

A Simple and Effective Framework for Pairwise Deep Metric Learning

Qi Qi¹, Yan Yan¹, Zixuan Wu², Xiaoyu Wang³, and Tianbao Yang^{✉1}

¹ Department of Computer Science, The University of Iowa, USA
{qi-qi, yan-yan-2, tianbao-yang}@uiowa.com

² Boston College, USA
wuzu@bc.edu

³ The Chinese University of Hong Kong (Shenzhen), China
fanghuaxue@gmail.edu

Abstract. Deep metric learning (DML) has received much attention in deep learning due to its wide applications in computer vision. Previous studies have focused on designing complicated losses and hard example mining methods, which are mostly heuristic and lack of theoretical understanding. In this paper, we cast DML as a simple pairwise binary classification problem that classifies a pair of examples as similar or dissimilar. It identifies the most critical issue in this problem—imbalanced data pairs. To tackle this issue, we propose a simple and effective framework to sample pairs in a batch of data for updating the model. The key to this framework is to define a robust loss for all pairs over a mini-batch of data, which is formulated by distributionally robust optimization. The flexibility in constructing the *uncertainty decision set* of the dual variable allows us to recover state-of-the-art complicated losses and also to induce novel variants. Empirical studies on several benchmark data sets demonstrate that our simple and effective method outperforms the state-of-the-art results.

Keywords: Deep Metric Learning, Distributed Robust Learning, Data Imbalance

1 Introduction

Metric Learning aims to learn a metric to measure the distance between examples that captures certain notion of human-defined similarity between examples. Deep metric learning (DML) has emerged as an effective approach for learning a metric by training a deep neural network. Simply speaking, a deep neural network can induce new feature embedding of examples and it is trained in such a way that the Euclidean distance between the induced feature embeddings of two similar examples shall be small and that between the induced feature embeddings of two dissimilar pairs shall be large. DML has been widely used in many tasks such as face recognition [2], image retrieval [1], and classification [15,8].

However, unlike training a deep neural network by minimizing the classification error, training a deep neural network for metric learning is notoriously

more difficult [16,23]. Many studies have attempted to address this challenge by focusing on several issues. The first issue is how to define a loss function over pairs of examples. A variety of loss functions have been proposed such as contrastive loss [4], binomial deviance loss [26], margin loss [25], lifted-structure (LS) loss [13], N-pair loss [22], triplet loss [19], multi-similarity (MS) loss [24].

The major difference between these pair-based losses lies at how the pairs interact with each other in a mini-batch. In simple pairwise loss such as binomial deviance loss, contrastive loss, and margin loss, pairs are regarded as independent of each other. In triplet loss, a positive pair only interacts with one negative pair. In N-pair loss, a positive pair interacts with all negative pairs. In LS loss and MS loss, a positive pair interacts with all positive pairs and all negative pairs. The trend is that the loss functions become increasingly complicated but are difficult to understand.

In parallel with the loss function, how to select informative pairs to construct the loss function has also received great attention. Traditional approaches that construct pairs or triplets over all examples before training suffer from prohibitive $O(N^2)$ or $O(N^3)$ sample complexity, where N is the total number of examples. To tackle this issue, constructing pairs within a mini-batch is widely used in practice. Although it helps mitigate the computational and storage burden, slow convergence and model degeneration with inferior performance still commonly exist when using all pairs in a mini-batch to update the model. To combat this issue, various *pair mining* methods have been proposed to complement the design of loss function, such as hard (semi-hard) mining for triplet loss [19], distance weighted sampling (DWS) for margin loss [25], MS sampling for MS loss [24]. These sampling methods usually keep all positive (similar) pairs and select roughly the same order of negative (dissimilar) pairs according to some criterion.

Regardless of these great efforts, existing studies either fail to explain the most fundamental problem in DML or fail to propose most effective approach towards addressing the fundamental challenge. It is evident that the loss functions become more and more complicated. But it is still unclear why these complicated losses are effective and how does the pair mining methods affect the overall loss within a mini-batch. In this paper, we propose a novel effective solution to DML and bring new insights from the perspective of learning theory that can guide the discovery of new methods. Our philosophy is simple: casting the problem of DML into a simple pairwise classification problem and focusing on addressing the most critical issue, i.e., the sheer imbalance between positive pairs and negative pairs. To this end, we employ simple pairwise loss functions (e.g., margin loss, binomial deviance loss) and propose a flexible distributionally robust optimization (DRO) framework for defining a robust loss over pairs within a mini-batch. The idea of DRO is to assign different weights to different pairs that are optimized by maximizing the weighted loss over an uncertainty set for the distributional variable. The model is updated by stochastic gradient descent with stochastic gradients computed based on the sampled pairs according to the founded optimal distributional variable.

The DRO framework allows us to (i) connect to advanced learning theories that already exhibit their power for imbalanced data, hence providing theoretical explanation for the proposed framework; (ii) to unify pair sampling and loss-based methods to provide a unified perspective for existing solutions; (iii) to induce simple and effective methods for DML, leading to state-of-the-art performance on several benchmark datasets. The contributions of our work are summarized as follows:

- We propose a general solution framework for DML, i.e., by defining a robust overall loss based on the DRO formulation and updating the model based on pairs sampled according to the optimized sampling probabilities. We provide theoretical justification of the proposed framework from the perspective of advanced learning theories.
- We show that the general DRO framework can recover existing methods based on complicated pair-based losses: LS loss and MS loss by specifying different uncertainty sets for the distributional variable in DRO. It verifies that our method is general and brings a unified perspective regarding pair sampling and complicated loss over all pairs within a batch.
- We also propose simple solutions under the general DRO framework for tackling DML. Experimental results show that our proposed variants of DRO framework outperform state-of-the-art methods on several benchmark datasets.

2 Related Work

Loss Design. The loss function is usually defined over the similarities or distances between the induced feature embeddings of pairs. There are simple pairwise losses that simply regard DML as binary classification problem using averaged loss over pairs, e.g., contrastive loss, binomial loss, margin loss. It is notable that the binomial loss proposed in [26] assigns asymmetric weights for positive and negative pairs, which can mitigate the issue of imbalance to certain degree. The principal of the newly designed complicated *pair-based* losses can be summarized as heuristically discovering specific kinds of relevant information between groups of pairs to boost the training. The key difference between these complicated losses lies at how to group the pairs. N-pair loss put one positive pair and all negative pairs together, Lifted-structure loss and MS-loss group all positive pairs together and all negative pairs together for each example. In contrast, our DRO framework employs simple pairwise loss but induce complicated overall loss in a systematic and interpretable way.

Pair Mining/Pair Weighting. [25] points out that pair mining plays an important role in distance metric learning. Different pair mining methods have been proposed, including semi-hard sampling for triplet loss, distance weighted sampling (DWS) for margin loss, MS mining for MS losses. These pair mining methods aim to select the hard positive and negative pairs for each anchor. For instance, [19] selects the hard negative pairs whose distance is smaller than that between the positive pairs in triplets, [21] selects the hardest positive pair

whose distance is smaller than that of the nearest negative pair in a batch, and MS mining [24] selects hard negative pairs whose distance is smaller than the largest distance between positive pairs and hard positive pairs whose distance is larger than the smallest distance between negative pairs at the same time. DWS method keeps all positive pairs but samples negative pairs according to their distance distribution within a batch. The proposed DRO framework induce a pair sampling method by using the optimal distributional variables that defines the robust loss over pairs within a mini-batch. As a result, the sampling probabilities induced by our DRO framework is automatically adaptive to the *pair-based losses*. There are other works that study the problem from the perspective of pair weighting instead of pair sampling. For example, [27] heuristically design exponential weights for the different pairs in a triplet, which is a special case of our DRO framework. Details are provided in the supplementary. However, since the quality of anchors varies very much, it may not be reasonable to sample the same number of pairs from all anchors.

Imbalance Data Classification. There are many studies in machine learning which have tackled the imbalanced issue. Commonly used tricks include over-sampling, under-sampling and cost-sensitive learning. However, these approaches do not take the differences between examples into account. Other effective approaches grounded on advanced learning theories include minimizing maximal losses [20], minimizing top-k losses [2] and minimizing variance-regularized losses [11]. However, these approaches are not efficient for deep learning with big data, which is a severe issue in DML. In contrast, the proposed DRO formulation is defined over a mini-batch of examples, which inherits the theoretical explanation from the literature and is much more efficient for DML. In addition, the induced loss by our DRO formulation include maximal loss, top-k loss and variance-regularized loss as special cases by specifying different uncertainty sets of the distributional variable.

3 DML as A DRO-Based Binary Classification Problem

In this section, we will first present a general framework for DML based on DRO with theoretical justification. We will then discuss three simple variants of the proposed framework and also show how the proposed framework recover existing complicated losses for DML.

Preliminaries. Let $\mathbf{x} \in \mathbb{R}^D$ denote an input data (e.g., image) and $f(\cdot; \theta) : \mathbb{R}^D \rightarrow \mathbb{R}^d$ denote the feature embedding function defined by a deep neural network parameterized by θ . The central task in DML is to update the model parameter θ by leveraging pairs of similar and dissimilar examples. Following most existing works, at each iteration we will sample a mini-batch of examples denoted by $\{\mathbf{x}_1, \dots, \mathbf{x}_B\}$. We can construct B^2 pairs between these examples⁴, and let y_{ij} denote the label of pairs, i.e., $y_{ij} = 1$ if the pair is similar (positive), and $y_{ij} = 0$ if the pair is dissimilar (negative). The label of pairs can be either defined by users

⁴ For simplicity, we consider all pairs including self-pair.

or derived from the class label of individual examples. Existing works of DML follow the same paradigm for learning the deep neural network i.e., a loss function $F(\theta)$ is first defined over the pairs within a mini-batch and the model parameter θ is updated by gradient-based methods. Various gradient-based methods can be used, including stochastic gradient descent (SGD), stochastic momentum methods and adaptive gradient methods (e.g. Adam). Taking SGD as an example, the model parameter θ can be updated by $\theta \leftarrow \theta - \eta \nabla F(\theta)$, where η denotes the learning rate. The focus here is to how to define the loss function $F(\theta)$ over all pairs within a mini-batch. As mentioned earlier, we will cast the problem as simple binary classification problem, i.e., classifying a pair into positive or negative. To this end, we use $l_{ij}(\theta) = l(f(\mathbf{x}_i; \theta), f(\mathbf{x}_j; \theta), y_{ij})$ denote the pairwise classification loss between \mathbf{x}_i and \mathbf{x}_j in the embedding space (e.g., margin loss [25], binomial loss [26]). A naive approach for DML is to use the averaged loss over all pairs, i.e., $F_{\text{avg}}(\theta) = \frac{1}{B^2} \sum_{i=1}^B \sum_{j=1}^B l_{ij}(\theta)$. However, this approach will suffer from the severe imbalanced issue, i.e., most pairs are negative pairs. The gradient of F_{avg} will be dominated by that of negative pairs.

3.1 General DRO-Based Framework

To address the imbalanced pair issue, we propose a general DRO formulation to compute a robust loss. The formulation of our DRO-based loss over all pairs within a mini-batch is given below:

$$F(\theta) = \max_{\mathbf{p} \in \mathcal{U}} \{g(\theta, \mathbf{p}) := \sum_{i=1}^B \sum_{j=1}^B p_{ij} l_{ij}(\theta)\}, \quad (1)$$

where $\mathbf{p} \in \mathbb{R}_+^{B^2}$ is a non-negative vector with each element p_{ij} representing a weight for an individual pair. $\mathcal{U} \subseteq \mathbb{R}^{B^2}$ denotes the decision set of \mathbf{p} , which encodes some prior knowledge about \mathbf{p} . In the literature of DRO [11], \mathbf{p} is interpreted as a probability vector such that $\sum_{ij} p_{ij} = 1$ called the distributional variable and \mathcal{U} denotes the uncertainty set that specifies how \mathbf{p} deviates from the uniform probabilities $(1/B^2, \dots, 1/B^2)$. In next subsection, we will propose simple variants of the above general framework by specifying different constraints or regularizations for \mathbf{p} . Below, we will provide some theoretical evidences to justify the above framework.

To theoretically justify the above loss, we connect (1) to exiting works in machine learning by considering three different uncertainty sets for \mathbf{p} . First, we can consider a simple constraint $\mathcal{U} = \Delta = \{p_{ij} \geq 0, \sum_{ij} p_{ij} = 1\}$. As a result, $F(\theta) = \max_{ij} l_{ij}(\theta)$ becomes the maximal loss over all pairs. [20] shows that minimizing maximum loss is robust to imbalanced data distributions and also derives better generalization error for imbalanced data with a rare class. However, the maximal loss is more sensitive to outliers [28]. To address this issue, top- K loss [2] and variance-regularized loss [11] are proposed, which can be induced by the above DRO framework. If we set $\mathcal{U} = \{\sum_{ij} p_{ij} = 1, 0 \leq p_{ij} \leq 1/K\}$, F will become the top- K loss $F(\theta) = \frac{1}{K} \sum_{i=1}^K l_{[i]}(\theta)$, where $l_{[i]}(\theta)$ denotes

the i -th largest loss over all pairs. If we set $\mathcal{U}_\phi = \{\mathbf{p} \in \Delta, D_\phi(\mathbf{p} \parallel \mathbf{1}/B^2) \leq \frac{\rho}{B^2}\}$, where $D_\phi(\mathbf{p} \parallel \mathbf{p}') = \int \phi(\frac{d\mathbf{p}}{d\mathbf{p}'})d\mathbf{p}'$ is the ϕ -divergence between two distributions \mathbf{p} and \mathbf{p}' with $\phi(t) = \frac{1}{2}(t-1)^2$, then the DRO-based loss becomes the variance-regularized loss under certain condition about the variance of the random loss, i.e., for a set of i.i.d random losses $\{\ell_1, \dots, \ell_n\} (n = B^2)$ we could have

$$\sup_{\mathbf{p} \in \mathcal{U}_\phi} \sum_{i=1}^n p_i \ell_i = \frac{1}{n} \sum_{i=1}^n \ell_i + \sqrt{\frac{2\rho \text{Var}_n(\ell)}{n}},$$

where $\text{Var}_n(\ell)$ denotes the empirical variance of ℓ_1, \dots, ℓ_n . We can see that the second term in R.H.S of the above equation involves the variance, which can play a role of regularization. The variance-regularized loss has been justified from advanced learning theory by [11], and its promising performance for imbalanced data has been observed as well.

Before ending this subsection, we will discuss how to update the model parameter θ based on the robust loss $F(\theta)$ defined by (1). A simple approach is to find an optimal distributional variable \mathbf{p}^* to (1) and then update θ according to the subgradient of weighted loss by $\partial_\theta g(\theta, \mathbf{p}^*) = \sum_{ij} p_{ij}^* \nabla l_{ij}(\theta)$, which is justified by the following lemma.

Lemma 1 *Assume that g is proper, lower-semicontinuous in θ and level-bounded in \mathbf{p} locally uniformly in θ . Then the subgradient $\partial F(\theta) \subset \bigcup_{\mathbf{p}^* \in P^*(\theta)} \partial_\theta g(\theta, \mathbf{p}^*)$, where $P^*(\theta)$ denotes the optimal solution set of the maximization problem in (1). Furthermore, when $l_{ij}(\theta)$ is smooth in θ and $P^*(\theta)$ is a singleton, i.e., $\mathbf{p}^* = \arg \max_{\mathbf{p}} g(\theta, \mathbf{p})$ is unique, we have $\partial F(\theta) = \partial_\theta g(\theta, \mathbf{p}^*)$.*

Remark 1. The above lemma can be proved by Theorem 10.13 in [17]. It shows that even if we may not directly compute $\partial F(\theta)$, our framework can at least obtain its superset $\partial_\theta g(\theta, \mathbf{p}^*)$. Particularly, if we have additional conditions, i.e., $l_{ij}(\theta)$ is smooth in θ and the optimal solution \mathbf{p}^* is unique (considering our regularized formulation below), it theoretically guarantees that our framework exactly computes $\partial F(\theta)$.

3.2 Proposed Three Variants of Our Framework

In order to contrast to other variants recovering existing complicated losses presented in next subsection, we introduce some notations and make some simplifications. For each example \mathbf{x}_i that serves as an anchor data, let $\mathcal{P}_i = \{j | y_{ij} = 1, j \in [B]\}$ and $\mathcal{N}_i = \{j | y_{ij} = 0, j \in [B]\}$ denote the index sets of its positive and negative pairs, respectively. Let $\mathcal{P} = \bigcup_{i=1}^B \mathcal{P}_i$ and $\mathcal{N} = \bigcup_{i=1}^B \mathcal{N}_i$. We denote the cardinality of a set by $P = |\mathcal{P}|$. For simplicity, we let $P_i = |\mathcal{P}_i|$, $N_i = |\mathcal{N}_i|$, $P = |\mathcal{P}|$ and $N = |\mathcal{N}|$. Since zero losses usually do not contribute to the computation of the subgradient for updating the model, we can simply eliminate those examples for consideration.

The first variant is to simply select the top- K pairs with K -largest losses, which is equivalent to the following DRO formulation:

$$\begin{aligned} \text{DRO-TopK:} \quad & \max_{\mathbf{p}} \sum_{i=1}^B \sum_{j \in \mathcal{P}_i \cup \mathcal{N}_i} p_{ij} l_{ij}(\theta) \\ & \text{s.t.} \sum_{i=1}^B \sum_{j \in \mathcal{P}_i \cup \mathcal{N}_i} p_{ij} = 1, 0 \leq p_{ij} \leq 1/K, \end{aligned}$$

where K is a hyper-parameter. The gradient of the robust loss can be simply computed by sorting the pairwise losses and computing the average of top- K losses.

The second variant is a variant of the variance-regularized loss. Instead of specifying the uncertainty set \mathcal{U}_ϕ , we use a regularization term for the ease of computation, which is defined by

$$\begin{aligned} \text{DRO-KL} \quad & \max_{\mathbf{p} \in \mathbb{R}_+^{P+N}} \sum_{i=1}^B \sum_{j \in \mathcal{P}_i \cup \mathcal{N}_i} p_{ij} l_{ij}(\theta) - \gamma D_{KL}(\mathbf{p} \parallel \frac{\mathbf{1}}{P+N}) \\ & \text{s.t.} \sum_{i=1}^B \sum_{j \in \mathcal{P}_i \cup \mathcal{N}_i} p_{ij} = 1, \end{aligned}$$

where $\gamma > 0$ is a hyper-parameter and D_{KL} denotes the KL divergence between two probabilities. The optimal solution to \mathbf{p} can be easily computed following [10]. It is notable that the optimal solution \mathbf{p}^* is not necessarily sparse. Hence, computing $\sum_{ij} p_{ij}^* \nabla l_{ij}(\theta)$ needs to compute the gradient of pairwise loss for all pairs, which could be prohibitive in practice when the mini-batch size is large. To alleviate this issue, we can simply sample a subset of pairs according to probabilities in \mathbf{p}^* and then compute the averaged gradient of these sampled pairs.

The third variant of our DRO framework is explicitly balancing the number of positive pairs and negative pairs by choosing top $K/2$ pairs for each class, which is denoted by DRO-TopK-PN:

$$\begin{aligned} \text{DRO-TopK-PN:} \quad & \max_{\mathbf{p} \in \{0,1\}^{P+N}} \sum_{i=1}^B \sum_{j \in \mathcal{P}_i \cup \mathcal{N}_i} p_{ij} l_{ij}(\theta) \\ & \text{s.t.} \sum_{i=1}^B \sum_{j \in \mathcal{P}_i} p_{ij} \leq \frac{K}{2}, \sum_{i=1}^B \sum_{j \in \mathcal{N}_i} p_{ij} \leq \frac{K}{2}. \end{aligned}$$

For implementation, we can simply select $K/2$ positive pairs with largest losses and $K/2$ negative pairs with largest loss respectively, and compute averaged gradient of the pairwise losses of the selected pairs for updating the model parameter.

3.3 Recovering the Method based on SOTA Pair-Based Loss

Next we show that proposed framework can recover the method based on SOTA complicated losses. With the induced feature vector $f(\mathbf{x}; \theta)$ normalized to have

unit norm, we define the similarity of two samples as $S_{ij} := \langle f(\mathbf{x}_i; \theta), f(\mathbf{x}_j; \theta) \rangle$, where $\langle \cdot, \cdot \rangle$ denotes dot product. Specifically, we consider two SOTA loss functions, LS and MS loss, which are defined below:

$$\mathcal{L}_{MS} = \frac{1}{n} \sum_{i=1}^n \left\{ \frac{1}{\alpha} \log \left[1 + \sum_{k \in \mathcal{P}_i} e^{-\alpha(S_{ik} - \lambda)} \right] + \frac{1}{\beta} \log \left[1 + \sum_{k \in \mathcal{N}_i} e^{\beta(S_{ik} - \lambda)} \right] \right\} \quad (2)$$

$$\mathcal{L}_{LS} = \sum_{i=1}^n \left[\log \sum_{k \in \mathcal{P}_i} e^{\lambda - S_{ik}} + \log \sum_{k \in \mathcal{N}_i} e^{S_{ik} - \lambda} \right]_+ \quad (3)$$

where α, β, λ are hyper-parameters of these losses.

The key to our argument is that the gradient computed based on these losses can be exactly computed according to our DRO framework by choosing appropriate constrained set \mathcal{U} and setting the pairwise loss as the margin loss. To this end, we first show the gradient based on the LS loss, which can be computed by [24]:

$$\frac{\partial \mathcal{L}(S)}{\partial \theta} = \frac{\partial \mathcal{L}(S)}{\partial S} \cdot \frac{\partial S}{\partial \theta} = \sum_{i=1}^B \sum_{j=1}^B \frac{\partial \mathcal{L}(S)}{\partial S_{ij}} \cdot \frac{\partial S_{ij}}{\partial \theta} \quad (4)$$

which can be alternatively written as

$$\frac{\partial \mathcal{L}(S)}{\partial \theta} = \sum_{i=1}^B \left(\sum_{j \in \mathcal{N}_i} w_{ij}^- \frac{\partial S_{ij}}{\partial \theta} - \sum_{j \in \mathcal{P}_i} w_{ij}^+ \frac{\partial S_{ij}}{\partial \theta} \right). \quad (5)$$

It can be shown that for LS loss, derivations are provided in the supplementary, we have

$$w_{ij}^+ = \frac{e^{\lambda - S_{ij}}}{\sum_{k \in \mathcal{P}_i} e^{\lambda - S_{ik}}} = \frac{1}{\sum_{k \in \mathcal{P}_i} e^{S_{ij} - S_{ik}}}, \quad w_{ij}^- = \frac{e^{S_{ij} - \lambda}}{\sum_{k \in \mathcal{N}_i} e^{S_{ik} - \lambda}} = \frac{1}{\sum_{k \in \mathcal{N}_i} e^{S_{ik} - S_{ij}}} \quad (6)$$

To recover the gradient of the LS loss under our DRO framework, we employ the pairwise margin loss for $l_{ij}(\theta)$, i.e., $l_{ij}(\theta) = [m + y_{ij}(\lambda - S_{ij})]_+$, where m and λ are two hyper-parameters and $[a]_+ = \max\{0, a\}$. Assume that the margin parameter m is sufficiently large such that $l_{ij}(\theta) > 0$ for all pairs. The key to deriving the same gradient of the LS loss under our framework is to group distributional variables in \mathbf{p} for the positive and negative pairs according to the anchor data. Let $\mathbf{p}_i^+ \in \mathbb{R}^{\mathcal{P}_i}$ and $\mathbf{p}_i^- \in \mathbb{R}^{\mathcal{N}_i}$ denote the corresponding vectors of positive and negative pairs for the anchor \mathbf{x}_i , respectively. Let us consider the following DRO formulation:

$$\begin{aligned} F(\theta) = & \max_{\mathbf{p} \in \mathbb{R}_+^{\mathcal{P}_+ \cup \mathcal{N}_+}} \sum_{i=1}^B \sum_{j \in \mathcal{P}_i \cup \mathcal{N}_i} p_{ij} l_{ij}(\theta) - \sum_{i=1}^B \left(\gamma_i^+ D_{KL}(\mathbf{p}_i^+ \| \frac{\mathbf{1}}{P_i}) + \gamma_i^- D_{KL}(\mathbf{p}_i^- \| \frac{\mathbf{1}}{N_i}) \right) \\ \text{s.t.} & \sum_{j \in \mathcal{P}_i} p_{ij} = 1, \sum_{k \in \mathcal{N}_i} p_{ik} = 1, \text{ for } i \in [B], \end{aligned} \quad (7)$$

where $\gamma_i^+ \geq 0$ and $\gamma_i^- \geq 0$ for $i \in [B]$ are hyper-parameters. we can easily derive the closed-form solution for \mathbf{p}^* , i.e., $p_{ij}^{+*} = \frac{1}{\sum_{k \in \mathcal{P}_i} e^{(S_{ij} - S_{ik})/\gamma_i^+}}$, and $p_{ij}^{-*} = \frac{1}{\sum_{k \in \mathcal{N}_i} e^{(S_{ik} - S_{ij})/\gamma_i^-}}$. Then computing the gradient of the robust loss with respect to θ by using the above optimal \mathbf{p}^* , we have:

$$\partial F(\theta) = \sum_{i=1}^B \left(\sum_{j \in \mathcal{N}_i} p_{ij}^{-*} \frac{\partial S_{ij}}{\partial \theta} - \sum_{j \in \mathcal{P}_i} p_{ij}^{+*} \frac{\partial S_{ij}}{\partial \theta} \right)$$

which exactly recover the gradient in (6) by setting $\gamma_i^+ = \gamma_i^- = 1$.

Finally, we can also recover the MS loss using another variants of our DRO framework that similar to equation (7). Besides losses, our DRO framework can also cover the heuristically designed exponential weights sampling strategy HAP2S_E in [27]. This verifies that our DRO framework is also able to provide hindsight for heuristic methods from advanced learning theories perspective. The detailed derivations of above two equivalence are provided in the supplementary.

4 Experiments

Our methods was implemented by pytorch and using BN-Inception network [6] pre-trained on ImageNet ILSVRC [18] to fairly compare with other works. The same as [24], a FC layer on the top of the model structure following the global pooling layer was added with randomly initialization for our task. Adam Optimizer with $1e - 5$ learning rate was used for all our experiments.

We verify our methods on the image retrieval task with three standard datasets, Cub-200-2011, Cars-196 and In-Shop. These three datasets are split according to the standard protocol. For Cub-200-2011, the first 100 classes with 5864 images are used for training, and the the other 100 classes with 5924 images are saved for testing. Cars-196 consists of 196 car models with 16,185 images. We use the first 98 classes with 8054 images for training and the remaining 98 classes with 8,131 images for testing. For In-Shop, 997 classes with 25882 images are used for training. The test set is further partitioned to a query set with 14218 images of 3985 classes, and a gallery set containing 3985 classes with 12612 images. Batches are constructed with the following rule: we first sample a certain number of classes and then randomly sample M instances for each class. The standard recall@ k evaluation metric is used in all our experiments, where $k = \{1, 2, 4, 8, 16, 32\}$ on Cub-200-2011 and Car-196, and $k = \{1, 10, 20, 30, 40, 50\}$ on In-Shop

We apply margin loss and binomial loss as our base loss for DRO framework. \mathcal{L}_M and \mathcal{L}_B denote margin loss and binomial loss ([26]) respectively. m denotes the margin in \mathcal{L}_M . λ is the threshold for both \mathcal{L}_M and \mathcal{L}_B . α and β are hyperparameters in \mathcal{L}_B .

4.1 Quantitative Results

In this experiment, we compare our DRO framework with other SOTA baselines on Cub-200-2011, Cars-196 and In-Shop, which includes [24,27,7,14,3,5,25,12]. Among them, mining-based methods are Clustring, HDC, Margin, Smart Mining and HDL. ABIER and ABE are ensemble methods. HAP2S_E and MS are sampling-based methods, which are highly related to our methods. For our DRO framework, we test all three variants which are proposed in section 3.2. We apply two loss functions, margin loss and binomial loss, respectively. Since DRO \mathbf{p} -sampling works on all pairs in a batch, the binomial variant may not directly apply to \mathbf{p} -sampling. Thus, it makes totally five variants of our DRO framework, denoted by DRO-TopK $_M$, DRO-TopK $_B$, DRO-TopK-PN $_M$, DRO-TopK-PN $_B$ and DRO-KL $_M$, where the subscript M and B represent the variants of our framework using margin loss and binomial loss, respectively. We set embedding space dimension $d = 1024$. The batchsize is set $B = 80$ on Cub-200-2011 and Cars-196, $B = 640$ on In-Shop. γ is tuned from the range = $\{0.1 : 0.2 : 0.9\}$ on all three datasets and K is tuned from $\{160, 200, 240, 280\}$ on Cub-200-2011 and Cars-196, and selected from $\{640, 960, 1280, 1600, 1920\}$ on In-Shop.

Table 1: Recall@ k on Cub-200-2011 and Cars-196

Recall@ k (%)	Cub-200-2011						Cars-196					
	1	2	4	8	16	32	1	2	4	8	16	32
Clustring([12])	48.2	61.4	71.8	81.9	-	-	58.1	70.6	80.3	87.8	-	-
HDC([12])	53.6	65.7	77.0	85.6	91.5	95.5	73.7	83.2	89.5	93.8	96.7	98.4
Margin([25])	63.6	74.4	83.1	90.0	94.2	-	79.6	86.5	91.9	95.1	97.3	-
Smart Mining([5])	49.8	62.3	74.1	83.3	-	-	64.7	76.2	84.2	90.2	-	-
HDL([3])	57.1	68.8	78.7	86.5	92.5	95.5	81.4	88.0	92.7	95.7	97.4	99.0
ABIER([14])	57.5	68.7	78.3	86.2	91.9	95.5	82.0	89.0	93.2	96.1	97.8	98.7
ABE([7])	60.6	71.5	79.8	87.4	-	-	85.2	90.5	94.0	96.1	-	-
HAP2S.E([27])	56.1	68.3	79.2	86.9	-	-	74.1	83.5	89.9	94.1	-	-
MS([24])	65.7	77.0	86.3	91.3	94.8	97.0	84.1	90.4	94.0	96.5	98.0	98.9
DRO-TopK $_M$ (Ours)	67.4	77.7	85.9	91.6	95.0	97.3	86.0	91.7	95.0	97.3	98.5	99.2
DRO-TopK $_B$ (Ours)	68.1	78.4	86.0	91.4	95.1	97.6	85.4	91.0	94.2	96.5	98.0	99.0
DRO-TopK-PN $_M$ (Ours)	67.3	77.6	85.7	91.2	95.0	97.7	86.1	91.7	95.1	97.1	98.4	99.1
DRO-TopK-PN $_B$ (Ours)	67.6	77.9	86.0	91.8	95.2	97.7	86.2	91.7	95.8	97.4	98.6	99.3
DRO-KL $_M$ (Ours)	67.7	78.0	86.1	91.8	95.6	97.8	86.4	91.9	95.4	97.5	98.7	99.3

Table 1 and 3 report the experiment results. We mark the best performer in bold in the corresponding evaluation measure on each column. For our framework, particularly, we mark those who outperform all other SOTA methods in bold. We can see that our five variants outperform other SOTA methods on recall@1 on all three datasets. Particularly on Cars-196, our five variants outperform other SOTA methods on all recall@ k measures. On Cub-200-2011, DRO-TopK $_B$ achieves a higher recall@1 (improving 2.4% in recall@1) than the best SOTA, MS. On Cars-196, DRO-KL $_M$ has the best performance, which improves 2.3% and 1.2% in recall@1 compared to the best non-ensemble SOTA, MS, and

the best ensemble SOTA, ABE. On In-Shop, DRO-TopK-PN_M improves 1.6% in recall@1 compared to the best results among SOTA, MS. The above results verify 1) the effectiveness of our DRO sampling methods and 2) the flexibility of our DRO framework to adopt different losses.

4.2 Ablation Study

Comparison with LS loss and MS loss In Section 3.3, we theoretically show that LS loss and MS loss can be viewed as special cases of our DRO framework. In this experiment, we aim to empirically demonstrate that our framework is general enough and recovers LS loss. Specifically, we would show 1) when $\gamma = 1$, our framework performs similarly to LS loss, as stated in Section 3.3, 2) our framework can be seen as a *generalized LS loss* by treating γ as a hyper-parameter, and 3) our *generalized LS loss* outperforms MS loss, even though the performance of the ordinary LS loss is inferior to that of MS loss.

We adopt the set up of embedding dimension and batchsize in the ablation study of [24], i.e., $d = 64$ and $B = 80$. Therefore, we report the existing results of MS and LS loss presented in [24] on Cars-196. For Cub-200-2011 and In-Shop, we implement MS and LS loss according to [24]. Following [24], we set $\alpha = 2, \beta = 50$ for MS loss. For our DRO framework, we apply grouping to \mathbf{p} by equation (7), and denote this variant of DRO framework as DRO-KL-G. We set $\gamma_i^+ = \gamma_i^- = \gamma = \{1, 0.1, 0.01, 0.001\}, i \in [B]$ for DRO-KL-G, $m = 0.2$ for the margin loss, and $\lambda = 0.5$ for all three losses (MS, LS and margin loss). As the pairs with zero loss will not contribute to the updates of model but affect the calculation of \mathbf{p} in DRO framework, we remove the pairs with zero loss to further promotes training.

Table 2 and 4 show experiment results on Cub-200-2011, Cars-196 and In-Shop, respectively. As can be seen, the performance of MS loss is better than LS loss on three datasets, particularly on Cars-196, which also verifies the results of ablation study in [24]. When $\gamma = 1$, our method performs similarly to LS loss, which verifies that our method recovers LS loss. Furthermore, when we treat γ as a hyper-parameter (especially $\gamma = 0.001$) and regard our framework as *generalized LS loss*, our method obtain improved performance compared to the ordinary LS loss. Lastly, even if MS loss exploits pseudo positive and negative pairs, our *generalized LS loss* outperforms MS loss.

Table 2: Recover of MS loss and LS loss on Cub-200-2011 and Cars-196

Recall@K(%)	Cub-200-2011						Cars-196					
	1	2	4	8	16	32	1	2	4	8	16	32
MS	55.6	67.7	77.4	86.3	92.1	95.8	73.2	81.5	87.6	92.6	-	-
LS	56.8	67.9	77.5	85.6	91.2	95.2	69.7	79.3	86.2	91.1	-	-
DRO-KL-G- $\gamma = 1$	56.4	68.3	78.9	86.3	91.7	95.8	70.5	79.8	86.6	91.6	94.9	97.1
DRO-KL-G- $\gamma = 0.1$	56.8	68.7	79.0	86.6	92.1	95.9	72.5	81.9	88.1	92.3	95.4	97.3
DRO-KL-G- $\gamma = 0.01$	57.0	69.4	79.9	87.0	92.3	95.9	73.1	82.2	88.8	93.4	96.2	98.0
DRO-KL-G- $\gamma = 0.001$	56.7	68.5	79.0	87.3	92.6	96.0	75.0	83.4	89.5	93.7	96.6	98.3

Table 3: Recall@ k on In-Shop

Recall@ K	1	10	20	30	40	50
FashionNet([9])	53.7	73.0	76.0	77.0	79.0	80.0
HDC([12])	62.1	84.9	89.0	91.2	92.3	93.1
HDL([3])	80.9	94.3	95.8	97.2	97.4	97.8
ABIER([14])	83.1	95.1	96.9	97.5	97.8	98.0
ABE([27])	87.3	96.7	97.9	98.2	98.5	98.7
MS([24])	89.7	97.9	98.5	98.8	99.1	99.2
DRO-TopK $_M$ (Ours)	91.0	98.1	98.7	99.0	99.1	99.2
DRO-TopK $_B$ (Ours)	90.7	97.7	98.4	98.8	99.0	99.1
DRO-TopK-PN $_M$ (Ours)	91.3	98.0	98.7	98.9	99.1	99.2
DRO-TopK-PN $_B$ (Ours)	91.1	98.1	98.6	98.8	99.0	99.2
DRO-KL $_M$ (Ours)	90.8	98.0	98.6	99.0	99.1	99.2



Table 4: Recover of MS loss and LS loss on In-Shop

Recall@ K (%)	1	10	20	30	40	50
MS	79.8	94.9	96.8	97.6	97.9	98.3
LS	82.6	94.1	95.6	96.4	96.9	97.4
DRO-KL-G- $\gamma = 1$	84.8	95.9	97.3	97.9	98.2	98.5
DRO-KL-G- $\gamma = 0.1$	85.1	96.1	97.5	98.0	98.3	98.5
DRO-KL-G- $\gamma = 0.01$	85.8	96.2	97.9	97.8	98.2	98.4
DRO-KL-G- $\gamma = 0.001$	85.7	96.1	97.4	97.9	98.2	98.5

Capacity to Handle Pair Imbalance. To show the effectiveness of DRO framework to handle pair imbalance, we conduct experiments under different positive and negative pair ratios and different dimensions of embedding space. By changing batchsize, the ratio of the numbers of positive and negative pairs also changes. For example, a larger batchsize leads to a smaller positive-negative (P-N) ratio.

In this experiment, we compare our DRO framework with two different SOTA sampling methods, i.e., semihard (SH) and DWS, in terms of sensitivity to P-N ratios under different embedding space dimensions. By setting different batchsizes $B \in \{80, 160, 320, 480, 640\}$, we have different P-N ratios $|\mathcal{P}| : |\mathcal{N}| \in \{0.053, 0.026, 0.012, 0.008, 0.006\}$. For all methods, we apply margin loss and set $M = 5$ for each class. The embedding space dimensions are $d = 1024$ and $d = 512$, respectively. SH mining is originally designed for triplet loss. Since there is no straightforward choice for the positive pair, we use λ as the upper bound to simulate the similarity of the positive pair in triplet loss. For DWS, we follow the parameter setting in the original paper [25]. We apply margin loss in the proposed three variants of our DRO framework, which are denoted by DRO-TopK $_M$, DRO-TopK-PN $_M$ and DRO-KL $_M$, respectively. We set $K = 2 * B$ for both DRO-TopK $_M$ and DRO-TopK-PN $_M$. We evaluate recall@1 of all methods, experimental results are reported in Figure 1 and Figure 2.

Experimental results illustrate that all three variants of our DRO framework has better or comparable performance than SH and DWS methods except for

Fig. 1: Recall vs Imbalance Ratio on Embedding Space $d = 1024$ Fig. 2: Recall vs Imbalance Ratio on Embedding Space $d = 512$

the largest P-N ratio 0.053 both on embedding space $d = 512$ and $d = 1024$. Among them, DRO-TopK-PN $_M$ constantly outperforms all other methods under embedding space dimension $d = 1024$ and achieve competitive results with DRO-TopK $_M$ under embedding space dimension $d = 512$. On the other hand, Figure 1 and 2 show that the DWS has similar performance when the P-N ratio is relatively large, and encounters a sharp decrease in recall@1 when the P-N ratio decreases on the embedding space dimension $d = 1024$. A smaller drop also exists on the embedding space dimension $d = 512$. The reason why DWS performs poorly when the P-N ratio is small may be that DWS aims to sample pairs uniformly in terms of distance [25], while our DRO framework and SH focus more on hard pairs.

Further, the fluctuations of recall@1 over three variants of our DRO framework are subtle when the feature embedding changes between 1024 and 512. For example, the recall@1 only changes, from 0.9046 to 0.9058 for DRO-TopK $_M$, from 0.9086 to 0.9073 for DRO-TopK-PN $_M$, from 0.9018 to 0.9021 for DRO-KL $_M$, when PN-Ratio is 0.006. However, DWS has a sharp decrease in recall@1 on the embedding space increasing from $d = 512$ to $d = 1024$. This implies that our DRO framework is not sensitive to embedding dimensions in comparison with DWS.

To sum up, above observations together verify that our methods are not sensitive to the embedding space dimensions in different batchsizes, and also outperform other SOTA mining methods in different embedding space dimensions.

Sensitivity of K in Top-K As we mentioned in section 1, selecting too many pairs within a batch will leads to poor performance of the model. On the other hand, when the number of selected pairs is too small, the model would be sensitive to outliers. In this experiment, we study the sensitivity of K in our DRO framework—how the performance of our DRO framework is affected by the value of K . We set the batchsize $B = 640$ and $M = 5$, which makes the number of positive pairs $|\mathcal{P}| = 1280$ and the number of negative pairs $|\mathcal{N}| = 198080$. We set K from the range $\{640, 960, 1280, 1600, 1920, 2560\}$ and evaluate recall@ k for models trained by different K . We choose the above range of K according to the number of pairs selected by DWS and SH in Section 4.2 (both selects 2560 pairs roughly).

Figure 3 illustrates how different values of K affect recall@ k on In-Shop. We can see that, DRO-TopK $_M$ performs best when $K = 1280$ and recall@ k is stable on the entire range of K . Our DRO framework is not sensitive to K when K is in a reasonably large range.

Runtime Comparison Next, we compare the running time of our proposed three variants of our DRO framework with different pair mining methods, MS and LS losses on In-shop. Our experiments conducted on eight GTX1080Ti GPU. The embedding dimension $d = 1024$, and results are compared under different batchsize $B = \{80, 160, 320, 480, 640\}$. The same as previous experiments, we




Fig. 3: The effects of K on recall@k on In-Shop




Fig. 4: Average running time of every iteration

set $K = 2 * B$ both for DRO-TopK_M and DRO-TopK-PN_M. $\gamma = 0.1$ for DRO-KL_M. SH is implemented according to the paper [19], [25]. DWS and MS are implemented based on the code provided by the author. LS loss is implemented following the code provided by [24].

Figure 4 reports the average running time of each iteration on 200 epochs. We can see that all of three proposed variants of DRO framework run faster than other *anchor-based* mining methods and losses. For all of our three variants, pairs are selected directly from all the pairs, while additional cost is required to select pairs anchor by anchor in other methods. LS loss is slower than MS loss, because MS mining is applied to MS loss, which would reduce the number of pairs for computing subgradients when updating the model. For DWS, the distance distribution of negative pairs is only calculated once for each iteration. It thus only needs to select pairs according to the pre-computed distribution for each anchor. In contrast, SH requires to compare negative pairs with the lower and upper bound of an interval at each iteration for each anchor, which increases the computational burden. It can be the reason why SH is slower than DWS.

5 Conclusion

In this paper, we cast DML as a simple pairwise binary classification problem and formulate it as a DRO framework. Compared to existing pairwise DML methods that leverage all pairs in a batch or employ heuristic approaches to sample pairs, our DRO framework constructs a robust loss to sample informative pairs, which also comes with theoretical justification from the perspective of learning theory. Our framework is general since it can include many novel designs in its uncertainty decision set. Its flexibility allows us to recover the state-of-the-art loss functions and exponential sample weighting strategy. Experiments show that our framework outperforms the state-of-the-art DML methods on benchmark datasets. We also empirically demonstrate that our framework is efficient, general and flexible in ablation study.

References

1. Chen, B., Deng, W.: Hybrid-attention based decoupled metric learning for zero-shot image retrieval. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 2750–2759 (2019)
2. Fan, Y., Lyu, S., Ying, Y., Hu, B.: Learning with average top-k loss. In: Advances in Neural Information Processing Systems. pp. 497–505 (2017)
3. Ge, W.: Deep metric learning with hierarchical triplet loss. In: Proceedings of the European Conference on Computer Vision (ECCV). pp. 269–285 (2018)
4. Hadsell, R., Chopra, S., LeCun, Y.: Dimensionality reduction by learning an invariant mapping. In: 2006 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'06). vol. 2, pp. 1735–1742. IEEE (2006)
5. Harwood, B., Kumar, B., Carneiro, G., Reid, I., Drummond, T., et al.: Smart mining for deep metric learning. In: Proceedings of the IEEE International Conference on Computer Vision. pp. 2821–2829 (2017)
6. Ioffe, S., Szegedy, C.: Batch normalization: Accelerating deep network training by reducing internal covariate shift. arXiv preprint arXiv:1502.03167 (2015)
7. Kim, W., Goyal, B., Chawla, K., Lee, J., Kwon, K.: Attention-based ensemble for deep metric learning. In: Proceedings of the European Conference on Computer Vision (ECCV). pp. 736–751 (2018)
8. Li, X., Yu, L., Fu, C.W., Fang, M., Heng, P.A.: Revisiting metric learning for few-shot image classification. arXiv preprint arXiv:1907.03123 (2019)
9. Liu, Z., Luo, P., Qiu, S., Wang, X., Tang, X.: Deepfashion: Powering robust clothes recognition and retrieval with rich annotations. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 1096–1104 (2016)
10. Namkoong, H., Duchi, J.C.: Stochastic gradient methods for distributionally robust optimization with f-divergences. In: Advances in Neural Information Processing Systems. pp. 2208–2216 (2016)
11. Namkoong, H., Duchi, J.C.: Variance-based regularization with convex objectives. In: Advances in Neural Information Processing Systems (NIPS). pp. 2975–2984 (2017)
12. Oh Song, H., Jegelka, S., Rathod, V., Murphy, K.: Deep metric learning via facility location. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 5382–5390 (2017)
13. Oh Song, H., Xiang, Y., Jegelka, S., Savarese, S.: Deep metric learning via lifted structured feature embedding. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 4004–4012 (2016)
14. Opitz, M., Waltner, G., Possegger, H., Bischof, H.: Deep metric learning with hier: Boosting independent embeddings robustly. IEEE transactions on pattern analysis and machine intelligence (2018)
15. Qian, Q., Jin, R., Zhu, S., Lin, Y.: Fine-grained visual categorization via multi-stage metric learning. In: The IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (June 2015)
16. Qian, Q., Tang, J., Li, H., Zhu, S., Jin, R.: Large-scale distance metric learning with uncertainty. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 8542–8550 (2018)
17. Rockafellar, R.T., Wets, R.J.B.: Variational analysis, vol. 317. Springer Science & Business Media (2009)
18. Russakovsky, O., Deng, J., Su, H., Krause, J., Satheesh, S., Ma, S., Huang, Z., Karpathy, A., Khosla, A., Bernstein, M., et al.: Imagenet large scale visual recognition challenge. International journal of computer vision **115**(3), 211–252 (2015)

19. Schroff, F., Kalenichenko, D., Philbin, J.: Facenet: A unified embedding for face recognition and clustering. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 815–823 (2015)
20. Shalev-Shwartz, S., Wexler, Y.: Minimizing the maximal loss: How and why. In: ICML. pp. 793–801 (2016)
21. Shi, H., Yang, Y., Zhu, X., Liao, S., Lei, Z., Zheng, W., Li, S.Z.: Embedding deep metric for person re-identification: A study against large variations. In: European conference on computer vision. pp. 732–748. Springer (2016)
22. Sohn, K.: Improved deep metric learning with multi-class n-pair loss objective. In: Advances in Neural Information Processing Systems. pp. 1857–1865 (2016)
23. Wang, J., Zhou, F., Wen, S., Liu, X., Lin, Y.: Deep metric learning with angular loss. In: Proceedings of the IEEE International Conference on Computer Vision. pp. 2593–2601 (2017)
24. Wang, X., Han, X., Huang, W., Dong, D., Scott, M.R.: Multi-similarity loss with general pair weighting for deep metric learning. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 5022–5030 (2019)
25. Wu, C.Y., Manmatha, R., Smola, A.J., Krahenbuhl, P.: Sampling matters in deep embedding learning. In: Proceedings of the IEEE International Conference on Computer Vision. pp. 2840–2848 (2017)
26. Yi, D., Lei, Z., Li, S.: Deep metric learning for practical person re-identification (2014). ArXiv e-prints
27. Yu, R., Dou, Z., Bai, S., Zhang, Z., Xu, Y., Bai, X.: Hard-aware point-to-set deep metric for person re-identification. In: Proceedings of the European Conference on Computer Vision (ECCV). pp. 188–204 (2018)
28. Zhu, D., Li, Z., Wang, X., Gong, B., Yang, T.: A robust zero-sum game framework for pool-based active learning. In: The 22nd International Conference on Artificial Intelligence and Statistics. pp. 517–526 (2019)