

Rethinking the Optimization of Average Precision: Only Penalizing Negative Instances before Positive Ones is Enough

Zhuo Li^{1,2*}, Weiqing Min^{1,2†}, Jiajun Song^{1,2}, Yaohui Zhu^{1,2}, Liping Kang³,
Xiaoming Wei³, Xiaolin Wei³ and Shuqiang Jiang^{1,2}

¹Key Lab of Intell. Info. Process., Inst. of Comput. Tech., CAS, Beijing, China

²University of Chinese Academy of Sciences, Beijing, China

³Meituan, Beijing, China

{zhuo.li, jiajun.song, yaohui.zhu}@vip.l.ict.ac.cn, {minweiqing, sqjiang}@ict.ac.cn,
{kangliping, weixiaoming, weixiaolin02}@meituan.com

Abstract

Optimising the approximation of Average Precision (AP) has been widely studied for image retrieval. Limited by the definition of AP, such methods consider both negative and positive instances ranking before each positive instance. However, we claim that only penalizing negative instances before positive ones is enough, because the loss only comes from these negative instances. To this end, we propose a novel loss, namely Penalizing Negative instances before Positive ones (PNP), which can directly minimize the number of negative instances before each positive one. In addition, AP-based methods adopt a fixed and sub-optimal gradient assignment strategy. Therefore, we systematically investigate different gradient assignment solutions via constructing derivative functions of the loss, resulting in PNP-I with increasing derivative functions and PNP-D with decreasing ones. PNP-I focuses more on the hard positive instances by assigning larger gradients to them and tries to make all relevant instances closer. In contrast, PNP-D pays less attention to such instances and slowly corrects them. For most real-world data, one class usually contains several local clusters. PNP-I blindly gathers these clusters while PNP-D keeps them as they were. Therefore, PNP-D is more superior. Experiments on three standard retrieval datasets show consistent results with the above analysis. Extensive evaluations demonstrate that PNP-D achieves the state-of-the-art performance. Code is available at <https://github.com/interestingzhuo/PNPloss>

Introduction

Image retrieval (Radenović, Tolias, and Chum 2019; Min et al. 2020) refers to finding all images containing relevant content with the query from the database. One important issue is to design the training objective. Recently, some works propose to directly optimize the Average Precision (AP) (Revaud et al. 2019; Brown et al. 2020; Cakir et al. 2019) to achieve end-to-end training. Considering AP is non-differentiable and can't be directly optimized by the gradient descent, these works leverage some transformations to approximate AP to make it differentiable. These methods greatly promote the development of retrieval.

*Work done during an internship at Meituan.

†Corresponding author.

Copyright © 2022, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.




Figure 1: The functional image of AP-based loss L_{AP} : $R(i, S_P)$ and $R(i, S_N)$ are the number of positive and negative instances before the instance i , respectively, where S_P is the similarity set between the query and positive instances and S_N is the similarity set between the query and negative instances. Arrows are the optimal path of positive instances.

However, there are two problems for AP-based methods: **(1) Redundancy exists in the optimization of AP.** Fig. 1 shows the functional image of AP-based loss L_{AP} . As shown in Fig. 1, the goal of minimizing L_{AP} is equivalent to minimize $R(i, S_N)$, because $L_{AP} = 0$ only when $R(i, S_N) = 0$. Therefore, the computation of $R(i, S_P)$ is redundant. **(2) Sub-optimal gradient assignment strategy.** Limited by the definition of AP, AP-based methods only adopt one specific gradient assignment strategy. Specifically, AP-based losses assign smaller gradient to larger $R(i, S_N)$, as shown in Fig. 1. Different gradient assignment strategies may result in different performance. Therefore, it leaves room for us to systematically explore different gradient assignment solutions to find potential better ones.

For the first problem, we propose a novel loss, namely Penalizing Negative instances before Positive ones (PNP), which directly minimizes the number of negative instances before each positive one. Therefore, each negative instance will receive penalties from all positive ones after it. In terms of each negative instance, if it ranks before more positive ones, it will receive more penalties and be corrected more quickly. As shown in Fig. 2, image (a) will receive more penalty because there are more positive instances after it

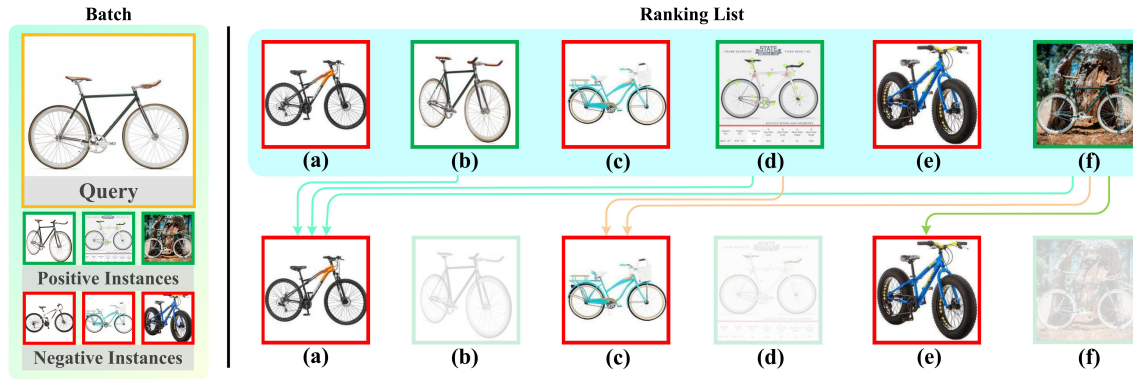


Figure 2: Penalizing Negative instances before Positive ones (PNP): PNP directly improves the retrieval performance by penalizing negative instances before positive ones. Each arrow indicates the penalty of each negative instance. Each negative instance will receive penalties from all positive ones after it, thus the negative instance before more positive ones will receive stronger penalty and be quickly corrected.

and it will be quickly corrected. In contrast, image (e) will be slowly corrected because that there are fewer positive instances (image (f)) after it.

For the second problem, we systematically investigate different gradient assignment solutions. We find that the derivative function of the loss function defines different gradient assignment solutions. To this end, we construct different derivative functions, which result in two different types of PNP, namely PNP-I with increasing derivative functions and PNP-D with decreasing ones. PNP-I assigns larger gradients to the positive instances with more negative ones before and tries to gather all relevant samples together. In contrast, PNP-D assigns smaller gradients to such positive instances. Therefore, PNP-D will slowly correct these instances, which probably belong to another center of the corresponding category. As shown in Fig. 2, PNP-I assigns larger gradients to image (f) and tries to quickly correct it, because there are more negative instances before it. In contrast, PNP-D assigns smaller gradient to image (f) by considering it probably belongs to another center of the corresponding category. For the real-world data, one class usually contains several local clusters, and PNP-D is more suitable for such case, thus may be more superior. Extensive evaluation on three standard retrieval datasets shows that PNP-D achieves the state-of-the-art performance.

In summary, our major contributions are as follows:

- We propose a novel loss function named PNP, which can improve the approximation of AP by simply ignoring positive instances before the target positive one.
- To our knowledge, we are the first research to design loss functions via the construction of their derivative functions. Based on the proposed method, we systematically investigate different gradient assignment solutions and obtain more appropriate loss functions.
- Experimental results on three benchmark datasets demonstrate that the proposed method achieves the state-of-the-art performance.

Related Work

Image Retrieval. How to obtain a compact image descriptor for retrieval has been widely studied (Radenović, Tolias, and Chum 2019; Min et al. 2020; Tolias, Jeníček, and Chum 2020; Cao, Araujo, and Sim 2020). Previous works bend their efforts for generating descriptors through an aggregation of local features, such as Fisher vectors (Perronnin et al. 2010) and VLAD (Arandjelović and Zisserman 2013; Jégou et al. 2010). More recently, CNNs have made impressive progress for image retrieval (Arandjelović et al. 2018; Babenko et al. 2014; Albert et al. 2016; Radenović, Tolias, and Chum 2016; Wang et al. 2014), because of their powerful nonlinear fitting capabilities. Besides studying the design of neural network structure to generate discriminative descriptors, there are also many researches on how to design the training objective (Revaud et al. 2019; Hadsell, Chopra, and Lecun 2006; Guiver and Snelson 2008; Sohn 2016; Sun et al. 2020; Qian et al. 2019; Gu, Ko, and Kim 2021), which brings out another research topic, deep metric learning.

Deep Metric Learning. Deep metric learning has been widely studied because of its important role in many tasks (Kim and Park 2021; Fu et al. 2021), such as image retrieval. How to design the objective for such task is one important topic. A wide variety of loss functions have been proposed in previous works. These loss functions (Schroff, Kalenichenko, and Philbin 2015; Ji et al. 2017; Engilberge et al. 2019) minimize the intra-class distance while maximize the inter-class distance by considering the constructed pairs, tuples, etc. For example, the triplet loss (Schroff, Kalenichenko, and Philbin 2015) directly optimizes the relative ranking of positive and negative instances given an anchor. However, there are too many pairs or tuples, and most of them are useless for training. To solve this problem, many methods have been proposed, such as pretraining for classification, combining multiple losses (Milbich et al. 2020a,b) and using complex sample mining strategies (Manmatha et al. 2017).

Optimizing Average Precision (AP). Directly optimising AP has been widely studied in the retrieval community. Various methods (Cakir et al. 2019; He et al. 2018; He, Lu, and

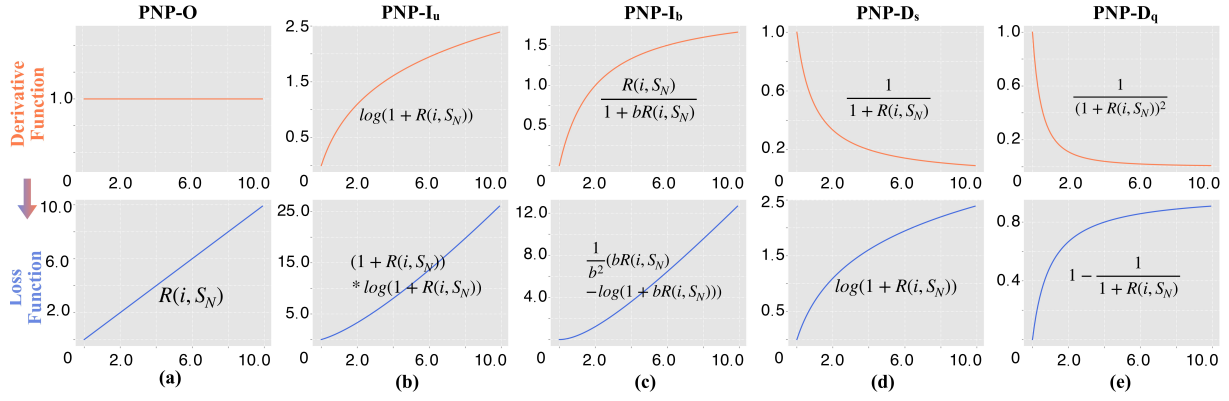


Figure 3: PNP Loss: Derivative functions and loss functions of different variants. PNP-I_u and PNP-I_b have increasing derivative functions, while PNP-D_s and PNP-D_q have decreasing derivative functions.

Sclaroff 2018; Revaud et al. 2019) have been proposed to overcome the non-differentiability in optimizing AP, such as optimizing a structured hinge-loss upper bound to the AP (Mohapatra et al. 2018), optimizing an approximation of AP derived from distance quantization (Cakir et al. 2019) and smoothing the path towards optimization (Brown et al. 2020). These approximations rely on the definition of AP to achieve better retrieval performance.

Different from above works, we propose the PNP loss, which directly minimizes the number of negative instances before each positive one to improve the retrieval performance. We further investigate different variants of PNP to explore potential better gradient assignment solutions via the construction of derivative functions. Extensive evaluations demonstrate that the proposed PNP loss achieves the state-of-the-art performance on three benchmark datasets.

Methodology

Given a query, the goal of retrieval systems is to rank all instances in a retrieval set $\Omega = \{I_i, i = 0, \dots, m\}$ based on their similarities to the query, where I_i is an instance and m is the size of the retrieval set. For each query I_q , the set is split into a positive set P_q and a negative set N_q , which are composed by instances of the same class and of different classes, respectively.

For the query I_q , the similarity set S_Ω of all instances in the retrieval set are measured via a chosen similarity. In this paper, we use the cosine similarity, and S_Ω can be defined as:

$$S_\Omega = \{s_i = \cos(v_q, v_i), i = 0, \dots, m\} \quad (1)$$

where v_q, v_i are the embeddings of I_q and I_i , respectively. S_Ω can be divided into S_P and S_N , i.e., $S_\Omega = S_P \cup S_N$, where $S_P = \{s_i, \forall i \in P_q\}$ and $S_N = \{s_i, \forall i \in N_q\}$ denote the positive and negative similarity set, respectively.

Based on the similarity set, the number of negative instances before I_i can be defined as:

$$R(i, S_N) = \sum_{j \in S_N, j \neq i} \mathbb{I}\{s_j - s_i > 0\} \quad (2)$$

where $\mathbb{I}\{\cdot\}$ is an Indicator function.

Note that $R(i, S_N)$ is not differentiable due to the Indicator function. Similar to (Brown et al. 2020), we relax the Indicator function with a sigmoid function $G(\cdot; \tau)$, where τ refers to the temperature for adjusting the sharpness:

$$G(x; \tau) = \frac{1}{1 + e^{-\frac{x}{\tau}}} \quad (3)$$

Substituting $G(\cdot; \tau)$ into Eq. 2, we can easily make $R(i, S_N)$ differentiable.

Definition of PNP

AP-based loss is defined as Eq. 4. The goal of minimizing L_{AP} is equivalent to minimize $R(i, S_N)$, and thus the computation of $R(i, S_P)$ in Eq. 4 is redundant.

$$L_{AP} = 1 - \frac{1}{|S_P|} \sum_{i \in S_P} \frac{1 + R(i, S_P)}{(1 + R(i, S_P) + R(i, S_N))} \quad (4)$$

Therefore, we propose a PNP loss to directly minimize $R(i, S_N)$. An intuitive loss function L_O can be defined as the number of negative instances before each positive one:

$$L_O = \frac{1}{|S_P|} \sum_{i \in S_P} R(i, S_N) \quad (5)$$

where $|S_P|$ is the normalization factor to remove the influence of the number of positive instances in each mini-batch. Substituting Eq. 2 into Eq. 5, we can get the original PNP. As shown in Fig. 2, each negative instance will receive penalties from all the positive ones ranking after it. Therefore, negative instance with more positive ones after will receive more penalties and will be quickly corrected. The PNP-O and its derivative function are shown in Fig. 3 (a), and the gradient assignment of PNP-O is constant.

Variants of PNP

Different gradient assignment strategies can result in different performance. In this section, we systematically investigate different gradient assignment solutions. Specifically, we design different variants of PNP by constructing their derivative functions, resulting in PNP-Increasing (PNP-I) with increasing derivative functions and PNP-Decreasing (PNP-D) with decreasing ones.

PNP-Increasing (PNP-I) According to the above definition, we should guarantee that the derivative function of PNP-I is increasing. Without losing generality, we can divide PNP-I into PNP-I_u with unbounded derivative function and PNP-I_b with bounded derivative functions.

As for PNP-I_u, we use Eq. 6 as its derivative functions. The value of this function will go down to zero when the loss is zero and go up to infinite when the loss tends to be infinity. Then we get PNP-I_u via infinitesimal analysis, shown in Eq. 7. L_{I_u} has no hyper-parameters, and thus reduces the difficulty of training.

$$\frac{\partial L_{I_u}}{\partial R(i, S_N)} = \log(1 + R(i, S_N)) \quad (6)$$

$$L_{I_u} = \frac{1}{|S_P|} \sum_{i \in S_P} (1 + R(i, S_N)) * \log(1 + R(i, S_N)) \quad (7)$$

As for PNP-I_b, in order to satisfy the two requirements: increasing and bounded, we use Eq. 8 as its derivative function. We introduce the parameter b to achieve an adjustable boundary. Then we can also get PNP-I_b as Eq. 9.

$$\frac{\partial L_{I_b}}{\partial R(i, S_N)} = \frac{R(i, S_N)}{1 + bR(i, S_N)} \quad (8)$$

$$L_{I_b} = \frac{1}{|S_P|} \sum_{i \in S_P} \frac{1}{b^2} (bR(i, S_N) - \log(1 + bR(i, S_N))) \quad (9)$$

Loss functions and their derivative functions are shown in Fig. 3 (b)&(c). The derivative functions of these two functions are both increasing, which means that positive instances with more negative ones before will receive more penalties and these two loss functions will try to make such hard positive instances closer to the query.

PNP-Decreasing (PNP-D) Specifically, we choose a common used decreasing function as the derivative functions of PNP-D, shown in Eq. 10 and Eq. 11. The introduced parameter α controls the speed of the descent, and larger α leads to faster speed. These two derivative functions result in PNP-D_s with slow speed and PNP-D_q with fast speed. After infinitesimal analysis, the PNP-D_s and PNP-D_q will be Eq. 12 and Eq. 13, respectively.

$$\frac{\partial L_{D_s}}{\partial R(i, S_N)} = \frac{1}{(1 + R(i, S_N))} \quad (10)$$

$$\frac{\partial L_{D_q}}{\partial R(i, S_N)} = \frac{\alpha}{(1 + R(i, S_N))^{\alpha+1}} \quad (\alpha \geq 1) \quad (11)$$

$$L_{D_s} = \frac{1}{|S_P|} \sum_{i \in S_P} \log(1 + R(i, S_N)) \quad (12)$$

$$L_{D_q} = 1 - \frac{1}{|S_P|} \sum_{i \in S_P} \frac{1}{(1 + R(i, S_N))^\alpha} \quad (\alpha \geq 1) \quad (13)$$

Loss functions and their derivative functions of PNP-D are shown in Fig. 3 (d)&(e). The derivative functions of these two functions are both decreasing, which means that positive instances with fewer negative ones before will receive more penalties and these two loss functions will quickly correct such instances.




Figure 4: Comparison between PNP-I and PNP-D: The dashed arrows represent the directions of optimization. Different colors represent different categories. b and y represent centers of different categories. It is worth mentioning that b_1 and b_2 are different centers of class b . PNP-I blindly makes b_1 and b_2 closer while PNP-D keeps such instances distribute into two clusters.




Figure 5: Relation between PNP and AP-based loss: The gradients of AP-based loss with respect to $R(i, S_N)$ when $R(i, S_P) = 0, 1, 2, 3$. The gradients of PNP-D ($\alpha = 2$) is steeper than AP-based loss.

Discussion PNP-I assigns larger gradients to positive instances with more negative ones before, while PNP-D assigns smaller gradients to such positive instances. PNP-I tries to make all the relevant instances together. In contrast, PNP-D only quickly corrects positive instances with fewer negative ones before, because such samples are considered to belong to the same center with the query. As shown in Fig. 4, if b_1 and b_2 are sampled into one batch, PNP-I will try to make them gather by assigning larger gradients. In contrast, considering too many negative instances (yellow) between them, PNP-D will assign smaller gradients to them and keep the multi-center distribution for the blue class. PNP-D can satisfy such distribution which is also mentioned by (Qian et al. 2019), and thus is superior to PNP-I. Such analysis has been further verified via our comprehensive experimental evaluation.

We just conduct above four variants of PNP. Actually, our proposed design strategy is general and similar strategies can be further explored to find more solutions.

Relation between PNP and AP

In this section, we investigate the relation between PNP and AP. The derivative function of AP-based loss with respect to $R(i, S_N)$ is shown in Eq. 14 and its functional image is shown in Fig. 5. AP-based methods have similar gradient

Method	R@1 (%)	R@10 (%)	R@100 (%)	R@1K (%)	Dists@intra	Dists@inter	NMI (%)
PNP-O	77.9	90.5	96.3	98.9	0.318	0.793	90.0
PNP-I _u	76.3	89.6	96.0	98.9	0.301	0.771	89.6
PNP-I _b	77.9	90.4	96.3	98.9	0.302	0.756	90.0
PNP-D _s	79.6	91.4	96.6	99.0	0.363	0.849	90.3
PNP-D _q	80.1	91.5	96.7	99.0	0.380	0.946	90.4

Table 1: Results on different variants of the PNP loss on SOP (BS = 112).

Method	Small (%)		Medium (%)		Large (%)		Dists@intra	Dists@inter	NMI (%)
	R@1	R@5	R@1	R@5	R@1	R@5			
PNP-O	93.6	97.1	91.6	95.9	89.2	95.4	0.165	0.791	90.5
PNP-I _u	92.4	96.9	90.4	95.4	87.6	94.8	0.137	0.727	90.0
PNP-I _b	94.0	97.3	92.0	96.0	89.9	95.4	0.144	0.674	90.6
PNP-D _s	94.3	97.3	92.7	96.4	91.0	95.9	0.203	0.865	90.8
PNP-D _q	94.6	97.4	92.9	96.3	91.4	95.9	0.204	0.974	91.0

Table 2: Results on different variants of the PNP loss on VehicleID (BS = 112).

assignment with PNP-D.

$$\frac{\partial L_{AP}}{\partial R(i, S_N)} = \frac{1 + R(i, S_P)}{(1 + R(i, S_P) + R(i, S_N))^2} \quad (14)$$

However, AP-based methods use different gradient assignments by considering the positions of positive instances. When too many positive instances rank before the target positive instance ($R(i, S_P) = 3$), the gradient respect to $R(i, S_N)$ is almost to be equal for different $R(i, S_N)$. It weakens the advantage of such gradient assignment, probably leading to worse retrieval performance for AP. In contrast, PNP can further enhance such advantage and finally achieve competitive performance.

Experiment

Dataset

We use the following three popular datasets as the benchmark datasets.

Stanford Online Products (SOP) (Song et al. 2016) contains 120,053 product images divided into 22,634 classes. The training set contains 11,318 classes with 59,551 images and the rest 11,316 classes with 60,502 images are for testing.

VehicleID (Liu et al. 2016) contains 221,736 images of 26,267 vehicle categories, where 13,134 categories with 110,178 images are used for training. Following the same test protocol as (Liu et al. 2016), three test sets of increasing sizes are used for evaluation (termed small, medium, large), which contain 800 classes (7,332 images), 1,600 classes (12,995 images) and 2,400 classes (20,038 images), respectively.

INaturalist (Van Horn et al. 2018) is a large-scale animal and plant species classification dataset with 461,939 images from 8,142 classes. We follow the setting from (Brown et al. 2020) by keeping 5,690 classes for training, and 2,452 unseen classes for testing.

Experimental Setup

We use the convolutional layers of ResNet-50 pretrained on ImageNet (He et al. 2016) to perform the training. The models are optimized using Adam (Kingma and Ba 2015), we set

the initial learning rate 10^{-5} , weight decay 4×10^{-4} . During training, we randomly sample k classes and $|P|$ samples per class to form each mini-batch. Following standard practice, we resize images to 256×256 , and randomly crop them to 224×224 as input. Random flipping ($p = 0.5$) is used during training for data augmentation, and a single center crop of 224×224 is used during evaluation. We use 0 as the fixed random seed for all experiments to avoid seed-based performance fluctuations. During training, we directly set τ in Eq. 3 to 0.01 which is explored in (Brown et al. 2020).

For all the datasets, every instance from each class is used in turn as the query I_q , and the retrieval set Ω is formed as all the remaining instances. Recall@k (R@k) is adopted as the main evaluation metric. In order to evaluate the generalization of the model, we also display dists@intra (Mean Intraclass Distance), dists@inter (Mean Interclass Distance) (Roth et al. 2020) and Normalized Mutual Information (NMI) (Schütze, Manning, and Raghavan 2008) for further performance analysis. BS represents mini-batch size.

Performance of Different PNP Variants

In this section, we investigate different variants of PNP on two popular benchmark datasets. The batch size of experiments in Table 1 and Table 2 are both 112. We report the best performance of different variants and the impact of hyper-parameters is shown in Fig. 8.

We present R@k, Distance and NMI in Table 1 and Table 2. PNP-D consistently outperforms PNP-O and PNP-I on two benchmark datasets for R@k and achieves larger dists@intra, dists@inter and NMI, which shows better generalization of the model (Roth et al. 2020). Specifically, PNP-D is 2.1% higher than PNP-O on SOP benchmark and 0.9% higher on VehicleID for R@1. In contrast, PNP-I is worse than PNP on R@1, and it is about 1% lower than PNP on SOP benchmark 1-2% on VehicleID for R@1.

PNP-I assigns larger gradients to positive instances which have more negative ones before, thus it will enforce relevant images to gather and result in smaller dists@intra and dists@inter. In contrast, PNP-D assigns smaller gradients




Figure 6: Illustration of the effect of different variants. (a) training and test data (b) The network trained with PNP-I fails to separate all test classes due to the aggregation of all relevant instances. (c) The PNP-D successfully separates the test classes by keeping the intra-class variance.



Figure 7: Image retrieval examples on three datasets. Images with green border are positive instances and these with red border are negative ones for the query.

to such instances and slowly optimizes them by considering that they have a high probability of belonging to other centers of the corresponding category, and therefore it has larger dists@intra and dists@inter . By this strategy, PNP-D can adaptively achieve “multi-centers” for each category if necessary and automatically adjust the number of clusters. Such strategy also solves the error correction of noise by considering the noisy data as another center.

Previous works have found that larger dists@intra and dists@inter will produce better generalization (Roth et al. 2020). The results shown in Table 1 and Table 2 further confirm such conclusion. Compared to PNP-I, PNP-D has larger dists@intra and dists@inter , and results in larger NMI, which shows better generalization of the model. Similar to (Roth et al. 2020), we conduct one experiment to illustrate the performance of two loss functions. Specifically, we use a fully-connected network with two layers and 30 neurons for each layer. The dimensions of input and embeddings are both 2D, and the embeddings are normalized on a unit circle. As shown in Fig. 6 (a), each of four training and test classes contains 150 samples, respectively. We train the networks using PNP-I_b and PNP-D_q, respectively. Fig. 6 illustrates the effect of two loss functions. PNP-I fails to separate all test classes in Fig. 6 (b) while PNP-D successfully separates them (with larger intra-class variance) in Fig. 6 (c). It shows that PNP-D enables the model to capture more information and exhibits stronger generalization to the unseen test classes by retaining the variance within the class in the training process.

Method	R@1	R@10	R@100	R@1K
Hist. (Ustinova and Lempitsky 2016)	72.4	86.1	94.1	98.3
Margin (Manmatha et al. 2017)	72.7	86.2	93.8	98.0
Divide (Sanakoyeu et al. 2019)	75.9	88.4	94.9	98.1
FastAP (Cakir et al. 2019)	76.4	89.0	95.1	98.2
MIC (Brattoli et al. 2019)	77.2	89.4	95.6	-
SoftTriplet (Qian et al. 2019)	78.3	90.3	95.9	-
RankMI (Kemertas et al. 2020)	74.3	87.9	94.9	98.3
Blackbox AP (Rolínek et al. 2020)	78.6	90.5	96.0	98.7
Cont. w/M (Wang et al. 2020)	80.6	91.6	96.2	98.7
Pnca++ (Teh, DeVries, and Taylor 2020)	80.7	92.0	96.7	98.9
Smooth-AP (Brown et al. 2020)	80.1	91.5	96.6	99.0
DCML-MDW (Zheng et al. 2021)	79.8	90.8	95.8	-
PNP-D _q (BS = 112)	80.1	91.5	96.7	99.0
PNP-D _q (BS = 384)	81.1	92.2	96.8	99.0

Table 3: Performance comparison on SOP (%).

Method	Small		Medium		Large	
	R@1	R@5	R@1	R@5	R@1	R@5
Divide (Sanakoyeu et al. 2019)	87.7	92.9	85.7	90.4	82.9	90.2
MIC. (Brattoli et al. 2019)	86.9	93.4	-	-	82.0	91.0
FastAP (Cakir et al. 2019)	91.9	96.8	90.6	95.9	87.5	95.1
Cont. w/M (Wang et al. 2020)	94.7	96.8	93.7	95.8	93.0	95.8
Smooth-AP (Brown et al. 2020)	94.9	97.6	93.3	96.4	91.9	96.2
PNP-D _q (BS = 112)	94.6	97.4	92.9	96.3	91.4	95.9
PNP-D _q (BS = 384)	95.5	97.8	94.2	96.9	93.2	96.6

Table 4: Performance comparison on VehicleID (%).

Comparison with State-of-the-art

We compare PNP loss to the recent AP-based methods and a series of state-of-the-art deep metric learning methods on three standard benchmarks.

SOP For fair comparison, we use the same setting with (Brown et al. 2020). Table 3 shows the performance on SOP, we observe that PNP-D_q achieves the state-of-the-art results. In particular, our model outperforms other methods on all evaluation metrics and outperforms AP approximating methods by 0.5 - 4% on R@1 when the same batch size (384) and dimension (512) are used. Note that, although current work (Cont. w/M (Wang et al. 2020)) uses memory techniques to sample from many mini-batches simultaneously for each iteration, our best model still outperforms this method, only using a single mini-batch on each iteration.

VehicleID We further conduct experiments on VehicleID to verify the performance of PNP on large-scale retrieval datasets. Table 4 shows the results on the VehicleID dataset. For the same batch size (384) and dimension (512), we observe that PNP-D_q again achieves state-of-the-art performance on the large-scale VehicleID dataset on all the evaluation metrics. Such result shows that PNP-D also works better on large-scale datasets.

INaturalist As shown in Table 5, our method also outperforms the best method by 0.7% on R@1 for the experiments on INaturalist with the same batch size (224) and dimension (512). These results demonstrate that PNP is particularly suit-

Method	R@1	R@4	R@16	R@32
Triplet (Manmatha et al. 2017)	58.1	75.5	86.8	90.7
Proxy NCA (Yair et al. 2017)	61.6	77.4	87.0	90.6
FastAP (Cakir et al. 2019)	60.6	77.0	87.2	90.6
Blackbox AP (Rolínek et al. 2020)	62.9	79.0	88.9	92.1
Smooth-AP (Brown et al. 2020)	65.9	80.9	89.8	92.7
PNP-D _q (BS = 224)	66.6	81.1	89.7	92.6

Table 5: Performance comparison on INaturalist (%).

Method	R@1	R@10	R@100	R@1K
Smooth-AP	73.2	86.4	93.6	97.5
PNP-I _u	67.5	82.7	92.3	97.5
PNP-D _q	73.8	87.1	94.2	98.0

Table 6: Evaluation of robustness on SOP (%).

able for large-scale retrieval datasets, which demonstrates its scalability to real-world retrieval problems.

Impact of Hyper-parameters

To investigate the effect of different hyper-parameter settings, *i.e.*, steepness α in Eq. 13, boundary b in Eq. 8, samples per class $|P|$, we train ResNet-50 on SOP and VehicleID. The random seed is fixed to 0 and the batch size is set to 112 for all experiments in this section.

Steepness α : The influence of α of PNP-D_q is shown in Fig. 8 (a). We can find that larger α results in larger R@1. It can be explained that larger α corresponds to more stable optimization (correct positive instances with sufficient confidence), and thus achieves better performance. However, when α is too large, the training phrase will crash due to the gradient exploding. Therefore, we use as larger α as possible to get better performance.

Boundary b : The influence of b for PNP-I_b is shown in Fig. 8 (b). $b = 4$ is the best on SOP while $b = 1$ is the best on VehicleID. b controls the biggest gradient of the PNP-I_b with respect to $R(i, S_N)$, and different datasets require different adjusting range. We report the best value for the comparison of different variants.

Samples Per Class $|P|$: The influence of the samples per class $|P|$ in a mini-batch is shown in Fig. 8 (c). We observe that $|P| = 4$ results in the highest R@1. The probable reason is that mini-batches are formed by sampling from each class, where a low value means a larger number of sampled classes and a higher probability of sampling hard-negative instances that violate the correct ranking order. Increasing the number of classes in the batch results in a better batch approximation of the true class distribution, allowing each training iteration to enforce a more optimally structured embedding space. However, too small value ($|P| = 2$) leads to worse performance because it can not give enough positive instances to violate the correct ranking order for training.

Evaluation of Robustness

In order to evaluate the robustness and the performance of multi-centers per class, we perform extensive experiments on constructed datasets in this section. Specifically, we use

Method	Small		Medium		Large	
	R@1	R@5	R@1	R@5	R@1	R@5
Smooth-AP	89.9	95.8	88.6	94.3	85.2	93.2
PNP-I _u	86.9	94.0	85.0	92.6	80.8	90.7
PNP-D _q	92.6	96.3	90.8	95.3	88.6	94.6

Table 7: Evaluation of robustness on VehicleID (%).




Figure 8: Impact of different hyper-parameters on SOP and VehicleID. (a) steepness α , (b) boundary b , (c) samples per class during mini batch sampling $|P|$.

the train set of SOP and VehicleID as benchmark datasets and merge every three categories to form one new class with multi-centers (at least three centers per class). As for evaluation, we still use original test set to evaluate the model. During the training, we use PNP-I_b, PNP-D_q and Smooth-AP as loss functions.

The results are shown in Table 6, 7. Compared with using original training set, these three methods all have performance loss on two benchmark datasets. Generally, PNP-D_q and Smooth-AP with similar gradient assignment strategy have less performance loss than PNP-I_b because of their adaptability to multi-centers, and PNP-D achieves the best performance. Specifically, for VehicleID, PNP-D exceeds the performance of Smooth-AP by 2 - 3% and PNP-I by a significant 6 - 8% on R@1.

The above results show better robustness of PNP-D. Noisy data is inevitable in the real-world data and it will leads to the instability of training. The above experiments show that PNP-D can achieve comparable performance, even if there are a lot of wrong annotations in datasets. This is because PNP-D can automatically remove the interference of noisy data by putting it in a single center.

Conclusions

In this paper, we propose a novel PNP loss, which directly improves the retrieval performance by penalizing negative instances before positive ones. Moreover, we find that different derivative functions of losses correspond to different gradient assignments. Therefore, we systematically investigate different gradient assignment solutions via constructing derivative functions of losses, resulting in PNP-I and PNP-D. PNP-D consistently achieve state-of-the-art performance on three benchmark datasets. We just provide four variants and will explore more forms via similar strategy in the future. In addition, we hope the proposed strategy for designing loss functions can also be applied to other tasks.

Acknowledgments

This work was supported in part by Beijing Natural Science Foundation under Grant L182054, in part by the National Natural Science Foundation of China under Grant 62125207, 61972378, 61532018, U1936203, U19B2040. This research was also supported by Meituan-Dianping Group.

References

- Albert, G.; Jon, A.; Jerome, R.; and Larlus, D. 2016. Deep Image Retrieval: Learning Global Representations for Image Search. In *European Conference on Computer Vision*, 241–257.
- Arandjelovic, R.; and Zisserman, A. 2013. All About VLAD. In *IEEE Conference on Computer Vision and Pattern Recognition*, 1578–1585.
- Arandjelović, R.; Gronat, P.; Torii, A.; Pajdla, T.; and Sivic, J. 2018. NetVLAD: CNN Architecture for Weakly Supervised Place Recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 40(6): 1437–1451.
- Babenko, A.; Slesarev, A.; Chigorin, A.; and Lempitsky, V. 2014. Neural Codes for Image Retrieval. In *European Conference on Computer Vision*, 584–599.
- Brattoli, B.; Roth, K.; Ommer, B.; and Bjorn, 2019. MIC: Mining Interclass Characteristics for Improved Metric Learning. In *IEEE International Conference on Computer Vision*, 7999–8008.
- Brown, A.; Xie, W.; Kalogeiton, V.; and Zisserman, A. 2020. Smooth-AP: Smoothing the Path Towards Large-Scale Image Retrieval. In *European Conference on Computer Vision*, 677–694.
- Cakir, F.; He, K.; Xia, X.; Kulis, B.; and Sclaroff, S. 2019. Deep Metric Learning to Rank. In *IEEE Conference on Computer Vision and Pattern Recognition*, 1861–1870.
- Cao, B.; Araujo, A.; and Sim, J. 2020. Unifying Deep Local and Global Features for Image Search. In *European Conference on Computer Vision*, 726–743.
- Engilberge, M.; Chevallier, L.; Perez, P.; and Cord, M. 2019. SoDeep: A Sorting Deep Net to Learn Ranking Loss Surrogates. In *IEEE Conference on Computer Vision and Pattern Recognition*, 10792–10801.
- Fu, Z.; Li, Y.; Mao, Z.; Wang, Q.; and Zhang, Y. 2021. Deep Metric Learning with Self-Supervised Ranking. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 1370–1378.
- Gu, G.; Ko, B.; and Kim, H.-G. 2021. Proxy Synthesis: Learning with Synthetic Classes for Deep Metric Learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 1460–1468.
- Guiver, J.; and Snelson, E. 2008. Learning to rank with SoftRank and Gaussian processes. In *Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, 259–266.
- Hadsell, R.; Chopra, S.; and Lecun, Y. 2006. Dimensionality Reduction by Learning an Invariant Mapping. In *IEEE Conference on Computer Vision and Pattern Recognition*, 1735 – 1742.
- He, K.; Cakir, F.; Bargal, S.; and Sclaroff, S. 2018. Hashing as Tie-Aware Learning to Rank. In *IEEE Conference on Computer Vision and Pattern Recognition*, 4023–4032.
- He, K.; Lu, Y.; and Sclaroff, S. 2018. Local Descriptors Optimized for Average Precision. In *IEEE Conference on Computer Vision and Pattern Recognition*, 596–605.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep Residual Learning for Image Recognition. In *IEEE Conference on Computer Vision and Pattern Recognition*, 770–778.
- Ji, X.; Wang, W.; Zhang, M.; and Yang, Y. 2017. Cross-Domain Image Retrieval with Attention Modeling. In *Proceedings of the ACM Multimedia Conference*, 1654–1662.
- Jégou, H.; Douze, M.; Schmid, C.; and Perez, P. 2010. Aggregating local descriptors into a compact image representation. In *IEEE Conference on Computer Vision and Pattern Recognition*, 3304 – 3311.
- Kemertas, M.; Pishdad, L.; Derpanis, K. G.; and Fazly, A. 2020. RankMI: A Mutual Information Maximizing Ranking Loss. In *IEEE Conference on Computer Vision and Pattern Recognition*, 14362–14371.
- Kim, Y.; and Park, W. 2021. Multi-level Distance Regularization for Deep Metric Learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 1827–1835.
- Kingma, D.; and Ba, J. 2015. Adam: A Method for Stochastic Optimization. In *International Conference on Learning Representations*, 1–15.
- Liu, H.; Tian, Y.; Wang, Y.; Pang, L.; and Huang, T. 2016. Deep Relative Distance Learning: Tell the Difference Between Similar Vehicles. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2167–2175.
- Manmatha, R.; Wu, C.-Y.; Smola, A.; and Krahenbuhl, P. 2017. Sampling Matters in Deep Embedding Learning. In *IEEE International Conference on Computer Vision*, 2859–2867.
- Milbich, T.; Roth, K.; Bharadhwaj, H.; Sinha, S.; Bengio, Y.; Ommer, B.; and Cohen, J. P. 2020a. DiVA: Diverse Visual Feature Aggregation for Deep Metric Learning. In *European Conference on Computer Vision*, 590–607.
- Milbich, T.; Roth, K.; Brattoli, B.; and Ommer, B. 2020b. Sharing Matters for Generalization in Deep Metric Learning. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, PP(99): 1–1.
- Min, W.; Mei, S.; Li, Z.; and Jiang, S. 2020. A Two-Stage Triplet Network Training Framework for Image Retrieval. *IEEE Transactions on Multimedia*, 22(12): 3128–3138.
- Mohapatra, P.; Rolinek, M.; Jawahar, C. V.; Kolmogorov, V.; and Kumar, M. P. 2018. Efficient Optimization for Rank-Based Loss Functions. In *IEEE Conference on Computer Vision and Pattern Recognition*, 3693–3701.
- Perronnin, F.; Liu, Y.; Sánchez, J.; and Poirier, H. 2010. Large-scale image retrieval with compressed Fisher vectors. In *IEEE Conference on Computer Vision and Pattern Recognition*, 3384–3391.
- Qian, Q.; Shang, L.; Sun, B.; Hu, J.; Tacoma, T.; Li, H.; and Jin, R. 2019. SoftTriple Loss: Deep Metric Learning

- Without Triplet Sampling. In *IEEE International Conference on Computer Vision*, 6449–6457.
- Radenović, F.; Tolias, G.; and Chum, O. 2016. CNN Image Retrieval Learns from BoW: Unsupervised Fine-Tuning with Hard Examples. In *European Conference on Computer Vision*, 3–20.
- Radenović, F.; Tolias, G.; and Chum, O. 2019. Fine-tuning CNN Image Retrieval with No Human Annotation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 41(7): 1655–1668.
- Revaud, J.; Almazan, J.; Rezende, R.; and De Souza, C. 2019. Learning With Average Precision: Training Image Retrieval With a Listwise Loss. In *IEEE International Conference on Computer Vision*, 5107–5116.
- Rolínek, M.; Musil, V.; Paulus, A.; Vlastelica, M.; Michaelis, C.; and Martius, G. 2020. Optimizing Rank-Based Metrics With Blackbox Differentiation. In *IEEE Conference on Computer Vision and Pattern Recognition*, 7617–7627.
- Roth, K.; Milbich, T.; Sinha, S.; Gupta, P.; Ommer, B.; and Cohen, J. P. 2020. Revisiting training strategies and generalization performance in deep metric learning. In *International Conference on Machine Learning*, 8242–8252.
- Sanakoyeu, A.; Tschernezki, V.; Büchler, U.; and Ommer, B. 2019. Divide and Conquer the Embedding Space for Metric Learning. In *IEEE Conference on Computer Vision and Pattern Recognition*, 471–480.
- Schroff, F.; Kalenichenko, D.; and Philbin, J. 2015. FaceNet: A unified embedding for face recognition and clustering. In *IEEE Conference on Computer Vision and Pattern Recognition*, 815–823.
- Schütze, H.; Manning, C. D.; and Raghavan, P. 2008. *Introduction to information retrieval*, volume 39. Cambridge University Press Cambridge.
- Sohn, K. 2016. Improved deep metric learning with multi-class n-pair loss objective. In *Advances in Neural Information Processing Systems*, 1857–1865.
- Song, H. O.; Xiang, Y.; Jegelka, S.; and Savarese, S. 2016. Deep metric learning via lifted structured feature embedding. In *IEEE Conference on Computer Vision and Pattern Recognition*, 4004–4012.
- Sun, Y.; Cheng, C.; Zhang, Y.; Zhang, C.; Zheng, L.; Wang, Z.; and Wei, Y. 2020. Circle Loss: A Unified Perspective of Pair Similarity Optimization. In *IEEE Conference on Computer Vision and Pattern Recognition*, 6397–6406.
- Teh, E. W.; DeVries, T.; and Taylor, G. W. 2020. ProxyNCA++: Revisiting and Revitalizing Proxy Neighborhood Component Analysis. In *European Conference on Computer Vision*, 448–464.
- Tolias, G.; Jeníček, T.; and Chum, O. 2020. Learning and Aggregating Deep Local Descriptors for Instance-Level Recognition. In *European Conference on Computer Vision*, 460–477.
- Ustinova, E.; and Lempitsky, V. 2016. Learning Deep Embeddings with Histogram Loss. In *Annual Conference on Neural Information Processing Systems*, 4170–4178.
- Van Horn, G.; Mac Aodha, O.; Song, Y.; Cui, Y.; Sun, C.; Shepard, A.; Adam, H.; Perona, P.; and Belongie, S. 2018. The inaturalist species classification and detection dataset. In *IEEE Conference on Computer Vision and Pattern Recognition*, 8769–8778.
- Wang, J.; Song, Y.; Leung, T.; Rosenberg, C.; Wang, J.; Philbin, J.; Chen, B.; and Wu, Y. 2014. Learning Fine-Grained Image Similarity with Deep Ranking. In *IEEE Conference on Computer Vision and Pattern Recognition*, 1386–1393.
- Wang, X.; Zhang, H.; Huang, W.; and Scott, M. R. 2020. Cross-Batch Memory for Embedding Learning. In *IEEE Conference on Computer Vision and Pattern Recognition*, 6387–6396.
- Yair, M.-A.; Toshev, A.; Leung, T.; Ioffe, S.; and Singh, S. 2017. No Fuss Distance Metric Learning Using Proxies. In *IEEE International Conference on Computer Vision*, 360–368.
- Zheng, W.; Wang, C.; Lu, J.; and Zhou, J. 2021. Deep Compositional Metric Learning. In *IEEE Conference on Computer Vision and Pattern Recognition*, 9320–9329.