

Denoising Multi-Similarity Formulation: A Self-Paced Curriculum-Driven Approach for Robust Metric Learning

Chenkang Zhang¹, Lei Luo², Bin Gu^{1,3*}

¹ School of Computer and Software, Nanjing University of Information Science and Technology, P.R.China

² School of Computer Science and Engineering, Nanjing University of Science and Technology, P.R.China

³ MBZUAI, United Arab Emirates

20201221058@nuist.edu.cn, luoleipitt@gmail.com, jsgubin@gmail.com

Abstract

Deep Metric Learning (DML) is a group of techniques that aim to measure the similarity between objects through the neural network. Although the number of DML methods has rapidly increased in recent years, most previous studies cannot effectively handle noisy data, which commonly exists in practical applications and often leads to serious performance deterioration. To overcome this limitation, in this paper, we build a connection between noisy samples and hard samples in the framework of self-paced learning, and propose a **Balanced Self-Paced Metric Learning (BSPML)** algorithm with a denoising multi-similarity formulation, where noisy samples are treated as extremely hard samples and adaptively excluded from the model training by sample weighting. Especially, due to the pairwise relationship and a new balance regularization term, the sub-problem *w.r.t.* sample weights is a nonconvex quadratic function. To efficiently solve this nonconvex quadratic problem, we propose a doubly stochastic projection coordinate gradient algorithm. Importantly, we theoretically prove the convergence not only for the doubly stochastic projection coordinate gradient algorithm, but also for our BSPML algorithm. Experimental results on several standard data sets demonstrate that our BSPML algorithm has better generalization ability and robustness than the state-of-the-art robust DML approaches.

Introduction

DML aims to learn an embedding space in which similar samples are pulled closer while dissimilar samples are encouraged to stay away from each other (Xing et al. 2002; Wang, Peng, and Lin 2021; Wang et al. 2022). Compared with traditional metric learning methods, which may not capture the nonlinear nature of data, DML utilizes the neural network to obtain representative and discriminative feature embeddings. Thus, DML has attracted increasing attention and been applied to various tasks, including visual tracking (Leal-Taixé, Canton-Ferrer, and Schindler 2016; Tao, Gavves, and Smeulders 2016), face recognition (Wen et al. 2016), image retrieval (Wohllhart and Lepetit 2015; He et al. 2018; Grabner, Roth, and Lepetit 2018), person re-identification (Hermans, Beyrer, and Leibe 2017; Yu et al.

2018a) and zero-shot learning (Zhang and Saligrama 2016; Yelamarthi et al. 2018; Bucher, Herbin, and Jurie 2016).

The superior performance of machine learning greatly depends on a large number of labeled data sets (Shi et al. 2021), and the performance of DML is no exception. However, manually generating clean data set would involve domain experts evaluating the quality of collected data and thus is very expensive and time-consuming (Frénay and Verleysen 2013). To address this issue, some researchers utilize the online key search engine method (Yu et al. 2018b) and the crowdsourcing method (Li et al. 2017) to gain required data sets at a low cost, but it is possible to introduce **noisy samples** that represent mislabeled ones (Wu et al. 2021; Yao et al. 2021; Zhai et al. 2020). As far as we know, most existing DML approaches are sensitive to noisy samples since they directly utilize sample labels to learn the similarity information between samples.

Constructing robust DML models against noisy samples is a challenging task, and some researchers have paid attention to this problem. Specifically, Wang *et al.* proposed a novel objective using the ℓ_1 -norm distance (Wang, Nie, and Huang 2014), and Al-Obaidi *et al.* utilized the rescaled hinge loss, which was a general form of the common hinge loss, to formulate the DML problem (Al-Obaidi, Zabihzadeh, and Hajiabadi 2020). Moreover, Kim *et al.* proposed the Proxy-Anchor loss that was robust against noisy labels and outliers because of the use of proxies (Kim et al. 2020). Different from the above technical routes, Yuan *et al.* proposed a robust distance metric based on the Signal-to-Noise Ratio (SNR) (Yuan et al. 2019).

Different from mislabeled noisy samples, **hard samples** are correctly labeled ones that are difficult to distinguish by the model (Nguyen et al. 2019; Zhu et al. 2021). In the image data set, hard samples can be anything from cats that look like dogs to images with slightly blurred resolution. Evidently, noisy samples and hard samples have different roles. Noisy samples are harmful because they would mislead the training direction. However, hard samples could force the model to learn more representative features, and thus properly training these hard samples could improve the model generalization ability (Zhu et al. 2021; Chen et al. 2020). Although noisy samples and hard samples are different, we would build the connection between them in the perspective of Self-Paced Learning (SPL) (Kumar, Packer, and Koller

*Corresponding Author

Copyright © 2023, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

2010), and further utilize SPL to filter out noisy samples from the training set.

Inspired by human cognitive mechanisms, (Kumar, Packer, and Koller 2010) proposed a novel learning strategy called Self-Paced Learning (SPL), which utilizes the loss value to describe the difficulty degree of the sample. Also, SPL starts learning from easy samples and gradually incorporates more difficult samples (*i.e.*, hard samples). Particularly, we could consider noisy samples as extremely hard samples because they usually own larger loss values than conventional hard samples (please refer to Fig. 1.(a)). Thus, SPL can assign smaller weights to noisy samples under the guidance of loss values, and then noisy samples would have little influence on the predicted model. Theoretically, (Meng, Zhao, and Jiang 2017) has proved that such a re-weighting learning process is equivalent to minimizing a latent noise-robust loss that would weaken the contribution of noisy samples. Experimentally, it has been declared that SPL is an effective method to improve the model robustness against noisy data (Zhang et al. 2020; Yin, Liu, and Sun 2021; Ren et al. 2020; Gu et al. 2021).

In this paper, we propose a Balanced Self-Paced Metric Learning (BSPML) algorithm with a novel denoising multi-similarity formulation. Benefiting from the mechanism of SPL, our BSPML algorithm could exclude noisy samples and emphasize the importance of clean samples. Because DML problems focus on a large number of classes, we have to face the challenge of the unbalanced average sample weights among classes. To overcome this problem, we introduce a new balance regularization term to punish the absolute difference between the average sample weights of different classes. Following the traditional SPL practice, our BSPML algorithm utilizes the alternative optimization strategy based on two key sub-problems *w.r.t.* model parameters and sample weights respectively. To efficiently solve the nonconvex sub-problem *w.r.t.* sample weights, we propose a doubly stochastic projection coordinate gradient algorithm. Theoretically, we prove the convergence of our proposed algorithms under mild assumptions. Experimental results show the advantages of our BSPML algorithm in generalization ability and robustness.

Preliminaries

In this section, we give a brief review of self-paced learning and multi-similarity loss.

Self-Paced Learning

Suppose $\{(\mathbf{x}_i, y_i)_{i=1}^N\}$ is a set of N labeled samples, where $y_i \in [C]$ is the corresponding label of the sample $\mathbf{x}_i \in \mathbb{R}^M$. SPL utilizes a sample weight vector $\mathbf{w} \in [0, 1]^N$ to indicate whether or not each training sample should be included in the current training. The classic SPL model is formed as:

$$\min_{\theta, \mathbf{w} \in [0, 1]^N} \sum_{i=1}^N w_i L_{\text{point}}(\mathbf{x}_i, y_i, \theta) - \lambda \sum_{i=1}^N w_i, \quad (1)$$

where θ means the model parameters, L_{point} represents one point-wise loss, *e.g.*, hinge loss, λ is the age parameter which controls the learning pace in SPL.

SPL methods utilize the alternative optimization strategy based on two key sub-problems *w.r.t.* model parameters θ and sample weights \mathbf{w} respectively. Especially, for the classic SPL formulation, the closed-form solution of the sub-problem *w.r.t.* \mathbf{w} can be obtained easily:

$$\begin{cases} w_i = 1, & L_{\text{point}}(\mathbf{x}_i, y_i, \theta) \leq \lambda \\ w_i = 0, & L_{\text{point}}(\mathbf{x}_i, y_i, \theta) > \lambda \end{cases} \quad (2)$$

The above solution implies that if the loss value of a sample is less than λ , the sample will be assigned the weight of value 1 and thus be selected to join the training. Otherwise, the sample will be excluded from the training process by reducing its weight to 0. In this case, with increased λ , hard samples with larger loss values will join the training. Thus, the age parameter λ controls the learning pace in SPL.

To better understand the mechanism of SPL and the roles of hard samples and noisy samples, we provide an example of the self-paced classification problem in Fig. 1. According to Eq. (2), when we set λ to a small value λ_1 , only easy samples are selected to join the training in Fig. 1.(a). As λ increases to λ_2 , hard samples with larger loss are also taken into consideration as shown in Fig. 1.(b). These hard samples could force the model to learn more representative features, and thus properly training these hard samples could improve the generalization ability. However, when λ increases to λ_3 , noisy samples with extremely excessive losses are considered and lead to the degeneration of the generalization ability as shown in Fig. 1.(c).

Multi-Similarity Loss

Let $f_{\theta}(\mathbf{x}_i)$ be an embedding vector of the sample \mathbf{x}_i where $f_{\theta} : \mathbb{R}^M \rightarrow \mathbb{R}^m$ means the model with parameters θ . For the training stability, embedding vectors have been normalized. Formally, we define the similarity between two samples as $S_{\mathbf{x}_i, \mathbf{x}_j} = \langle f_{\theta}(\mathbf{x}_i), f_{\theta}(\mathbf{x}_j) \rangle$, where $\langle \cdot, \cdot \rangle$ denotes dot product. A high similarity between two samples implies that these two samples are close to each other in the embedding space. Besides, in DML, a positive pair means a pair of samples from the same class and a negative pair represents a pair of samples from different classes.

Similar to contrastive loss (Bromley et al. 1993) and triplet loss (Hermans, Beyer, and Leibe 2017), Multi-Similarity (MS) loss (Wang et al. 2019) is one pair-based DML loss, which pays attention to the distance information between pairs of samples. Specifically, the MS objective is formulated as:

$$L_{\text{MS}}(\theta) = \sum_{c=1}^C \frac{1}{N^c} \sum_{a=1}^{N^c} (\xi_{\theta}^{+}(\mathbf{x}_a^c) + \xi_{\theta}^{-}(\mathbf{x}_a^c)), \quad (3)$$

$$\begin{aligned} \text{where } \xi_{\theta}^{+}(\mathbf{x}_a^c) &= \frac{1}{\alpha} \log \left[1 + \sum_{\substack{p \in [N^c] \\ p \neq a}} e^{-\alpha(S_{\mathbf{x}_a^c, \mathbf{x}_p^c} - \rho)} \right], \\ \xi_{\theta}^{-}(\mathbf{x}_a^c) &= \frac{1}{\beta} \log \left[1 + \sum_{\substack{k \in [C] \\ k \neq c}} \sum_{n=1}^{N^k} e^{\beta(S_{\mathbf{x}_a^c, \mathbf{x}_n^k} - \rho)} \right]. \end{aligned} \quad \text{Here, } \alpha, \beta$$

and ρ are fixed hyper-parameters, \mathbf{x}_a^c means the a -th sample in class $c \in [C]$, N^c is the number of samples in class c




Figure 1: Self-paced classification problem with hard and noisy samples. (λ is the age parameter and $\lambda_1 < \lambda_2 < \lambda_3$.)

and $N = \sum_{c=1}^C N^c$. Moreover, an informative pair mining method is also proposed. For an anchor sample \mathbf{x}_a^c , if a negative pair $\{\mathbf{x}_a^c, \mathbf{x}_n^k\}$, $k \in [C], k \neq c, n \in [N^k]$ satisfies the following condition:

$$S_{\mathbf{x}_a^c, \mathbf{x}_n^k} > \min_{\substack{p \in [N^c] \\ p \neq a}} S_{\mathbf{x}_a^c, \mathbf{x}_p^c} - \epsilon, \quad (4)$$

where ϵ is a given margin, the negative pair can be considered as an informative negative pair. And an informative positive pair $\{\mathbf{x}_a^c, \mathbf{x}_p^c\}$, $p \neq a$ should satisfy:

$$S_{\mathbf{x}_a^c, \mathbf{x}_p^c} < \max_{\substack{k \in [C], k \neq c \\ n \in [N^k]}} S_{\mathbf{x}_a^c, \mathbf{x}_n^k} + \epsilon. \quad (5)$$

Proposed Algorithm

In this section, we provide our balanced self-paced metric learning algorithm followed by its objective function.

Objective Function

Inspired by the classic SPL formulation, we propose the new denoising multi-similarity formulation:

$$\begin{aligned} & \mathcal{L}(\theta, \mathbf{w}; \lambda) \\ &= \sum_{c=1}^C \sum_{a=1}^{N^c} w_a^c \left(\frac{\zeta(c) - w_a^c}{N^c - 1} \xi_{\theta}^+(\mathbf{x}_a^c) + \frac{\sum_{k \in [C], k \neq c} \zeta(k)}{C - 1} \frac{\zeta(k)}{N^k} \xi_{\theta}^-(\mathbf{x}_a^c) \right) \\ & - \lambda \sum_{c=1}^C \frac{\zeta(c)}{N^c} + \frac{\mu}{C - 1} \sum_{c=1}^C \sum_{k=c+1}^C \left(\frac{\zeta(c)}{N^c} - \frac{\zeta(k)}{N^k} \right)^2 \\ & \text{s.t. } w_a^c \in [0, 1], \forall c \in [C], a \in [N^c] \end{aligned} \quad (6)$$

where λ is the age parameter that controls the learning pace in SPL and $\zeta(c) = \sum_{i=1}^{N^c} w_i^c$ means the sum of sample weights in class c . Importantly, the last term in Eq. (6) is our new balance regularization term. Specifically, DML problems usually focus on a large number of classes. Thus, we have to face the challenge of the unbalanced average sample weights among classes, e.g., all samples in some classes are assigned the weights of value 1, while all samples in other classes are assigned the weights of value 0. To overcome this challenge, we propose our balance regularization term, which punishes the absolute difference between the average sample weights of different classes.

Balanced Self-Paced Metric Learning

Following the traditional SPL practice, our BSPML utilizes the alternative optimization strategy based on two key sub-problems *w.r.t.* θ and \mathbf{w} as shown in Algorithm 1, where c is the multiplier of λ and λ_{∞} is the maximum value of λ . In the following, we present Algorithms 2 and 3 used for solving the two key sub-problems *w.r.t.* θ and \mathbf{w} respectively.

Mini-Batch Gradient Algorithm for Weighted Multi-Similarity Loss When we fix \mathbf{w} to optimize θ , our objective function degenerates into a weighted multi-similarity formulation that can be optimized by the common practice in DML.

Stochastic sampling: We first stochastically sample P classes and then stochastically sample K samples per class. *Informative pair mining:* We select informative pairs from these PK samples according to Eqs. (4) and (5). In fact, when we use the original MS formulation, we inevitably utilize noisy pairs that would produce noisy gradients as Fig. 2.(a) shows. However, in our BSPML algorithm, we exclude noisy samples by reducing these weights under the guidance of the loss value and thus obtain an embedding space with satisfactory performance, where there exists enough distance between samples from different classes as shown in Fig. 2.(d).

Weighted informative batch loss: For an anchor sample \mathbf{x}_i , we define $\hat{P}_{\mathbf{x}_i}$ as the index set of informative positive pairs and $\hat{N}_{\mathbf{x}_i}$ as the index set of informative negative pairs. Then, the weighted informative batch loss is formed as:

$$\begin{aligned} & L_{\text{WBL}}(\theta) \\ &= \frac{1}{PK} \sum_{i=1}^{PK} w_i \left\{ \frac{\sum_{p \in \hat{P}_{\mathbf{x}_i}} w_p}{|\hat{P}_{\mathbf{x}_i}|^{\alpha}} \log \left[1 + \sum_{p \in \hat{P}_{\mathbf{x}_i}} e^{-\alpha(S_{\mathbf{x}_i, \mathbf{x}_p} - \rho)} \right] \right. \\ & \left. + \frac{\sum_{n \in \hat{N}_{\mathbf{x}_i}} w_n}{|\hat{N}_{\mathbf{x}_i}|^{\beta}} \log \left[1 + \sum_{n \in \hat{N}_{\mathbf{x}_i}} e^{\beta(S_{\mathbf{x}_i, \mathbf{x}_n} - \rho)} \right] \right\}. \end{aligned} \quad (7)$$

Finally, we summarize the above procedures in Algorithm 2, where the gradient-based method means the optimization method utilizing the gradient information such as stochastic

Algorithm 1: Balanced Self-Paced Metric Learning

Input: Maximum iteration T , initial age parameter λ^0 , multiplier c and maximum age parameter λ_∞ .

Output: The optimal θ .

- 1: Initialize $\mathbf{w}^0 = \mathbf{1}_N$.
 - 2: **for** $t = 1, \dots, T$ **do**
 - 3: Fix \mathbf{w} to optimize θ through Algorithm 2.
 - 4: Fix θ to optimize \mathbf{w} through Algorithm 3.
 - 5: $\lambda^t = \min(c\lambda^{t-1}, \lambda_\infty)$.
 - 6: **end for**
-

Algorithm 2: Mini-Batch Gradient Algorithm for Weighted Multi-Similarity Loss

Input: Maximum iteration T , initial model parameter matrix θ^0 , class batch size P and sample batch size K .

Output: The optimal θ .

- 1: **for** $t = 1, \dots, T$ **do**
 - 2: Sample P classes stochastically and then sample K samples per class stochastically.
 - 3: Select informative pair index sets \hat{P} and \hat{N} according to Eqs. (4) and (5).
 - 4: Update θ using gradient-based method with the loss L_{WBL} , *i.e.*, Eq. (7).
 - 5: **end for**
-

Algorithm 3: Doubly Stochastic Projection Coordinate Gradient Method

Input: Maximum iteration T , initial sample weight vector \mathbf{w}^0 , learning rate γ , class batch size P and weight batch size K .

Output: The optimal \mathbf{w} .

- 1: **for** $t = 1, \dots, T$ **do**
 - 2: Sample one weight $w_a^c, c \in [C], a \in [N^c]$ stochastically.
 - 3: Sample K weights different from w_a^c in class c stochastically, sample P classes different from class c and K weights per class stochastically.
 - 4: Calculate the stochastic gradient $G(w_a^c)$.
 - 5: Update $\mathbf{w}^t = \mathcal{P}_{[0,1]^N}(\mathbf{w}^{t-1} - \gamma^t G(w_a^c) \mathbf{e}_a^c)$.
 - 6: **end for**
-

gradient descent (SGD) (Ketkar 2017) and adaptive moment estimation (Adam) (Kingma and Ba 2014).

Doubly Stochastic Projection Coordinate Gradient Method

When we fix θ to optimize \mathbf{w} , the sub-problem *w.r.t.* \mathbf{w} is a more complex quadratic problem compared with the one in existing SPL problems, *e.g.*, Eq. (1). To solve this complex sub-problem efficiently, we propose a doubly stochastic projection coordinate gradient method.

Stochastic sampling: Specifically, we first sample an anchor weight $w_a^c, c \in [C], a \in [N^c]$ stochastically. Then, for the anchor weight w_a^c , we select K weights different from w_a^c in class c , and select P classes different from c and select K weights per class.

Stochastic gradients: Under the circumstance, we show the

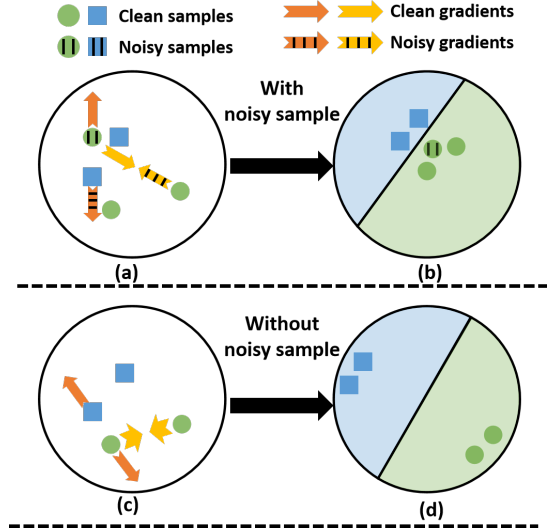


Figure 2: The difference between the two-class DML tasks with and without the noisy sample. (A pair of pointers of the same color in one sub-figure represents a pair of gradients.)

stochastic gradient $G(w_a^c)$ regarding to a particular coordinate w_a^c as follows:

$$\begin{aligned}
 G(w_a^c) &= \frac{1}{N^c} (G_p(w_a^c) + G_n(w_a^c) + G_b(w_a^c) - \lambda), \\
 G_p(w_a^c) &= \frac{1}{K} \sum_{p=1}^K w_p^c (\xi_\theta^+(\mathbf{x}_p^c) + \xi_\theta^+(\mathbf{x}_a^c)), \\
 G_n(w_a^c) &= \frac{1}{P} \sum_{k=1}^P \frac{1}{K} \sum_{n=1}^K w_n^k (\xi_\theta^-(\mathbf{x}_n^k) + \xi_\theta^-(\mathbf{x}_a^c)), \\
 G_b(w_a^c) &= 2\mu \left(\frac{\sum_{i=1}^{N^c} w_i^c}{N^c} - \frac{1}{C-1} \sum_{\substack{k \in [C] \\ k \neq c}} \frac{\sum_{i=1}^{N^k} w_i^k}{N^k} \right).
 \end{aligned} \tag{8}$$

The gradient G in Eq. (8) is composed of four terms. Both G_p and G_n are related to the similarity between samples, and they imply that if one sample owns the high similarity to samples from different classes and the low similarity to samples from the same class, our algorithm would exclude this sample from the training by gradually reducing its weight. Moreover, the G_b is generated by our proposed balance regularization term. As we expected, a sample will be assigned a larger weight if the average sample weight $\frac{\sum_{i=1}^{N^c} w_i^c}{N^c}$ of its class is low, or the average sample weight $\frac{1}{C-1} \sum_{k \in [C], k \neq c} \frac{\sum_{i=1}^{N^k} w_i^k}{N^k}$ of all other classes is high. Finally, the last term of Eq. (8) is the age parameter λ , which controls the learning pace in SPL. With increased λ , our algorithm tends to assign samples larger weights.

The doubly stochastic projection coordinate gradient method is summarized in Algorithm 3, where γ means the learning rate, \mathcal{P}_S is the projection operation to the set S and \mathbf{e}_a^c is one unit vector where the coordinate with the same index as w_a^c holds the value of 1.

Theoretical Analysis

In this section, we prove the convergence of Algorithms 3 and 1. All the proof details are available in Appendix.

Convergence of Algorithm 3: For Algorithm 3, let $w_{j(t)}$ mean the coordinate selected in t -iteration and $\mathcal{B}(t)$ mean the mini-batch sampled in t -iteration. Then, we define

$$\nabla_{j(t)} \mathcal{L}(\mathbf{w}^t; \theta, \lambda) = \frac{\partial \mathcal{L}(\mathbf{w}^t; \theta, \lambda)}{\partial w_{j(t)}}.$$

Next, we introduce the necessary assumption and the definition of the projected gradient.

Assumption 1. For any $t \in \mathbb{N}$, we have

$$\mathbb{E}_{\mathcal{B}(t)} [|G(w_{j(t)}) - \nabla_{j(t)} \mathcal{L}(\mathbf{w}^t; \theta, \lambda)|^2] \leq (\sigma^t)^2, \quad (9)$$

where $\sigma^t > 0$ is some constant.

Definition 1 (Projected gradient (Ghadimi, Lan, and Zhang 2016)). Let S be a closed convex set with dimension N and the projected gradient is defined as:

$$\mathcal{G}_S(\mathbf{w}, \mathbf{g}, \gamma) = \frac{1}{\gamma}(\mathbf{w} - \mathcal{P}_S(\mathbf{w} - \gamma \mathbf{g})), \quad (10)$$

where $\mathbf{w} \in S$, $\mathbf{g} \in \mathbb{R}^N$ and $\gamma \in \mathbb{R}$.

Assumption 1 is a common assumption in stochastic optimization (Gu, Huo, and Huang 2019), which bounds the error between the stochastic gradient and the full gradient. Besides, based on Definition 1, we have that

$$\mathcal{G}_{[0,1]^N}(\mathbf{w}^t, G(w_{j(t)}) \mathbf{e}_{j(t)}, \gamma^t) = \frac{\mathbf{w}^t - \mathbf{w}^{t+1}}{\gamma^t} := \hat{D}^t \quad (11)$$

means the projected gradient generated in t -iteration, where \mathbf{e}_i is one unit vector with i -th coordinate holding the value of 1. Based on these, we provide the convergence of Algorithm 3 in Theorem 1.

Theorem 1. When Assumption 1 holds and we use Algorithm 3 to optimize \mathbf{w} , $L_{\max} > 0$ is the maximum lipschitz constant of sub-problems w.r.t. single coordinate $w_i, i \in [N]$, stepsizes $\{\gamma^t\}_{t=1}^{\infty}$ satisfy $0 < \gamma^{t+1} \leq \gamma^t < \frac{2}{L_{\max}}$ and

$$\sum_{t=1}^{\infty} \gamma^t = +\infty, \quad \sum_{t=1}^{\infty} \gamma^t (\sigma^t)^2 < \infty, \quad (12)$$

then there exists an index sub-sequence \mathcal{K} such that

$$\lim_{\substack{t \rightarrow \infty \\ t \in \mathcal{K}}} \mathbb{E} \|\hat{D}^t\|_2 = 0. \quad (13)$$

Remark 1. The above theorem shows that Algorithm 3 approaches to a stationary point of the sub-problem w.r.t. sample weights. It indicates that our algorithm can assign appropriate sample weights to guide the model training.

Convergence of Algorithm 1: Before providing the convergence of our BSPML algorithm, *i.e.*, Algorithm 1, we introduce the following assumption:

Assumption 2. If we call the gradient-based algorithm to minimize $F(\mathbf{X})$ with an initial solution \mathbf{X}^0 , we have that $F(\mathbf{X}^0) \geq F(\mathbf{X}^t)$ when the iteration number t is large enough.

It is easy to verify that Assumption 2 is a basic requirement for a gradient-based optimizer no matter whether the optimization objective is convex or not. At this time, we give the convergence of our BSPML as follows.

Theorem 2. If Assumption 2 holds, the objective function sequence $\{\mathcal{L}(\theta^t, \mathbf{w}^t; \lambda^t)\}_{t=1}^T$ generated by the BSPML algorithm converges with the following property:

$$\lim_{t \rightarrow \infty} \|\mathcal{L}(\theta^t, \mathbf{w}^t; \lambda^t) - \mathcal{L}(\theta^{t-1}, \mathbf{w}^{t-1}; \lambda^{t-1})\| = 0.$$

Remark 2. The above theorem shows that when the iteration number is large enough, our objective function can converge to a fixed value by using our BSPML algorithm.

Experiments

In this section, we present experimental results to demonstrate the superiority of our BSPML algorithm.

Experimental Setup

Data Sets: We conduct experiments on four standard data sets: CUB200, Cars-196, In-Shop and Stanford Online Products (SOP). For all data sets, we use half of the classes for training and the rest for testing. In addition, to test the robustness of all methods, we construct artificial noisy data sets. Specifically, we stochastically select samples from each class in the training set and change their labels to other classes. Since there are only a few samples in most classes of the SOP data set, making it noisy would seriously break the nature of this data set. Thus, we did not conduct experiments on the noisy SOP data set.

Compared Algorithms: We compare our algorithm with classic DML algorithms, *i.e.*, contrastive loss (Hadsell, Chopra, and LeCun 2006) and triplet_{smooth} loss (Hermans, Beyer, and Leibe 2017), state-of-the-art DML algorithms, *i.e.*, margin loss (Wu et al. 2017), FastAP loss (Cakir et al. 2019) and MS loss (Wang et al. 2019), and robust DML algorithms, *i.e.*, RDML loss (Al-Obaidi, Zabihzadeh, and Hajjabad 2020), Proxy-Anchor loss (Kim et al. 2020) and SNR loss (Yuan et al. 2019).

Design of Experiments: All algorithms are implemented based on the open PyTorch package (Musgrave, Belongie, and Lim 2020) using the same network structure with the embedding size 512. All the network parameters are optimized by SGD with the learning rate $5e-6$ and the batch size 64. Specifically, one batch is constructed by first sampling $P = 16$ classes, and then sampling $K = 4$ images for each class. For contrastive loss and SNR loss, the positive margin is set to 1 and the negative margin is set to 0. As original paper shows, we set $\alpha = 0.2, \beta^{(0)} = 1.2$ and $\beta^{(class)} = \beta^{(img)} = 0$ for margin loss. For FastAP loss, the number of soft histogram bins is set to 10 recommended by the authors. η is selected from $\{0.5, 1, 2, 3, 4\}$ for RDML loss. For Proxy-Anchor loss, $\alpha = 32, \delta = 0.1$ and all proxies are initialized using a normal distribution. For MS loss and our BSPML, ϵ is set to 0.1 and $\alpha = 2, \rho = 1, \beta = 50$.

Considering hyper-parameters μ and λ , we design experiments to analyze their roles. Note that we introduce the

	CUB							Cars						
	R@1	R@2	R@4	R@8	R@16	R@32	NMI	R@1	R@2	R@4	R@8	R@16	R@32	NMI
Contrastive	49.6	62.1	73.3	82.7	89.8	94.6	57.7	56.4	68.1	77.2	85.7	91.3	95.6	51.8
Triplet _{smooth}	42.0	54.5	66.4	77.6	86.5	92.7	52.3	37.4	49.9	62.3	73.7	83.6	90.7	47.9
Margin	38.0	49.9	62.7	74.6	84.8	91.4	51.4	42.8	54.6	66.5	77.7	86.7	93.1	49.4
FastAP	42.9	55.0	67.3	78.0	86.8	92.9	52.7	33.6	46.7	59.5	71.6	82.3	90.6	49.4
MS	49.2	61.5	73.5	83.1	89.5	94.6	58.2	58.1	70.2	80.0	87.5	92.6	96.1	54.4
RDML	43.5	56.8	69.1	80.1	88.5	93.9	53.1	37.7	51.5	64.7	76.1	85.5	92.3	49.1
Proxy-Anchor	48.9	61.6	72.9	82.7	89.4	94.6	57.9	58.5	70.4	80.2	87.4	92.7	96.5	54.6
SNR	47.8	61.5	72.9	82.6	89.3	94.3	56.4	56.0	67.9	78.1	86.1	91.9	95.5	51.2
BSPML	50.6	62.8	74.1	83.5	89.7	94.9	58.9	59.6	71.4	80.3	87.6	92.9	96.4	55.4
	In-shop							SOP						
	R@1	R@2	R@4	R@8	R@16	R@32	NMI	R@1	R@8	R@16	R@32	R@64	R@128	NMI
Contrastive	73.1	80.9	86.8	90.8	93.4	95.5	82.9	65.0	79.5	83.1	86.2	88.7	91.1	88.3
Triplet _{smooth}	67.6	77.0	84.7	89.8	93.5	96.0	83.2	61.6	77.4	81.6	85.3	88.4	91.1	87.8
Margin	55.7	66.5	75.9	83.2	88.9	93.0	80.1	51.4	69.9	75.3	80.4	84.8	88.7	86.2
FastAP	64.6	74.2	82.3	88.1	92.1	94.7	82.3	58.6	76.0	80.5	84.6	88.0	91.0	87.5
MS	73.4	80.7	86.4	90.6	93.4	95.3	84.7	65.8	79.4	83.0	86.2	89.2	91.7	88.6
RDML	69.6	78.7	85.4	90.3	93.5	95.7	83.0	62.8	78.0	82.1	85.6	88.7	90.7	88.1
Proxy-Anchor	74.2	81.2	86.9	91.0	93.7	95.5	85.0	66.3	80.2	83.3	86.2	88.7	90.9	88.7
SNR	72.7	80.6	87.1	91.3	94.1	96.2	83.6	65.3	79.0	82.3	85.5	88.3	90.6	88.3
BSPML	75.0	82.1	87.8	91.6	94.2	96.0	85.9	66.3	80.9	84.4	87.5	90.1	92.5	89.1

Table 1: Retrieval and clustering performance (%) on original data sets.

	NR	Contrastive	Triplet _{smooth}	Margin	FastAP	MS	RDML	Proxy-Anchor	SNR	BSPML
CUB	10%	48.26±0.21	39.69±0.17	37.72±0.18	41.57±0.13	48.46±0.16	41.59±0.15	48.31±0.12	45.95±0.19	50.08 ± 0.18
	20%	47.79±0.19	36.25±0.17	35.06±0.20	38.80±0.15	47.68±0.16	39.63±0.16	47.62±0.14	45.16±0.17	49.38 ± 0.20
	30%	46.25±0.24	33.92±0.20	34.36±0.18	37.10±0.15	46.25±0.14	37.57±0.16	46.66±0.15	44.43±0.21	48.74 ± 0.19
Car	10%	55.24±0.26	36.58±0.18	40.24±0.18	32.33±0.12	57.11±0.15	37.09±0.16	57.84±0.15	55.04±0.17	58.78 ± 0.18
	20%	54.55±0.24	35.76±0.17	38.49±0.21	31.68±0.10	56.37±0.17	36.32±0.15	57.18±0.15	53.49±0.18	58.24 ± 0.16
	30%	52.94±0.22	34.81±0.19	37.67±0.17	31.04±0.10	55.14±0.18	35.55±0.15	56.06±0.16	51.51±0.20	57.38 ± 0.19
In-shop	10%	69.59±0.29	59.69±0.21	50.33±0.20	56.28±0.18	70.33±0.25	62.50±0.22	70.51±0.17	67.24±0.25	71.44 ± 0.23
	20%	63.34±0.27	53.20±0.18	45.74±0.21	52.03±0.13	65.68±0.20	57.73±0.20	66.38±0.17	60.36±0.23	67.91 ± 0.25
	30%	56.76±0.22	48.35±0.17	43.16±0.15	49.37±0.13	58.85±0.24	53.97±0.21	60.25±0.15	53.62±0.23	62.04 ± 0.23

Table 2: Recall@1 results (%) with the corresponding standard deviation on noisy data sets with different Noise Ratios (NR).

Mean of Average Weights of classes (MAW) and the Standard Deviation of Average Weights of classes (SDAW):

$$\begin{aligned}
 \text{MAW} &= \frac{1}{C} \sum_{c=1}^C \frac{\sum_{i=1}^{N^c} w_i^c}{N^c}, \\
 \text{SDAW} &= \left(\frac{1}{C} \sum_{c=1}^C \left(\frac{\sum_{i=1}^{N^c} w_i^c}{N^c} - \text{MAW} \right)^2 \right)^{\frac{1}{2}}.
 \end{aligned} \tag{14}$$

MAW represents the average weight of all samples and SDAW implies the balance degree between the average sample weights of classes. The smaller the value of SDAW, the more balanced the average sample weights of classes. Moreover, to verify the feasibility of each part of our BSPML, we carry out ablation experiments with varying embedding sizes {64, 128, 256, 512} on noisy data sets.

For the retrieval task, all algorithms are evaluated by the standard metric Recall@K. To calculate Recall@K, each testing sample first retrieves K nearest neighbors from the test set and receives a score 1 if a sample of the same class is retrieved among the K nearest neighbors. Considering the clustering performance, we utilize the normalized mutual information (NMI) score: $\text{NMI}(\Omega, \mathcal{C}) = \frac{2I(\Omega, \mathcal{C})}{H(\Omega) + H(\mathcal{C})}$, where Ω

denotes the real clustering result and \mathcal{C} denotes the set of clusters obtained by K-means. Here, $I(\cdot)$ represents the mutual information and $H(\cdot)$ represents the entropy.

Results and Discussion

Table 1 presents Recall@K and NMI performance on four standard data sets. Benefiting from the SPL strategy, our BSPML avoids getting stuck into one bad local optimal solution and thus achieves better generalization ability than MS loss that is the non-SPL version of our BSPML. Meanwhile, compared with other state-of-the-art DML algorithms, our BSPML also obtains better performance.

Table 2 shows Recall@1 performance on three noisy data sets with different noise ratios (from 10% to 30%). These results show that our BSPML obviously achieves better performance than non-robust DML algorithms. Compared with robust DML algorithms, our BSPML also has sufficient advantages. Specifically, while SNR loss seems to be helpful against the noisy feature, it has limited robustness against the noisy label that is a much trickier challenge. Proxy-Anchor loss and RDML loss attempt to reduce the influence of noisy samples, but they are still sensitive to large noise ratios. Our BSPML is able to exclude noisy samples from the training




Figure 3: MAW results with different values of μ and λ on data sets with 20% noisy samples.




Figure 4: SDAW results with different values of μ and λ on data sets with 20% noisy samples.




Figure 5: Recall@1 results (%) with different embedding sizes on data sets with 20% noisy samples.

process directly and thus has better robustness.

Fig. 3 and Fig. 4 show the effect of λ and μ on the sample weight. Note that MAW (14) represents the average weight of all samples and SDAW (14) implies the balance degree between the average sample weights of classes. Fig. 3 clearly reveals that with the increase of λ , our BSPML tends to assign samples larger weights and thus allows hard samples with larger losses to join the training. Obviously, this phenomenon is consistent with the classic SPL strategy. Moreover, as shown in Fig. 4, the average sample weights of classes are balanced as much as possible when μ is set to a large value. This phenomenon shows the critical role of our balance regularization term.

Fig. 5 shows the Recall@1 performance of ablation experiments with varying embedding sizes $\{64, 128, 256, 512\}$ on the data sets with 20% noisy samples, where the SPMS is generated by introducing SPL strategy into MS loss without the balance regularization term. Compared with the original MS loss, the performance of SPMS is improved regard-

less of the embedding size. However, SPMS still faces the challenge of the unbalanced average sample weights among classes. Benefiting from the balance regularization term, our BSPML achieves the best performance.

Conclusion

In this paper, we build a connection between noisy samples and hard samples in the framework of self-paced learning and propose a Balanced Self-Paced Metric Learning (BSPML) algorithm with a novel denoising multi-similarity formulation to deal with noisy data in DML effectively. Specifically, our BSPML algorithm treats noisy samples as extremely hard samples and adaptively excludes them from the model training by sample weighting. Experimental results demonstrate that our BSPML algorithm has better robustness than the state-of-the-art robust DML approaches.

Acknowledgments

Bin Gu was partially supported by the National Natural Science Foundation of China (No:61573191).

References

- Al-Obaidi, S. A. R.; Zabihzadeh, D.; and Hajiabadi, H. 2020. Robust metric learning based on the rescaled hinge loss. *International Journal of Machine Learning and Cybernetics*, 11: 2515–2528.
- Bromley, J.; Bentz, J. W.; Bottou, L.; Guyon, I.; LeCun, Y.; Moore, C.; Säckinger, E.; and Shah, R. 1993. Signature verification using a “siamese” time delay neural network. *International Journal of Pattern Recognition and Artificial Intelligence*, 7(04): 669–688.
- Bucher, M.; Herbin, S.; and Jurie, F. 2016. Improving semantic embedding consistency by metric learning for zero-shot classification. In *European Conference on Computer Vision*, 730–746. Springer.
- Cakir, F.; He, K.; Xia, X.; Kulis, B.; and Sclaroff, S. 2019. Deep metric learning to rank. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 1861–1870.
- Chen, K.; Chen, Y.; Han, C.; Sang, N.; and Gao, C. 2020. Hard sample mining makes person re-identification more efficient and accurate. *Neurocomputing*, 382: 259–267.
- Frénay, B.; and Verleysen, M. 2013. Classification in the presence of label noise: a survey. *IEEE transactions on neural networks and learning systems*, 25(5): 845–869.
- Ghadimi, S.; Lan, G.; and Zhang, H. 2016. Mini-batch stochastic approximation methods for nonconvex stochastic composite optimization. *Mathematical Programming*, 155(1-2): 267–305.
- Grabner, A.; Roth, P. M.; and Lepetit, V. 2018. 3d pose estimation and 3d model retrieval for objects in the wild. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 3022–3031.
- Gu, B.; Huo, Z.; and Huang, H. 2019. Scalable and Efficient Pairwise Learning to Achieve Statistical Accuracy. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, 3697–3704.
- Gu, B.; Zhai, Z.; Li, X.; and Huang, H. 2021. Finding Age Path of Self-Paced Learning. In *2021 IEEE International Conference on Data Mining (ICDM)*, 151–160. IEEE.
- Hadsell, R.; Chopra, S.; and LeCun, Y. 2006. Dimensionality reduction by learning an invariant mapping. In *2006 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR’06)*, volume 2, 1735–1742. IEEE.
- He, X.; Zhou, Y.; Zhou, Z.; Bai, S.; and Bai, X. 2018. Triplet-center loss for multi-view 3d object retrieval. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 1945–1954.
- Hermans, A.; Beyer, L.; and Leibe, B. 2017. In defense of the triplet loss for person re-identification. *arXiv preprint arXiv:1703.07737*.
- Ketkar, N. 2017. Stochastic gradient descent. In *Deep learning with Python*, 113–132. Springer.
- Kim, S.; Kim, D.; Cho, M.; and Kwak, S. 2020. Proxy anchor loss for deep metric learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 3238–3247.
- Kingma, D. P.; and Ba, J. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- Kumar, M. P.; Packer, B.; and Koller, D. 2010. Self-paced learning for latent variable models. In *Advances in neural information processing systems*, 1189–1197.
- Leal-Taixé, L.; Canton-Ferrer, C.; and Schindler, K. 2016. Learning by tracking: Siamese CNN for robust target association. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, 33–40.
- Li, W.; Wang, L.; Li, W.; Agustsson, E.; and Van Gool, L. 2017. Webvision database: Visual learning and understanding from web data. *arXiv preprint arXiv:1708.02862*.
- Meng, D.; Zhao, Q.; and Jiang, L. 2017. A theoretical understanding of self-paced learning. *Information Sciences*, 414: 319–328.
- Musgrave, K.; Belongie, S.; and Lim, S.-N. 2020. A metric learning reality check. In *European Conference on Computer Vision*, 681–699. Springer.
- Nguyen, D. T.; Mummadi, C. K.; Ngo, T. P. N.; Nguyen, T. H. P.; Beggel, L.; and Brox, T. 2019. Self: Learning to filter noisy labels with self-ensembling. *arXiv preprint arXiv:1910.01842*.
- Ren, Y.; Huang, S.; Zhao, P.; Han, M.; and Xu, Z. 2020. Self-paced and auto-weighted multi-view clustering. *Neurocomputing*, 383: 248–256.
- Shi, W.; Gu, B.; Li, X.; Deng, C.; and Huang, H. 2021. Triply stochastic gradient method for large-scale nonlinear similar unlabeled classification. *Machine Learning*, 110(8): 2005–2033.
- Tao, R.; Gavves, E.; and Smeulders, A. W. 2016. Siamese instance search for tracking. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 1420–1429.
- Wang, C.; Peng, G.; and Lin, W. 2021. Robust local metric learning via least square regression regularization for scene recognition. *Neurocomputing*, 423: 179–189.
- Wang, H.; Nie, F.; and Huang, H. 2014. Robust distance metric learning via simultaneous l_1 -norm minimization and maximization. In *International conference on machine learning*, 1836–1844. PMLR.
- Wang, X.; Han, X.; Huang, W.; Dong, D.; and Scott, M. R. 2019. Multi-similarity loss with general pair weighting for deep metric learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 5022–5030.
- Wang, Y.; Zhou, W.; Lv, Q.; and Yao, G. 2022. MetricMask: Single Category Instance Segmentation by Metric Learning. *Neurocomputing*.

Wen, Y.; Zhang, K.; Li, Z.; and Qiao, Y. 2016. A discriminative feature learning approach for deep face recognition. In *European conference on computer vision*, 499–515. Springer.

Wohlhart, P.; and Lepetit, V. 2015. Learning descriptors for object recognition and 3d pose estimation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 3109–3118.

Wu, C.-Y.; Manmatha, R.; Smola, A. J.; and Krahenbuhl, P. 2017. Sampling matters in deep embedding learning. In *Proceedings of the IEEE International Conference on Computer Vision*, 2840–2848.

Wu, Z.-F.; Wei, T.; Jiang, J.; Mao, C.; Tang, M.; and Li, Y.-F. 2021. NGC: a unified framework for learning with open-world noisy data. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 62–71.

Xing, E. P.; Ng, A. Y.; Jordan, M. I.; and Russell, S. 2002. Distance metric learning with application to clustering with side-information. In *NIPS*, volume 15, 12. Citeseer.

Yao, Y.; Sun, Z.; Zhang, C.; Shen, F.; Wu, Q.; Zhang, J.; and Tang, Z. 2021. Jo-src: A contrastive approach for combating noisy labels. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 5192–5201.

Yelamathi, S. K.; Reddy, S. K.; Mishra, A.; and Mittal, A. 2018. A zero-shot framework for sketch based image retrieval. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 300–317.

Yin, T.; Liu, N.; and Sun, H. 2021. Self-paced active learning for deep CNNs via effective loss function. *Neurocomputing*, 424: 1–8.

Yu, R.; Dou, Z.; Bai, S.; Zhang, Z.; Xu, Y.; and Bai, X. 2018a. Hard-aware point-to-set deep metric for person re-identification. In *Proceedings of the European conference on computer vision (ECCV)*, 188–204.

Yu, X.; Liu, T.; Gong, M.; and Tao, D. 2018b. Learning with biased complementary labels. In *Proceedings of the European conference on computer vision (ECCV)*, 68–83.

Yuan, T.; Deng, W.; Tang, J.; Tang, Y.; and Chen, B. 2019. Signal-to-noise ratio: A robust distance metric for deep metric learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 4815–4824.

Zhai, Z.; Gu, B.; Li, X.; and Huang, H. 2020. Safe sample screening for robust support vector machine. In *Proceedings of the AAAI conference on artificial intelligence*, volume 34, 6981–6988.

Zhang, X.; Wu, X.; Chen, F.; Zhao, L.; and Lu, C.-T. 2020. Self-Paced Robust Learning for Leveraging Clean Labels in Noisy Data. In *AAAI*, 6853–6860.

Zhang, Z.; and Saligrama, V. 2016. Zero-shot learning via joint latent similarity embedding. In *proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 6034–6042.

Zhu, C.; Chen, W.; Peng, T.; Wang, Y.; and Jin, M. 2021. Hard Sample Aware Noise Robust Learning for Histopathology Image Classification. *IEEE Transactions on Medical Imaging*, 41(4): 881–894.