset, which contains 151,128 3D CAD models belonging to 660 unique object categories [62]. There are two subsets of *ModelNet* can be used for evaluation. *ModelNet10* contains 4,899 3D shapes from 10 categories while *ModelNet40* has 12,311 shapes from 40 categories. In our experiments, we used the same training and test splits as in [62]. Specifically, we randomly selected 100 unique shapes per category, where 80 shapes were chosen for training and the remaining 20 shapes for the test.

Evaluations: For retrieval, we used Euclidean distance to measure the similarity of the shapes based on their learned feature vectors output by the metric network as shown in Figure 2. Given a query from the test set, a ranked list of the remaining test samples was returned according to their distances to the query sample. We used the evaluation metrics for retrieval as in [63] when presenting our results. The metrics include nearest neighbor (NN) [14], first tier (FT) [54], second tier (ST) [13], E-measure (E) [11], discounted cumulated gain (DCG) [27], and mean average precision (mAP) [40]. For classification, we trained one-vs-all linear SVMs [8] to classify 2D images and 3D shapes using their features. The average category accuracy [62] was used to evaluate the classification performance.

Parameters settings: In our 2D image recognition, the learning rate and batch size were 0.01 and 64 respectively. Our optimizer had a momentum of 0.9 and 0 weight decay rate. The regularized parameter λ in Equation (5) was set to be 5.0 while the re-scaling parameter γ in Equation (6) being 10. In the sketch-based 3D shape retrieval and 3D shape recognition experiments, the batch size was reduced to 32. Meanwhile, the learning rate, weight decay and momentum remained the same as what has been used in 2D experiments. We increased the regularized parameter λ to 10, which is the same as the re-scaling parameter γ .

5.2. Evaluation of Our Proposed Method

5.2.1 2D Image Recognition

Firstly, we empirically evaluated the effect of our proposed method on two broadly used 2D images benchmark datasets, i.e., *MNIST* and *CIFAR10*. Training images were randomly shuffled at the start of each epoch. In each training step, the optimal transportation plan \mathbf{T}^* between two batches of image features was approximated by iterating the

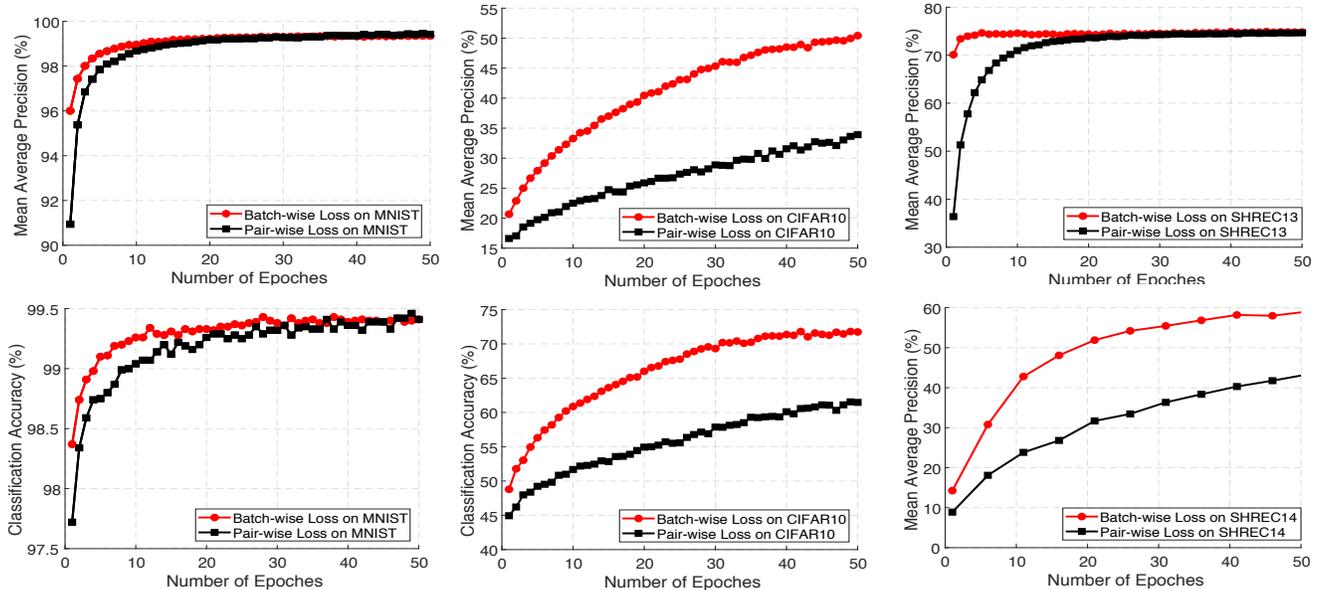


Figure 3. *Left*: Mean average precision (mAP) and classification accuracy curves of batch-wise optimal transport loss and pair-wise contrastive loss on 2D *MNIST* dataset. *Middle*: Comparison of their mAP and accuracy curves on 2D *CIFAR10* dataset. *Right*: Comparison of their mAP curves on sketch-based 3D shapes *SHREC13* and *SHREC14* dataset.

Method	NN	FT	ST	E	DCG	mAP
CDMR	0.279	0.203	0.296	0.166	0.458	0.250
SBR-VC	0.164	0.097	0.149	0.085	0.348	0.116
SP	0.017	0.016	0.031	0.018	0.240	0.026
FDC	0.110	0.069	0.107	0.061	0.307	0.086
Siamese	0.405	0.403	0.548	0.287	0.607	0.469
LWBR	0.712	0.725	0.725	0.369	0.814	0.752
N-pair	0.634	0.623	0.702	0.335	0.750	0.653
Lifted	0.688	0.705	0.774	0.363	0.803	0.734
Our Method	0.713	0.728	0.788	0.366	0.818	0.754

Method	NN	FT	ST	E	DCG	mAP
CDMR	0.109	0.057	0.089	0.041	0.328	0.054
SBR-VC	0.095	0.050	0.081	0.037	0.319	0.050
DB-VLAT	0.160	0.115	0.170	0.079	0.376	0.131
Siamese	0.239	0.212	0.316	0.140	0.496	0.228
DCML	0.272	0.275	0.345	0.171	0.498	0.286
LWBR	0.403	0.378	0.455	0.236	0.581	0.401
N-pair	0.300	0.270	0.321	0.150	0.484	0.289
Lifted	0.513	0.538	0.634	0.300	0.711	0.573
Our Method	0.536	0.564	0.629	0.305	0.712	0.591

Sinkhorn’s algorithm for 20 times. After each epoch, we computed all image features with the symmetrical network trained so far for classification and retrieval. The categorical accuracies provided by one-vs-rest linear SVMs and the retrieval mAPs given by the similarity measure based on the testing samples were recorded.

The left-hand and middle subfigures in Figure 3 present accuracy and mAP curves of the batch-wise optimal transport loss learning concerning the number of epochs. These figures illustrate the relationship between the convergence rate and recognition performance. Comparing with the pair-wise contrastive loss, our method possesses a significantly faster convergence rate. On *CIFAR10*, it provides a retrieval mAP and classification accuracy which are approximately 15% and 10% higher than the corresponding values achieved by the pair-wise loss at the end of 50 epochs. The empirical results indicate that the importance-driven distance metric learning can effectively adjust the distribution of weights. It pays more attention to the hard positive and

negative samples during the training process.

5.2.2 Sketch-based 3D Shape Retrieval

We then evaluated our method for sketch-based 3D shape retrieval on two large-scale benchmark datasets, i.e., *SHREC13* and *SHREC14*. The right-hand two subfigures in Figure 3 demonstrate the mAP curves of our batch-wise optimal transport loss as compared to the pair-wise loss objective. As illustrated, our method is about 5 times and 3 times faster than pair-wise loss on *SHREC13* and *SHREC14* respectively. Meanwhile, the retrieval performance is remarkably higher than the compared loss objective.

We also compared our method with several mainstream approaches for 3D shape retrieval, including CDMR [22], SBR-VC [33], SP [51], FDC [51], Siamese network [58], DCML [19], DB-VLAT [55], LWBR [63], N-pair loss [50], and Lifted loss [37]. The evaluation criteria include NN, FT, ST, E, DCG, and mAP.

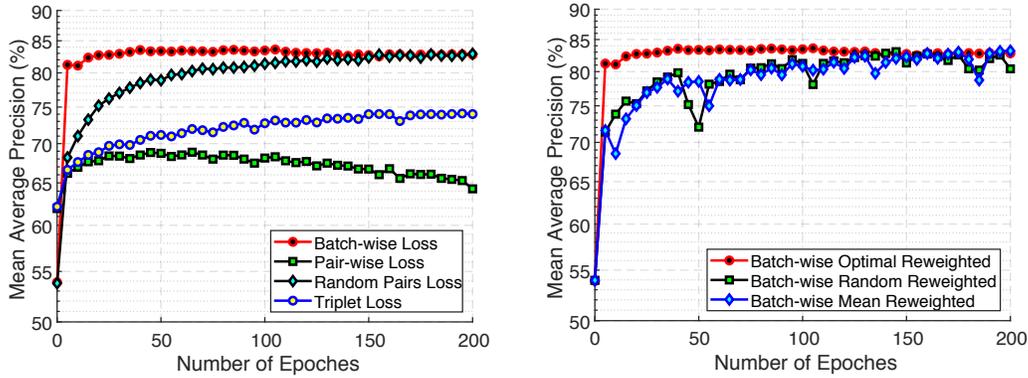


Figure 4. Mean average precision (mAP) curve with respect to the number of epochs evaluated of various methods on the *ModelNet40* dataset. Left subfigure illustrates the mAP curves of four loss objectives for 3D shape retrieval, and right subfigure illustrates the mAP curves of three weighting modesc. The mAP have been observed every five epochs for 200 epochs in figures.

Table 3. Comparisons of batch-wise optimal transport loss with other benchmark methods on the *ModelNet40* dataset.

		Evaluation Criteria						
Methods		NN	FT	ST	DCG	E	mAP (%)	Accuracy (%)
Pair-wise	Individual	0.8287	0.6544	0.7891	0.8562	0.5668	69.3%	88.6%
	Triplets	0.8324	0.6968	0.8029	0.8629	0.5927	74.1%	89.1%
	Random	0.8688	0.7948	0.9048	0.9140	0.6601	83.1%	89.5%
	Mean Weighted	0.8750	0.7986	0.9032	0.9158	0.6589	83.3%	89.7%
Batch-wise	Random Reweighted	0.8688	0.7673	0.8846	0.9051	0.6445	83.1%	89.0%
	Optimal Reweighted	0.8762	0.8013	0.8991	0.9178	0.6560	83.8%	90.3%

Table 4. Retrieval and classification results on the *ModelNet10* and *ModelNet40* datasets.

Methods	Shape Descriptor	<i>ModelNet10</i>		<i>ModelNet40</i>	
		mAP (%)	Accuracy (%)	mAP (%)	Accuracy (%)
(1) MVCNN [52]	2D View-based Descriptor (#Views=12)	N/A	N/A	80.2%	89.5%
	2D View-based Descriptor (#Views=80)	N/A	N/A	79.5%	90.1%
(2) GIFT [4]	2D View-based Descriptor (#Views=64)	91.1 %	92.3%	81.9%	83.1%
(3) 3DShapeNets [62]	3D Voxel Grid (30 × 30 × 30)	68.3%	83.5%	49.2%	77.0%
(4) Geometry Image [48]	2D Geometry Image	74.9%	88.4%	51.3%	83.9%
(5) PANORAMA-NN [45]	2D Panoramic View	87.4%	91.1%	83.5%	90.7 %
(6) DeepPano [46]	2D Panoramic View	84.1%	85.4%	76.8%	77.6%
(7) N-pair [50]	2D View-based Descriptor (#Views=12)	77.4%	86.8%	67.7%	79.1%
(8) Lifted [37]	2D View-based Descriptor (#Views=12)	88.8%	92.5%	83.0%	91.0%
Our Method	2D View-based Descriptor (#Views=12)	87.5%	93.7%	83.8%	90.3%

As summarized in Table 1 and Table 2, our batch-wise optimal transport loss based method achieved the best retrieval performance with respect to almost all evaluation metrics on *SHREC13* and *SHREC14*. Among compared methods, CDMR, DCML, Siamese network, and LWBR are all deep metric learning based approaches. They measured similarity based on pairs of samples and mapped data into an embedding metric space through different pooling schemes. In contrast, our proposed batch-wise optimal

transport loss objective can correctly re-weight the importance value of samples, mainly focus on those *hard samples*. Thus, our approach obtains better retrieval performance. Its mAP reaches 0.754, which is slightly better than LBWR and significantly better than other methods. Furthermore, the advantage of our approach is enlarged on *SHREC14* because this dataset has more severe intra-class and cross-modality variations. As a consequence, the mAP of our proposed method is 0.591, which is 0.190, 0.302, and 0.018

higher than LBWR, N-pair, and Lifted loss, respectively.

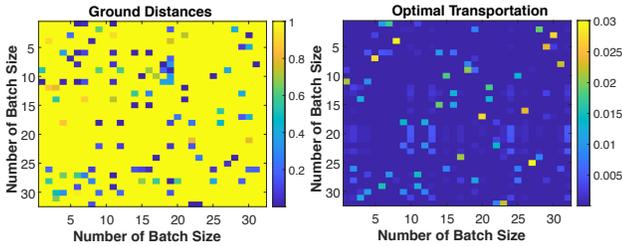


Figure 5. Illustration of relationship between the ground distances \mathbf{G}^* and the optimal transportation plan \mathbf{T}^* on the *ModelNet40* dataset. For two batches (each with batch size 32) of samples, we visualize the values of the batch-wise ground distances matrix (i.e., 32×32) and the corresponding optimal transportation plan.

5.2.3 3D Shapes Recognition

We finally verified the proposed method for 3D shape recognition on two large-scale 3D shape datasets, i.e., *ModelNet10* and *ModelNet40*. Pair-wise loss and triplet loss suffer from slow convergence rate because they are not capable of exploring all available semantical information within training batches simultaneously. To alleviate this problem, we used random sampling techniques (i.e., recurrently shuffle the training batches during each epoch) to loop over as many randomly sampled pairs as possible. It is expected that the random pairs based loss objective could make full use of all information so that the finally learned semantic metric could be balanced correctly. The left-hand subfigure in Figure 4 presents the mAP curves of batch-wise optimal transport loss and other compared loss objectives for 3D shape retrieval. Similarly, the batch-wise optimal transport loss objective still has the significantly faster convergence rate and can achieve a decent retrieval performance within a small number of epochs (i.e., 5 epochs).

We also examined two different probability distributions, i.e., uniformly distributed mean value ($\nu = \frac{1}{n^2}$) and random numbers in the interval $(0, 1)$, as alternatives of the optimal transportation plan. Uniformly distributed mean value weights in the batch-wise loss imply that samples are equally important for later metric learning. Uniformly distributed random weights randomly mark some samples as *hard samples* within a pair of batches during the learning process. The right-hand subfigure in Figure 4 illustrates comparison results of the retrieval performance concerning the number of epochs for these three re-weighting strategies. It demonstrates that the convergence rate of optimal re-weighted is much faster than the others.

The detailed comparison results are summarized in Table 3. We compared the batch-wise optimal transport loss with other designed benchmark methods using NN, FT, ST, E, DCG and mAP on the *ModelNet40* dataset. As illustrated in Figure 4 and Table 3, learning with batch-wise optimal transport loss objective has considerably faster convergence

rate than other benchmark methods. It takes only a few epochs (i.e., 5 epochs) to achieve mAP at 83.8% and accuracy at 90.3%, which are better than others after 200 epochs. It demonstrates that the learned optimal transportation plan can correctly re-weight the training samples according to their importance values during the metric learning process. Moreover, solving Equation (5) to learn optimal transportation plan \mathbf{T}^* is not computational expensive in practice. The average running time required by one epoch of individual pair loss objective is 2.51 seconds, and that of batch-wise optimal transport loss objective takes 9.02 seconds.

Here, we analyzed the role of the semantical information embedded ground distances $\mathbf{G}^* \in \mathbb{R}^{n \times n}$ and the optimal transportation plan $\mathbf{T}^* \in \mathbb{R}^{n \times n}$ in our work. The element \mathbf{G}_{ij}^* is filled with the distances of batch-wise similar positive samples \mathbf{G}_{ij}^+ and the distances of batch-wise dissimilar negative samples \mathbf{G}_{ij}^- . The right-hand subfigure in Figure 5 shows that far similar positive and adjacent dissimilar negative samples (i.e., *hard samples*) are sparsely distributed under our predefined ground distance metric \mathbf{G}^* . The left-hand subfigure is the optimal transportation plan which is actually a probability distribution [16]. The color map reveals the learned optimal transportation plan \mathbf{T}^* ensures higher importance weights to those few samples with small ground distance while giving less on the remaining ones.

In the end, we compared our method with state-of-the-art approaches for shape retrieval and classification, including MVCNN [52], GIFT [4], and DeepPano [46], et al. The detailed comparison results are summarized in Table 4. Compared to these approaches, our method based on batch-wise optimal transport loss learning can achieve the (almost) state-of-the-art performance on both tasks.

6. Conclusion

In this paper, we proposed a novel batch-wise optimal transport loss objective to learn an importance-driven distance metric. The learned distance metric can effectively emphasize *hard samples* according to their importance weights. We then formulated the proposed loss objective into an end-to-end deep metric learning network for recognition. We evaluated the performance and versatility of our method with various visual recognition tasks, including 2D image recognition, 2D sketch-based 3D shape cross-modality retrieval, and multiview based 3D shape recognition. The empirical results verified the proposed method could generically accelerate the convergence rate while achieving excellent recognition performance. Our future work will involve facilitating such a trend and applying this importance-driven distance metric learning to more widespread applications, for example, 3D point cloud classification, segmentation, 3D scene reconstruction, cross-modality correspondence among visual, audio, and text.

References

- [1] The Google 3d Warehouse. <https://3dwarehouse.sketchup.com/>, 2006.
- [2] The Turbosquid. <https://www.turbosquid.com/>, 2000.
- [3] The Princeton ModelNet. <http://modelnet.cs.princeton.edu/>, 2015.
- [4] S. Bai, X. Bai, Z. Zhou, Z. Zhang, and L. Jan Latecki. Gift: A real-time and scalable 3d shape search engine. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5023–5032, 2016.
- [5] S. Bell and K. Bala. Learning visual similarity for product design with convolutional neural networks. *ACM Transactions on Graphics (TOG)*, 34(4):98, 2015.
- [6] J. Bromley, I. Guyon, Y. LeCun, E. Säckinger, and R. Shah. Signature verification using a “siamese” time delay neural network. In *Advances in Neural Information Processing Systems*, pages 737–744, 1994.
- [7] A. M. Bronstein, M. M. Bronstein, L. J. Guibas, and M. Ovsjanikov. Shape google: Geometric words and expressions for invariant shape retrieval. *ACM Transactions on Graphics (TOG)*, 30(1):1, 2011.
- [8] C.-C. Chang and C.-J. Lin. Libsvm: a library for support vector machines. *ACM transactions on intelligent systems and technology (TIST)*, 2(3):27, 2011.
- [9] S. Chaudhuri and V. Koltun. Data-driven suggestions for creativity support in 3d modeling. *ACM Transactions on Graphics (TOG)*, 29(6):183, 2010.
- [10] G. Chechik, V. Sharma, U. Shalit, and S. Bengio. Large scale online learning of image similarity through ranking. *Journal of Machine Learning Research*, 11(Mar):1109–1135, 2010.
- [11] T. Y. Chen, H. Leung, and I. Mak. Adaptive random testing. In *Annual Asian Computing Science Conference*, pages 320–329. Springer, 2004.
- [12] S. Chopra, R. Hadsell, and Y. LeCun. Learning a similarity metric discriminatively, with application to face verification. In *Computer Vision and Pattern Recognition, 2005. CVPR 2005. IEEE Computer Society Conference on*, volume 1, pages 539–546. IEEE, 2005.
- [13] N. D. Cornea, M. F. Demirci, D. Silver, S. Dickinson, P. Kantor, et al. 3d object retrieval using many-to-many matching of curve skeletons. In *Shape Modeling and Applications, 2005 International Conference*, pages 366–371. IEEE, 2005.
- [14] T. Cover and P. Hart. Nearest neighbor pattern classification. *IEEE transactions on information theory*, 13(1):21–27, 1967.
- [15] Y. Cui, F. Zhou, Y. Lin, and S. Belongie. Fine-grained categorization and dataset bootstrapping using deep metric learning with humans in the loop. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1153–1162, 2016.
- [16] M. Cuturi. Sinkhorn distances: Lightspeed computation of optimal transport. In *Advances in neural information processing systems*, pages 2292–2300, 2013.
- [17] M. Cuturi and D. Avis. Ground metric learning. *Journal of Machine Learning Research*, 15(1):533–564, 2014.
- [18] M. Cuturi and A. Doucet. Fast computation of wasserstein barycenters. In *International Conference on Machine Learning*, pages 685–693, 2014.
- [19] G. Dai, J. Xie, F. Zhu, and Y. Fang. Deep correlated metric learning for sketch-based 3d shape retrieval. In *AAAI*, pages 4002–4008, 2017.
- [20] J. Donahue, Y. Jia, O. Vinyals, J. Hoffman, N. Zhang, E. Tzeng, and T. Darrell. Decaf: A deep convolutional activation feature for generic visual recognition. In *International conference on machine learning*, pages 647–655, 2014.
- [21] C. Frogner, C. Zhang, H. Mobahi, M. Araya, and T. A. Poggio. Learning with a wasserstein loss. In *Advances in Neural Information Processing Systems*, pages 2053–2061, 2015.
- [22] T. Furuya and R. Ohbuchi. Ranking on cross-domain manifold for sketch-based 3d model retrieval. In *Cyberworlds (CW), 2013 International Conference on*, pages 274–281. IEEE, 2013.
- [23] R. Girshick, J. Donahue, T. Darrell, and J. Malik. Rich feature hierarchies for accurate object detection and semantic segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 580–587, 2014.
- [24] R. Hadsell, S. Chopra, and Y. LeCun. Dimensionality reduction by learning an invariant mapping. In *Computer vision and pattern recognition, 2006 IEEE computer society conference on*, volume 2, pages 1735–1742. IEEE, 2006.
- [25] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.
- [26] E. Hoffer and N. Ailon. Deep metric learning using triplet network. In *International Workshop on Similarity-Based Pattern Recognition*, pages 84–92. Springer, 2015.
- [27] K. Järvelin, S. L. Price, L. M. Delcambre, and M. L. Nielsen. Discounted cumulated gain based evaluation of multiple-query ir sessions. In *European Conference on Information Retrieval*, pages 4–15. Springer, 2008.
- [28] J. Knopp, M. Prasad, G. Willems, R. Timofte, and L. Van Gool. Hough transform and 3d surf for robust three dimensional classification. *Computer vision—ECCV 2010*, pages 589–602, 2010.
- [29] I. Kokkinos, M. M. Bronstein, R. Litman, and A. M. Bronstein. Intrinsic shape context descriptors for deformable shapes. In *Computer Vision and Pattern Recognition (CVPR), 2012 IEEE Conference on*, pages 159–166. IEEE, 2012.
- [30] A. Krizhevsky, V. Nair, and G. Hinton. The cifar-10 dataset. online: <http://www.cs.toronto.edu/kriz/cifar.html>, 2014.
- [31] A. Krizhevsky, I. Sutskever, and G. E. Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*, pages 1097–1105, 2012.
- [32] Y. LeCun, L. Bottou, Y. Bengio, and P. Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998.
- [33] B. Li, Y. Lu, A. Godil, T. Schreck, M. Aono, H. Johan, J. M. Saavedra, and S. Tashiro. *SHREC’13 track: large scale sketch-based 3D shape retrieval*. 2013.

- [34] B. Li, Y. Lu, C. Li, A. Godil, T. Schreck, M. Aono, M. Burtscher, Q. Chen, N. K. Chowdhury, B. Fang, et al. A comparison of 3d shape retrieval methods based on a large-scale benchmark supporting multimodal queries. *Computer Vision and Image Understanding*, 131:1–27, 2015.
- [35] D. Maturana and S. Scherer. Voxnet: A 3d convolutional neural network for real-time object recognition. In *Intelligent Robots and Systems (IROS), 2015 IEEE/RSJ International Conference on*, pages 922–928. IEEE, 2015.
- [36] M. Norouzi, D. J. Fleet, and R. R. Salakhutdinov. Hamming distance metric learning. In *Advances in neural information processing systems*, pages 1061–1069, 2012.
- [37] H. Oh Song, Y. Xiang, S. Jegelka, and S. Savarese. Deep metric learning via lifted structured feature embedding. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4004–4012, 2016.
- [38] R. Osada, T. Funkhouser, B. Chazelle, and D. Dobkin. Shape distributions. *ACM Transactions on Graphics (TOG)*, 21(4):807–832, 2002.
- [39] O. Pele and M. Werman. Fast and robust earth mover’s distances. In *Computer vision, 2009 IEEE 12th international conference on*, pages 460–467. IEEE, 2009.
- [40] J. Philbin, O. Chum, M. Isard, J. Sivic, and A. Zisserman. Object retrieval with large vocabularies and fast spatial matching. In *Computer Vision and Pattern Recognition, 2007. CVPR’07. IEEE Conference on*, pages 1–8. IEEE, 2007.
- [41] D. Pickup, X. Sun, P. L. Rosin, R. Martin, Z. Cheng, Z. Lian, M. Aono, A. Ben Hamza, A. Bronstein, M. Bronstein, et al. Shrec’14 track: Shape retrieval of non-rigid 3d human models. In *Proceedings of the 7th Eurographics workshop on 3D Object Retrieval*, volume 1, page 6. Eurographics Association, 2014.
- [42] Q. Qian, R. Jin, S. Zhu, and Y. Lin. Fine-grained visual categorization via multi-stage metric learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3716–3724, 2015.
- [43] Y. Rubner, C. Tomasi, and L. J. Guibas. The earth mover’s distance as a metric for image retrieval. *International journal of computer vision*, 40(2):99–121, 2000.
- [44] F. Schroff, D. Kalenichenko, and J. Philbin. Facenet: A unified embedding for face recognition and clustering. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 815–823, 2015.
- [45] K. Sfikas, T. Theoharis, and I. Pratikakis. Exploiting the panorama representation for convolutional neural network classification and retrieval. In *Eurographics Workshop on 3D Object Retrieval*, 2017.
- [46] B. Shi, S. Bai, Z. Zhou, and X. Bai. Deeppano: Deep panoramic representation for 3-d shape recognition. *IEEE Signal Processing Letters*, 22(12):2339–2343, 2015.
- [47] K. Simonyan and A. Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014.
- [48] A. Sinha, J. Bai, and K. Ramani. Deep learning 3d shape surfaces using geometry images. In *European Conference on Computer Vision*, pages 223–240. Springer, 2016.
- [49] R. Sinkhorn. Diagonal equivalence to matrices with prescribed row and column sums. *The American Mathematical Monthly*, 74(4):402–405, 1967.
- [50] K. Sohn. Improved deep metric learning with multi-class n-pair loss objective. In *Advances in Neural Information Processing Systems*, pages 1857–1865, 2016.
- [51] P. Sousa and M. J. Fonseca. Sketch-based retrieval of drawings using spatial proximity. *Journal of Visual Languages & Computing*, 21(2):69–80, 2010.
- [52] H. Su, S. Maji, E. Kalogerakis, and E. Learned-Miller. Multi-view convolutional neural networks for 3d shape recognition. In *Proceedings of the IEEE international conference on computer vision*, pages 945–953, 2015.
- [53] C. Szegedy, W. Liu, Y. Jia, P. Sermanet, S. Reed, D. Anguelov, D. Erhan, V. Vanhoucke, and A. Rabinovich. Going deeper with convolutions. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1–9, 2015.
- [54] J. W. Tangelder and R. C. Veltkamp. A survey of content based 3d shape retrieval methods. In *Shape Modeling Applications, 2004. Proceedings*, pages 145–156. IEEE, 2004.
- [55] A. Tatsuma, H. Koyanagi, and M. Aono. A large-scale shape benchmark for 3d object retrieval: Toyohashi shape benchmark. In *Signal & Information Processing Association Annual Summit and Conference (APSIPA ASC), 2012 Asia-Pacific*, pages 1–10. IEEE, 2012.
- [56] S. Vallender. Calculation of the wasserstein distance between probability distributions on the line. *Theory of Probability & Its Applications*, 18(4):784–786, 1974.
- [57] C. Villani. *Optimal transport: old and new*, volume 338. Springer Science & Business Media, 2008.
- [58] F. Wang, L. Kang, and Y. Li. Sketch-based 3d shape retrieval using convolutional neural networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1875–1883, 2015.
- [59] J. Wang, Y. Song, T. Leung, C. Rosenberg, J. Wang, J. Philbin, B. Chen, and Y. Wu. Learning fine-grained image similarity with deep ranking. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1386–1393, 2014.
- [60] K. Q. Weinberger and L. K. Saul. Distance metric learning for large margin nearest neighbor classification. *Journal of Machine Learning Research*, 10(Feb):207–244, 2009.
- [61] Y. Wu and S. Verdú. Witsenhausen’s counterexample: A view from optimal transport theory. In *Decision and Control and European Control Conference (CDC-ECC), 2011 50th IEEE Conference on*, pages 5732–5737. IEEE, 2011.
- [62] Z. Wu, S. Song, A. Khosla, F. Yu, L. Zhang, X. Tang, and J. 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.
- [63] J. Xie, G. Dai, F. Zhu, and Y. Fang. Learning barycentric representations of 3d shapes for sketch-based 3d shape retrieval. In *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3615–3623. IEEE, 2017.