

Rethinking Loss Design for Large-scale 3D Shape Retrieval

Zhaoqun Li^{1*}, Cheng Xu^{1*} and Biao Leng^{1,2,3†}

¹School of Computer Science and Engineering, Beihang University, Beijing, 100191

²Beijing Advanced Innovation Center for Big Data and Brain Computing, Beihang University, Beijing, 100191

³State Key Laboratory of Virtual Reality Technology and Systems, Beihang University, Beijing, 100191
 {lizhaoqun, cxu, lengbiao}@buaa.edu.cn

Abstract

Learning discriminative shape representations is a crucial issue for large-scale 3D shape retrieval. In this paper, we propose the Collaborative Inner Product Loss (CIP Loss) to obtain ideal shape embedding that discriminative among different categories and clustered within the same class. Utilizing simple inner product operation, CIP loss explicitly enforces the features of the same class to be clustered in a linear subspace, while inter-class subspaces are constrained to be at least orthogonal. Compared to previous metric loss functions, CIP loss could provide more clear geometric interpretation for the embedding than Euclidean margin, and is easy to implement without normalization operation referring to cosine margin. Moreover, our proposed loss term can combine with other commonly used loss functions and can be easily plugged into existing off-the-shelf architectures. Extensive experiments conducted on the two public 3D object retrieval datasets, ModelNet and ShapeNetCore 55, demonstrate the effectiveness of our proposal, and our method has achieved state-of-the-art results on both datasets.

1 Introduction

3D shape retrieval is a fundamental problem in 3D shape analysis communities, spanning wide applications from medical imaging to robot navigation. With the development of Convolutional Neural Network (CNN) and the emergence of large-scale 3D repositories [Chang *et al.*, 2015; Wu *et al.*, 2015], numerous approaches are proposed [Johns *et al.*, 2016; Su *et al.*, 2015; Xu *et al.*, 2018; Li *et al.*, 2019] which significantly boost the performance in 3D shape retrieval task. Among these methods, view-based methods have achieved the best performance so far. In view-based 3D shape retrieval, images are first projected from different viewpoints of a 3D shape, and then they are passed into CNNs to obtain the discriminative and informative shape representation. The crucial

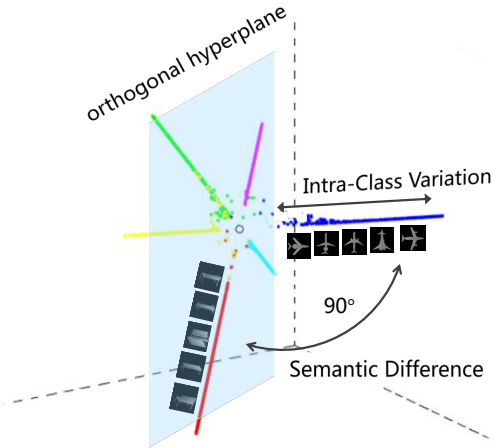


Figure 1: Illustration of CIP loss effect in three-dimension. Different colors represent different classes which divide equally the whole space. In this distribution, the features of different classes are at least orthogonal and that of the same class are in a line. The norm of features corresponds to the intra-class variation and the angle corresponds to the semantic difference.

issue of 3D shape retrieval is how to obtain ideal shape representations: *discriminative among different categories and clustered within the same class.*

The majority of view-based methods like MVCNN [Su *et al.*, 2015] train a CNN to extract shape representations with the standard softmax loss. While softmax loss does not explicitly enforce the distance between features in the embedding. To obtain ideal shape representations, [He *et al.*, 2018] recently introduce deep metric loss functions to the 3D shape retrieval task. Such metric learning techniques, like center loss [Wen *et al.*, 2016], triplet loss [Schroff *et al.*, 2015] and triplet-center loss (TCL) [He *et al.*, 2018], are conducted in the Euclidean space and could remarkably boost the retrieval performance. However, these loss functions are designed by Euclidean margin, which is difficult to be determined because of the large span of Euclidean distance and has weak geometric constraint due to the variance of feature vector magnitude. With more clear geometric interpretation, metric loss functions based on cosine margin like coco loss [Liu *et al.*, 2017]

*They contributed equally to this work

†Corresponding author

are particularly popular in the image retrieval and face verification communities. Considering the superior geometric constraint of cosine margin, we also introduce the coco loss to 3D shape retrieval in our experiment. The drawbacks with these methods lie in the complicated back-propagation calculation and the unstable parameter update, especially for large-scale 3D shape datasets. Therefore, we seek for a simple and stable metric loss with clear geometric interpretation to improve the discriminability of shape representations.

In this work, our intuition is that, for an ideal embedding space, the features of the same class should be in a line while at least orthogonal to other features. Meanwhile, the simplicity and stability of the loss function must be taken into consideration in the loss design for large-scale dataset. To this end, we propose a new Collaborative Inner Product Loss function (CIP Loss) to jointly optimize intra-class similarity and inter-class margin of the learned shape representations. In the loss design, we directly adopt the simple inner product which could provide distinct geometric constraints (see Fig. 1) and enforce learned visual features f to satisfy two conditions:

$$Dot(f_i, f_j) \rightarrow \infty \quad \text{if } y_i = y_j \quad (1)$$

$$Dot(f_i, f_j) \leq 0 \quad \text{if } y_i \neq y_j \quad (2)$$

where y is the label and $Dot(\cdot, \cdot)$ means the inner product between two vectors. On one hand, CIP loss encourages visual features of the same class to cluster in a linear subspace where the inner product between features tends to infinity, indicating the maximum extent of linear correlation among features. On the other hand, inter-class subspaces are requested to be at least orthogonal, meaning that for each category subspace, other category subspaces are squeezed to another half-space of its orthogonal hyperplane. In particular, if there are six categories in three-dimensional feature space, CIP loss would make all categories divide the whole embedding space equally as shown in Fig. 1. Specially, in order to save computation, the inner product between visual features is replaced by the inner product between visual features and category centerlines in our method.

We show in Fig. 2 the effect of different loss functions on the feature distribution. Compared with previous metric learning techniques in 3D shape retrieval community, the proposed method has the following key advantages. 1) CIP loss could provide more explicit geometric interpretation for the embedding than approaches based on Euclidean distance, see Fig. 2(b)(c)(e). 2) Compared with popular cosine margin, CIP loss is simple and elegant to implement without margin design and normalization operation, which can stabilize the optimization process. 3) CIP loss is flexible to be plugged into existing off-the-shelf architectures, where it can work standalone or in combination with other loss functions.

In summary, our main contributions are as follows.

- We propose a novel metric loss function, namely collaborative inner product loss (CIP Loss), which adopts elegant inner product between features to perform more explicit geometric constraints on the shape embedding.
- Two components of CIP loss are proposed, namely Cluster loss and Ortho loss. Cluster loss guarantees visual

features of the same class to cluster in a linear subspace, while Ortho loss enforces inter-class subspaces are at least orthogonal, making shape representations more discriminative.

- Our method achieves the state-of-the-art in the large-scale datasets, ModelNet and ShapeNetCore55, showing the effectiveness of CIP loss.

2 Related Work

In this section, we briefly review some view-based approaches of 3D shape retrieval and introduce some typical deep metric learning methods in this domain.

2.1 Multi-view 3D Shape Retrieval

3D shape retrieval methods could be roughly divided into two categories: model-based methods and view-based methods. Although the model-based methods [Xie *et al.*, 2017; Qi *et al.*, 2017; Shen *et al.*, 2018] can capture the 3D structure of original shape by various data form, their performances are typically lower than view-based methods due to the complexity of computation and the noise in shape representation.

The view-based 3D shape retrieval methods use rendered images to represent the original shape. MVCNN [Su *et al.*, 2015] is a typical method adopting CNN to aggregate rendered images. In this method, the rendered images from different views are pooled to generate the shape feature. [Huang *et al.*, 2018] adopts a local MVCNN shape descriptors, which use a local shape descriptor of each point on the shape to analyze the shape. Recently proposed GVCNN [Feng *et al.*, 2018] separates views in groups by intrinsic hierarchical correlation and discriminability among them, which largely improves the performance on the 3D shape retrieval. Another technique used in 3D shape retrieval is re-ranking. GIFT [Bai *et al.*, 2016] adopts CNN with GPU acceleration to extract each single view feature and proposes the inverted file to reduce computation for fast retrieval.

The aforementioned methods are trained under the supervision of softmax loss which is not always consistent with retrieval task. [He *et al.*, 2018] recently puts forward TCL that combines the center loss and triplet loss, which achieve state-of-the-art results on various datasets. This shows that deep metric learning plays an important role in the 3D retrieval task.

2.2 Deep Metric Learning

As a key part of deep learning framework, loss design has been studied widely in retrieval and other domains. Most commonly used loss functions in 3D shape retrieval are designed in Euclidean space, such as triplet loss, center loss and TCL. Triplet loss [Schroff *et al.*, 2015] forces inter-class distance exceed intra-class distance by a positive Euclidean margin which is widely applied in face recognition. However, the original triplet loss has the problem in computation cost that the number of triplets grows cubically with the dataset. To tackle this problem, many improved version [Song *et al.*, 2016; 2016] are proposed in various domains. [Song *et al.*, 2016] is an efficient approach that proposes an algorithm for

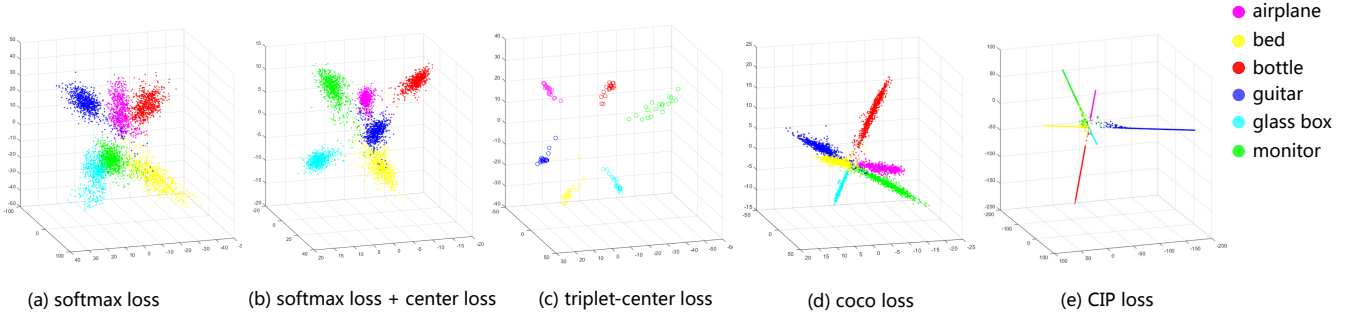


Figure 2: Illustrative comparison between CIP loss and other loss functions on ModelNet test set. In three dimension, CIP loss can converge with at most 6 classes because Ortho loss constrains the features of different classes to be at least orthogonal. So we select 6 classes from ModelNet and train CNNs on the corresponding train set. We show all the view features in (a)(b)(d)(e). For (c), only model features are available since triplet-center loss adopts MVCNN architecture. Compared with other loss functions, CIP loss imposes the embedding space a more strict geometric constraint that makes the features orthogonal between different classes and clustered in a linear subspace with the same class. Best viewed in color.

taking full advantage of the training samples by using the matrix of pairwise distances and achieves high performance on image retrieval. Another popular metric loss is center loss [Wen *et al.*, 2016] that tries to gather the same class features in one point which is called center of the class. Combining the advantages of triplet loss and center loss, TCL [He *et al.*, 2018] is proposed for 3D shape retrieval which gains a better performance than the other loss functions. However, they adopt the Euclidean margin which is difficult to design and has a weak geometric constraint.

On the other hand, cosine margin is recently proposed which are popular in face recognition. As a typical method, coco loss [Liu *et al.*, 2017] advocates that the weakness of softmax loss is in the final classification layer. By weight normalization and feature normalization, it introduces cosine margin to reduce intra-class variation. Although the cosine margin has more clear geometric interpretation than Euclidean margin, it also brings instability in the optimization.

3 Methods

Deep metric learning aims to train a neural network, denoted by $f_\theta(\cdot)$, which maps the original data onto a new embedding space $\mathbb{F} \in \mathbb{R}^n$. For writing convenience, we use $f_i \in \mathbb{R}^n$ to represent $f_\theta(x_i)$ which is the extracted feature from CNN. In this work, our goal is to make \mathbb{F} ideal, which means that the intra-class angular distance is 0 while the inter-class angular distance is at least $\frac{\pi}{2}$. To fulfill this goal, the loss design needs to consider two necessary tasks: (1) *Enlarging the inter-class distances between features* and (2) *Reducing the intra-class distances between features*. Then, the CIP loss could be formalized as the sum of pull term and push term corresponding to Cluster loss and Ortho loss respectively:

$$L_{CIP} = L_{cluster} + \lambda \cdot L_{ortho} \quad (3)$$

where λ is a trade-off hyperparameter.

To achieve the aforementioned goal, a natural idea is adopting cosine distance which needs normalization of feature vectors, while this operation would bring complexity in back-propagation and instability in optimization that we will discuss in Sec. 3.3. In this article, we use the inner product of

two vectors as the pairwise similarity for the sake of simplicity and stability:

$$s(f_i, f_j) = f_i^\top \cdot f_j \quad (4)$$

3.1 Proposed Loss Functions

Given a training dataset $\mathcal{X} \in \mathbb{R}^D$, let $\mathcal{Y} = \{y_1, y_2, \dots, y_{|\mathcal{X}|}\}$ denote the corresponding label set, $y_i \in \{1, 2, \dots, K\}$ where K is the number of classes.

The Cluster loss constrains the same class features in a linear subspace, *i.e.* a line in \mathbb{F} . Inspired by center loss, we define this line, denoted by $c_k \in \mathbb{R}^n$, as the centerline of the k^{th} class. Different from the interpretation of “center” in center loss, the centerline here represents the direction of the corresponding class features. Given a batch of training data with M samples, the pull term named Cluster loss is defined as:

$$L_{cluster} = \sum_{i=1}^M \frac{1}{f_i^\top \cdot c_{y_i} + d} \quad (5)$$

where $d = 2$ is a constant value for numerical stability. Cluster loss encourages f_i to pull its corresponding centerline c_{y_i} . And in turn, it forces the centerline to concentrate its corresponding features.

The design of Ortho loss considers that the different class features are at least orthogonal, according to condition 2. As centerline represents the direction of the corresponding class, we penalize the features which are not orthogonal to other negative centerlines. The push term named Ortho loss is formulated as:

$$L_{ortho} = \sum_{i=1}^M \left[\sum_{k=1, k \neq y_i}^K \max(f_i^\top \cdot c_k, 0) \right] \quad (6)$$

Ortho loss pushes f_i to be at least orthogonal to all other centers so that the inter-class distance will increase. The generated feature distribution of our method is shown in Fig. 2(e).

3.2 Backpropagation

Gradient of Feature Vector

Owing to the simple formulation of loss functions and adoption of the inner product, the formulas of gradient for feature are also simple. For writing convenience, we denote $(\cdot)_+ = \max(\cdot, 0)$. We use $\delta(cdt) = 1$ to indicate cdt is true and $\delta(cdt) = 0$ otherwise. The gradient of f_i for L_{ortho} is formulated as follows:

$$\frac{\partial L_{ortho}}{\partial f_i} = \sum_{k=1, k \neq y_i}^K \delta(f_i^T \cdot c_k > 0) \cdot c_k \quad (7)$$

As for $L_{cluster}$, the existence of $d > 0$ is to adjust the gradient of function $\frac{1}{x}$ which explodes on $x = 0$. We set $d = 2$ in our implementation. However the gradient function still has this problem when $f_i^T \cdot c_{y_i}$ is near to $-d$. In fact, the original formulation is:

$$\left(\frac{\partial L_{cluster}}{\partial f_i} \right)_{origin} = -\frac{1}{(f_i^T \cdot c_{y_i} + d)^2} \cdot c_{y_i} \quad (8)$$

Eq. 8 becomes ∞ when $f_i^T \cdot c_{y_i}$ tends to be $-d$. Thus we clip the term $f_i^T \cdot c_{y_i}$ by 0 so that the optimization becomes stable. Concretely, we use a surrogate gradient in backpropagation:

$$\left(\frac{\partial L_{cluster}}{\partial f_i} \right)_{srg} = -\frac{c_{y_i}}{((f_i^T \cdot c_{y_i})_+ + d)^2} \quad (9)$$

Gradient of Centerline

The gradient form is similar to that of feature vector since the inner product is a symmetric function. In the same way, we adopt a surrogate gradient to evade gradient explosion for $L_{cluster}$.

$$\left(\frac{\partial L_{cluster}}{\partial c_i} \right)_{srg} = \sum_{j=1}^M -\frac{\delta(y_j = i) \cdot f_j}{((f_j^T \cdot c_i)_+ + d)^2} \quad (10)$$

For L_{ortho} , we replace the gradient of L_{ortho} with respect to c_i by a ‘‘average’’ version which can further stabilize the update of center:

$$\left(\frac{\partial L_{ortho}}{\partial c_i} \right)_{srg} = \frac{\sum_{j=1}^M \delta(y_j \neq i) \cdot \delta(f_j^T \cdot c_i > 0) \cdot f_j}{1 + \sum_{j=1}^M \delta(y_j \neq i) \cdot \delta(f_j^T \cdot c_i > 0)} \quad (11)$$

3.3 Discussion

The loss function is essential for guiding the optimization that has a direct influence on the feature distribution. And in the training process, the convergence and the stability should be guaranteed, which poses challenges to the loss design. In this section, we will illustrate the motivation and the procedure of loss design in detail.

Inner Product without Normalization

In a typical CNN, the similarity between a deep feature vector f_i and a layer weight w_j is encoded by the inner product: $s = f_i^T \cdot w_j = \|f_i\|_2 \|w_j\|_2 \cos(f_i, w_j)$. In order to eliminate the effect of norms on similarity, some methods adopt normalization operations. However, weight normalization leads to instability. We can infer from the gradient of weight after normalization:

$$\frac{\partial}{\partial w} \left(\frac{w^T \cdot f}{\|w\|_2} \right) = \frac{f}{\|w\|_2} - \frac{(w^T \cdot f) \cdot w}{\|w\|_2^3} \quad (12)$$

From Eq. 12 we can see that small value of w will lead to gradient explosion. In comparison, for the inner product, the gradient is $\frac{\partial(w^T \cdot f)}{\partial w} = f$ which is more stable and computationally efficient. So we choose the inner product for the sake of stabilization in the training process. As for leaving out feature normalization, we take the same consideration.

In addition, this form doesn’t need margin design or, from another perspective, we can regard the angular margin as 0 and $\frac{\pi}{2}$ in two terms corresponding that the inner product is ∞ and 0.

Combination with Other Loss Functions

Both of Cluster loss and Ortho loss can combine solely with softmax loss and bring improvement on performance. It is worth pointing out that Ortho loss cannot maintain the centerlines solely since it focuses on enlarging inter-class distance. When we combine L_{ortho} with other loss functions, the centerlines will diverge. In this case, we employ the batch version:

$$L_{ortho, batch} = \sum_{i=1}^M \left[\sum_{j=1, y_j \neq y_i}^M \max(f_i^T \cdot f_j, 0) \right] \quad (13)$$

This batch version directly conducts the optimization on feature distribution without centerlines. The backpropagation formula of batch version is similar to Eq. 7 that we only need to replace the role of c_k by f_j . We will discuss the performance difference of these two versions in the Sec. 4.2.

4 Experiment

In this section, we evaluate the performance of our method on two representative large-scale 3D shape datasets. We first compare the results with the state-of-art methods. We also provide massive experiments to discuss the comparison and the combination with other metric loss functions. In the last part, we investigate the influence of the hyper-parameters and visualize the experiment results by dimension reduction.

4.1 Retrieval on Large-scale 3D Datasets

Implementation detail. The experiments are conducted on Nvidia GTX1080Ti GPU and our methods are implemented by Caffe. For the structure of the CNN, we use VGG-M [Chatfield *et al.*, 2014] pre-trained on ImageNet [Deng *et al.*, 2009] as the base network in all our experiments. We use the stochastic gradient descent (SGD) algorithm with momentum $2e-4$ to optimize the loss and the batch size is 100. The initial learning rate is 0.01 and is divided by 5 at the 20th epoch. The

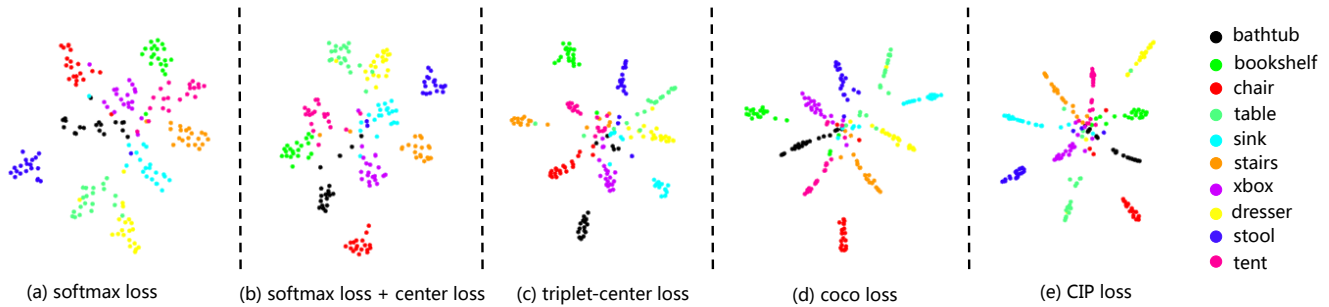


Figure 3: Visualization of experiment results on ModelNet40. We randomly select 10 classes from the dataset and use t-SNE to reduce the dimension of shape features for visualization. Best viewed in color.

Methods	AUC	MAP
ShapeNets	49.94%	49.23%
DeepPano	77.63%	76.81%
MVCNN	-	80.20%
GIFT	83.10%	81.94%
Siamese CNN-BiLSTM	-	83.30%
PANORAMA-NN	87.39%	83.45%
GVCNN	-	85.70%
RED	87.03%	86.30%
TCL	87.60%	86.70%
L_{CIP}	87.56%	86.49%
L_{CIP} + softmax loss	88.07%	87.08%
L_{CIP} + center loss	88.21%	87.22%

Table 1: The performance comparison with state-of-the-arts on ModelNet40.

total training epoch is 30. The centerlines are initialized by a Gaussian distribution of mean value 0 and standard deviation 0.01. Firstly, we render 36 views of each model by Phong reflection to generate depth images which compose our image dataset. The size of each image is 224x224 pixels in our experiment. Then, we randomly select images from the image dataset to train our CNN. In the test phase, the features of the penultimate layer, *i.e.* $fc7$, are extracted for the evaluation. We obtain the shape feature vector by averaging all view features of this shape. The cosine distance is adopted as the evaluation metric.

Dataset. We select two representative datasets, ModelNet and ShapeNetCore55, to conduct the evaluation of our methods. 1) ModelNet Dataset: this dataset is a large-scale 3D CAD model dataset composed of 127,915 3D CAD models from 662 categories. ModelNet40 and ModelNet10 are two subsets which contain 40 categories and 10 categories respectively. In our experiment, we follow the training and testing split as mentioned in [Wu *et al.*, 2015]. 2) ShapeNetCore55: this dataset from SHape REtrieval Contest 2016 [Savva *et al.*, 2016] is composed of 51,190 3D shapes from 55 common categories divided into 204 sub-categories. We follow the official training and testing split to conduct our experiment. The evaluation metrics used in this paper include mean average precision (MAP), area under curve (AUC), F-measure (F1)

Methods	Micro			Macro		
	F1	MAP	NDCG	F1	MAP	NDCG
Wang	24.6	60.0	77.6	16.3	47.8	69.5
Li	53.4	74.9	86.5	18.2	57.9	76.7
Kd-network	45.1	61.7	81.4	24.1	48.4	72.6
MVCNN	61.2	73.4	84.3	41.6	66.2	79.3
GIFT	66.1	81.1	88.9	42.3	73.0	84.3
TCL	67.9	84.0	89.5	43.9	78.3	86.9
TCL(VGG)	64.5	82.1	88.4	36.5	71.0	82.7
Our	67.4	83.6	89.7	46.1	75.4	85.8

Table 2: The performance (%) comparison on SHREC16 perturbed dataset.

and normalized discounted cumulative gain (NDCG). Refer to [Wu *et al.*, 2015; Savva *et al.*, 2016] for their detailed definitions.

Comparison with state-of-the-arts. On ModelNet40 dataset, we choose 3D ShapeNets [Wu *et al.*, 2015], DeepPano [Shi *et al.*, 2015], MVCNN [Su *et al.*, 2015], PANORAMA-NN [Sfikas *et al.*, 2017], RED [Song *et al.*, 2017], Siamese CNN-BiLSTM [Dai *et al.*, 2018], GVCNN [Feng *et al.*, 2018], TCL [He *et al.*, 2018] and GIFT [Bai *et al.*, 2016] methods for comparison. The experimental results and comparison among different methods are presented in Tab. 1. Our method (L_{CIP} + center loss) achieves retrieval AUC of 88.21% and MAP of 87.22% which is the best among different methods. Besides, as the current state-of-the-art view-based method on ModelNet40, TCL is trained on VGG_11 which has 3 more convolution layers than VGG-M and adopts batch normalization [Ioffe and Szegedy, 2015]. Compared with it, L_{CIP} only loses 0.21% of MAP on performance. We also re-conduct TCL using VGG-M according to the parameter settings in the paper (refer to Tab. 3). For the evaluation in ShapeNetCore55 dataset, we carry out our method (L_{CIP} + softmax loss) on the more challenging perturbed version. For a fair comparison, we re-implement “TCL + softmax loss” on VGG-M which is indicated as TCL(VGG) in Tab. 2. We also choose Wang [Savva *et al.*, 2016], Li [Savva *et al.*, 2016], K-d network [Klokov and Lempitsky, 2017], MVCNN [Su *et al.*, 2015] and GIFT [Bai *et al.*, 2016] methods for comparison.

Loss function	AUC	MAP
softmax loss	81.16%	79.91%
center loss + softmax loss	83.20%	82.04%
triplet loss + softmax loss	83.51%	82.40%
triplet-center loss(VGG)	85.70%	84.64%
coco loss(VGG)	86.76%	85.72%
$L_{cluster}$ + softmax loss	83.11%	81.94%
$L_{ortho,batch}$ + softmax loss	87.16%	86.16%
$L_{cluster}$ + $L_{ortho,batch}$	87.21%	86.17%
L_{CIP} ($L_{cluster}$ + L_{ortho})	87.56%	86.49%
L_{CIP} + softmax loss	88.07%	87.08%
L_{CIP} + center loss	88.21%	87.22%

Table 3: The performance comparison with different loss functions on ModelNet40.

We take two types of results in the competition, namely macro and micro. As shown in Tab. 2, our method achieves the state-of-the-art performance.

4.2 Comparison with Other Loss Functions

To demonstrate the efficiency of our loss functions, we set various comparison experiments on ModelNet40 dataset. Besides the commonly used loss functions, we also carry out two well-known methods, coco loss and triplet-center loss, on VGG-M. The experiment results are shown in Tab. 3. Our comparison contains 3 parts:

The first is the comparison with commonly used metric loss functions. As a well-designed deep metric loss, L_{CIP} dramatically improves the performance in MAP by 6.58% from softmax loss supervision. And our method outperforms TCL and coco loss which are improved efficient methods. Moreover, for the comparison of L_{ortho} and $L_{ortho,batch}$, we can see that they achieve similar results when combining with $L_{cluster}$. It also should be noted that the optimization will diverge if the CNN is trained solely with $L_{cluster}$ or L_{ortho} because both two terms are necessary.

The second part explores the combination with other loss functions. Both of our loss functions can combine softmax loss separately and the results show that the combinations bring a significant improvement. Compared with center loss, $L_{cluster}$ achieves 81.94% in MAP when jointly trained with softmax loss, which is evidence that $L_{cluster}$ has a comparable polymerization ability. And to prove the efficiency of L_{ortho} , we conduct “ L_{ortho} + softmax loss”. This combination brings a large improvement of 5.31% in MAP compared with single softmax loss, which means L_{ortho} has a powerful capacity in enlarging inter-class distance.

Finally, we train the CNN with three loss functions and obtain an impressive performance. “ L_{CIP} + softmax loss” achieves a MAP of 87.08% and “ L_{CIP} + center loss” achieves a MAP 87.22% on Modelnet40. Note that the loss weight of softmax loss and center loss is 0.1 and 0.0003 respectively in the combination with our loss function. All these experiments demonstrate our loss design can generate robust and discriminative features.

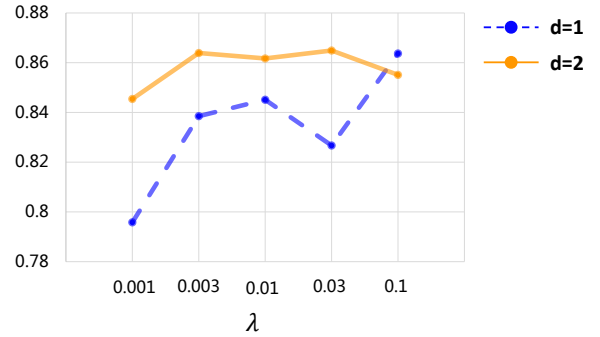


Figure 4: The retrieval performance (MAP) of different hyper-parameter settings on ModelNet40.

4.3 Discussion

Sensitiveness of hyper-parameter. In Eq. 5, d is a constant value for stability. In our method, we choose $d = 2$ because this choice can reduce the sensitiveness of trade-off parameter λ . To investigate the influence of d and λ on the performance, we conduct experiments supervised by CIP loss on ModelNet40. The experiment results are shown in Fig. 4. We can see that CIP loss can obtain state-of-the-art results with different values of d by adjusting λ , but the sensitiveness of λ is different. The performance varies a lot when $d = 1$. With $d = 2$, CIP loss can converge stably with a wide range of λ , which means our loss design guarantees stability in the optimization.

Visualization of experiment results. We use t-SNE [van der Maaten and Hinton, 2008] to visualize the experiment results with different loss functions in Fig. 3. From the figure, we can see that under the supervision of CIP loss, the distance between classes is increased compared with softmax loss and center loss. Compared with TCL and coco loss, the intra-class variance is further decreased with CIP loss and the geometric separation become more clear.

5 Conclusion

In this paper, we propose a novel loss function named Collaborative Inner Product Loss with regard to large-scale 3D shape retrieval task. The Inner product is employed in the loss design which not only imposes a more strict constraint but also guarantees stability in optimization. The proposed loss function consists of Cluster loss and Ortho loss that play different roles: one for reducing the intra-class distance and the other for enlarging the inter-class margin, and both of them can combine with other commonly used loss functions. Experimental results on two large-scale datasets have proven the superiority of our loss functions.

Acknowledgements

This work is supported by the Beijing Municipal Natural Science Foundation (No.L182014), and the Open Project Program of State Key Laboratory of Virtual Reality Technology and Systems, Beihang University (No.VRLAB2019C05).

References

- [Bai *et al.*, 2016] Song Bai, Xiang Bai, Zhichao Zhou, Zhaoxiang Zhang, and Longin Jan Latecki. Gift: A real-time and scalable 3d shape search engine. In *CVPR*, pages 5023–5032, 2016.
- [Chang *et al.*, 2015] 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.
- [Chatfield *et al.*, 2014] Ken Chatfield, Karen Simonyan, Andrea Vedaldi, and Andrew Zisserman. Return of the devil in the details: Delving deep into convolutional nets. 2014.
- [Dai *et al.*, 2018] Guoxian Dai, Jin Xie, and Yi Fang. Siamese cnn-bilstm architecture for 3d shape representation learning. In *IJCAI*, pages 670–676, 2018.
- [Deng *et al.*, 2009] Jia Deng, Wei Dong, Richard Socher, Li Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *CVPR*, pages 248–255, 2009.
- [Feng *et al.*, 2018] Yifan Feng, Zizhao Zhang, Xibin Zhao, Rongrong Ji, and Yue Gao. Gvcnn: Group-view convolutional neural networks for 3d shape recognition. In *CVPR*, pages 264–272, 2018.
- [He *et al.*, 2018] Xinwei He, Yang Zhou, Zhichao Zhou, Song Bai, and Xiang Bai. Triplet-center loss for multi-view 3d object retrieval. In *CVPR*, 2018.
- [Huang *et al.*, 2018] Haibin Huang, Evangelos Kalogerakis, Siddhartha Chaudhuri, Duygu Ceylan, Vladimir G Kim, and Ersin Yumer. Learning local shape descriptors from part correspondences with multiview convolutional networks. *ACM Transactions on Graphics (TOG)*, 2018.
- [Ioffe and Szegedy, 2015] Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. *CoRR*, 2015.
- [Johns *et al.*, 2016] Edward Johns, Stefan Leutenegger, and Andrew J Davison. Pairwise decomposition of image sequences for active multi-view recognition. In *CVPR*, pages 3813–3822, 2016.
- [Klokov and Lempitsky, 2017] Roman Klokov and Victor Lempitsky. Escape from cells: Deep kd-networks for the recognition of 3d point cloud models. In *ICCV*, pages 863–872, 2017.
- [Li *et al.*, 2019] Zhaoqun Li, Cheng Xu, and Biao Leng. Angular triplet-center loss for multi-view 3d shape retrieval. In *AAAI*, 2019.
- [Liu *et al.*, 2017] Yu Liu, Hongyang Li, and Xiaogang Wang. Rethinking feature discrimination and polymerization for large-scale recognition. *arXiv preprint arXiv:1710.00870*, 2017.
- [Qi *et al.*, 2017] Charles Ruizhongtai Qi, Li Yi, Hao Su, and Leonidas J Guibas. Pointnet++: Deep hierarchical feature learning on point sets in a metric space. In *NeurIPS*, pages 5099–5108, 2017.
- [Savva *et al.*, 2016] Manolis Savva, Fisher Yu, Hao Su, M Aono, B Chen, D Cohen-Or, W Deng, Hang Su, Song Bai, Xiang Bai, et al. Shrec’16 track large-scale 3d shape retrieval from shapenet core55. In *Proceedings of the eurographics workshop on 3D object retrieval*, 2016.
- [Schroff *et al.*, 2015] Florian Schroff, Dmitry Kalenichenko, and James Philbin. Facenet: A unified embedding for face recognition and clustering. *CVPR*, pages 815–823, 2015.
- [Sfikas *et al.*, 2017] Konstantinos Sfikas, Theoharis Theoharis, and Ioannis Pratikakis. Exploiting the PANORAMA Representation for Convolutional Neural Network Classification and Retrieval. In Ioannis Pratikakis, Florent Dupont, and Maks Ovsjanikov, editors, *Eurographics Workshop on 3D Object Retrieval*, pages 1–7. The Eurographics Association, 2017.
- [Shen *et al.*, 2018] Yiru Shen, Chen Feng, Yaoqing Yang, and Dong Tian. Mining point cloud local structures by kernel correlation and graph pooling. In *CVPR*, 2018.
- [Shi *et al.*, 2015] Baoguang Shi, Song Bai, Zhichao Zhou, and Xiang Bai. Deeppano: Deep panoramic representation for 3-d shape recognition. *IEEE Signal Processing Letters*, 22(12):2339–2343, 2015.
- [Song *et al.*, 2016] Hyun Oh Song, Xiang Yu, Stefanie Jegelka, and Silvio Savarese. Deep metric learning via lifted structured feature embedding. In *CVPR*, pages 4004–4012, 2016.
- [Song *et al.*, 2017] Bai Song, Zhichao Zhou, Jingdong Wang, Bai Xiang, and Tian Qi. Ensemble diffusion for retrieval. In *ICCV*, pages 774–783, 2017.
- [Su *et al.*, 2015] Hang Su, Subhransu Maji, Evangelos Kalogerakis, and Erik Learned-Miller. Multi-view convolutional neural networks for 3d shape recognition. In *ICCV*, pages 945–953, 2015.
- [van der Maaten and Hinton, 2008] L.J.P. van der Maaten and G.E. Hinton. Visualizing high-dimensional data using t-sne. In *Journal of Machine Learning Research* 9(Nov), 2008.
- [Wen *et al.*, 2016] Yandong Wen, Kaipeng Zhang, Zhifeng Li, and Yu Qiao. A discriminative feature learning approach for deep face recognition. In *ECCV*, pages 499–515, 2016.
- [Wu *et al.*, 2015] Zhirong Wu, Shuran Song, Aditya Khosla, Fisher Yu, Linguang Zhang, Xiaoou Tang, and Jianxiong Xiao. 3d shapenets: A deep representation for volumetric shapes. In *CVPR*, pages 1912–1920, 2015.
- [Xie *et al.*, 2017] Jin Xie, Guoxian Dai, Fan Zhu, Edward K Wong, and Yi Fang. Deepshape: Deep-learned shape descriptor for 3d shape retrieval. *IEEE transactions on pattern analysis and machine intelligence*, pages 1335–1345, 2017.
- [Xu *et al.*, 2018] Cheng Xu, Biao Leng, Cheng Zhang, and Xiaochen Zhou. Emphasizing 3d properties in recurrent multi-view aggregation for 3d shape retrieval. In *AAAI*, pages 7428–7435, 2018.