

SUPERVISED METRIC LEARNING FOR RETRIEVAL VIA CONTEXTUAL SIMILARITY OPTIMIZATION

Anonymous authors

Paper under double-blind review

ABSTRACT

Existing deep metric learning approaches fall into three general categories: contrastive learning, average precision (AP) maximization, and classification. We propose a novel alternative approach, *contextual similarity optimization*, inspired by work in unsupervised metric learning. Contextual similarity is a discrete similarity measure based on relationships between neighborhood sets, and is widely used in the unsupervised setting as pseudo-supervision. Inspired by this success, we propose a framework which optimizes *a combination of contextual and cosine similarities*. Contextual similarity calculation involves several non-differentiable operations, including the heaviside function and intersection of sets. We show how to circumvent non-differentiability to explicitly optimize contextual similarity, and we further incorporate appropriate similarity regularization to yield our novel metric learning loss. The resulting loss function achieves state-of-the-art Recall @ 1 accuracy on standard supervised image retrieval benchmarks when combined with the standard contrastive loss.

1 INTRODUCTION

Supervised metric learning aims to learn a transformation from data to an embedding space, where similar samples are close together and dis-similar samples are far apart. Three types of metric learning approaches currently exist: contrastive learning, average precision (AP) maximization, and classification. Contrastive losses roughly follow the idea proposed by Hadsell et al. (2006): minimize the distance between positive pairs and maximize the distance between negative pairs. Alternatively, approaches such as ProxyNCA (Movshovitz-Attias et al. (2017)) use a classification loss to optimize a linear layer on top of the embedding. The linear layer and the embedding are trained jointly, and the former is discarded. More recently, there is strong interest in average precision (AP) maximization. AP is a metric which takes into account the rank of samples relative to each other and is non-differentiable. Several studies tackle this non-differentiability issue: Fast-AP (Cakir et al. (2019)) uses a soft-binning approach; Smooth-AP (Brown et al. (2020)) and ROADMAP (Ramzi et al. (2021)) use sigmoid-like functions to approximate the heaviside function used for ranking. In the current work, we take a different approach to metric learning, using the notion of contextual similarity from unsupervised metric learning (Zhong et al. (2017) and Kim et al. (2022)).

Unsupervised metric learning aims to learn a good representation without using labels. Most existing work in this area focus on a semi-supervised pseudo-labeling approach. In particular, a teacher model is used to guess the true similarity between unlabeled samples; the guess is then used as artificial supervision for training a student model. Instead of directly using the cosine similarity between teacher features as the pseudo-similarity label, most state-of-the-art methods rely on a weighted combination of the cosine similarity and contextual similarity (Ge et al. (2020a), Ge et al. (2020b), Kim et al. (2022)). Inspired by this success, we endeavor to explicitly optimize contextual similarity.

The definition of contextual similarity varies. In the person re-ID literature, most papers use the k -reciprocal re-rank distance proposed in Zhong et al. (2017). Contextual similarity is based on the following intuition: (1) neighbors are likely to have the same label regardless of absolute similarity, (2) two samples are likely to have the same label if they share the same neighbors, (3) reciprocal relationships are important: sample i is likely to be similar to sample j if i is a neighbor of j and j is a neighbor of i , (4) neighbors of neighbors can be treated as neighbors. These observations lead directly to a rigorous definition of contextual similarity presented in Section 3.1. Existing supervised

metric learning approaches already implicitly minimize some notion of contextual similarity – this is why contextual similarity works as a pseudo-label. We hypothesize that *explicitly* optimizing contextual similarity boosts recall performance when combined with standard contrastive learning.

We use the term “context” to refer to the set of neighbors around a sample. Roughly speaking, the contextual similarity between samples i and j is a fraction between 0 and 1 indicating the degree of overlap between the contexts of i and j . We seek to minimize the contextual similarity between samples with different labels, and vice versa. Clearly, this requires the contextual similarity to be differentiable with respect to the embeddings. Contextual similarity is non-differentiable because it involves manipulation of indicator functions, e.g. to calculate the intersection between contexts. We find that existing tools to circumvent non-differentiability, such as using sigmoid approximations, are insufficient for our problem. Therefore, we develop a novel optimization approach for contextual similarity in Section 3.2.

We make two technical and one experimental contribution to the metric learning literature:

1. We derive a highly non-trivial differentiable contextual loss function.
2. We propose a simple but novel similarity regularizer which significantly improves the standard contrastive loss and may be of independent interest.
3. Our framework improves the Recall @ 1 performance significantly over the current state-of-the-art across four diverse image retrieval benchmarks.

2 LITERATURE REVIEW

Supervised Metric Learning There are three main approaches: (1) contrastive (2) classification and (3) AP maximization. The contrastive loss was first proposed by Hadsell et al. (2006). The triplet loss (Weinberger et al. (2005)) operates on triplets (a pair of positive samples with a negative sample) instead of pairs and is competitive with the contrastive loss when combined with an effective triplet sampling strategy (Wu et al. (2017)). While contrastive and triplet losses are well-known, they remain the go-to method for metric learning and Musgrave et al. (2020a) show that they are comparable in performance to many recent methods. Variations of the contrastive loss include the lifted structure loss (Oh Song et al. (2016)) and multi-similarity loss (MS) (Wang et al. (2019)). Standard classification methods can achieve state-of-the-art metric learning results when appropriate tricks are used (Boudiaf et al. (2020), Zhai & Wu (2018), Teh et al. (2020)). AP maximization methods learn to rank samples within a mini-batch correctly. The AP objective is non-differentiable because of the heaviside used for ranking. As a workaround, Fast-AP uses soft-binning; Smooth-AP uses a low-temperature sigmoid; ROADMAP uses an upper bound on the heaviside instead of an approximation. Our work is distinct from all three categories while having similarities to learning to rank (see Figure 2).

Unsupervised Metric Learning There is extensive interest in unsupervised and domain adaptive metric learning, mainly in the context of person re-ID (see survey Ye et al. (2021)). Contextual similarity is a cornerstone for most recent work in this area. Most unsupervised person re-ID papers use the k -reciprocal re-rank distance, which is a weighted combination of Euclidean distance and Jaccard distance between reciprocal-neighbor sets. Their formulation has three disadvantages: (1) it operates on the entire dataset instead of a batch; (2) it involves a series of set expansion operations which are hard to vectorize, much less make differentiable; (3) it is optimized for the person re-ID application. More recently, STML (Kim et al. (2022)) propose an unsupervised metric learning framework for image retrieval which uses a simpler batch-wise contextual similarity measure. We mostly follow STML’s contextual similarity definition, with some minor tweaks. We emphasize that all work cited above use contextual similarity as a pseudo-label for unsupervised learning, while our work optimizes the contextual similarity in the supervised setting.

3 METHOD

This section is organized as follows: (1) We state the mathematical definition of contextual similarity, which loosely follows Kim et al. (2022), then (2) we explain how to optimize the objective.

Notation Denote the normalized output of the embedding network as $f_i \in \mathbb{R}^d$. $s_{ij} \in [-1, 1]$ denotes the cosine similarity between the samples i and j ($s_{ij} = \langle f_i, f_j \rangle$). $y_{ij} \in \{0, 1\}$ denotes the true similarity between i and j . $y_{ij} = 1$ if samples i and j share the same label; otherwise, $y_{ij} = 0$. We use uppercase letters to denote matrices, math script to denote sets, and lowercase letters to denote scalars. i, j and p are reserved for sample indices.

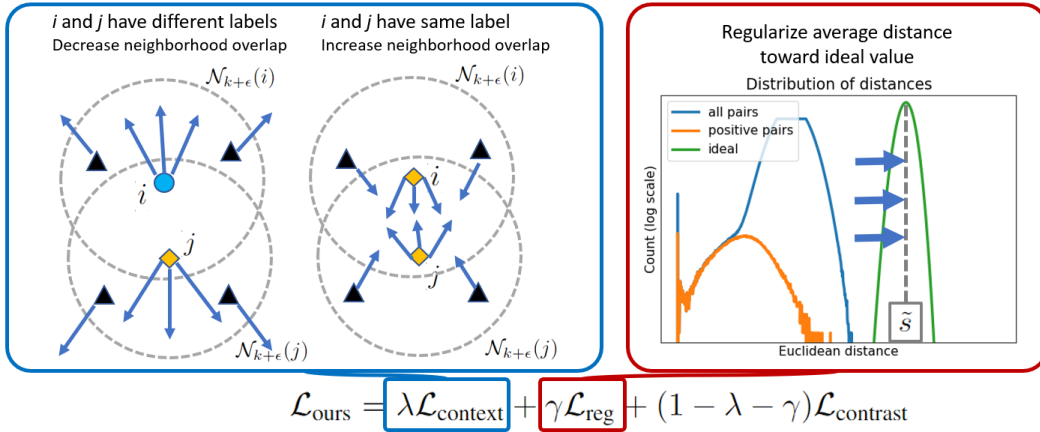


Figure 1: Method Overview. Our final framework consists of three loss terms. $\mathcal{L}_{\text{context}}$ pulls apart the contexts of dis-similar samples and pushes together the contexts of similar samples. This is illustrated on the left, where black triangles indicate samples with any label. See also analysis in Section 4.2. \mathcal{L}_{reg} is a regularizer that encourages utilization of the entire embedding space by regularizing average similarities to a fixed value. $\mathcal{L}_{\text{contrast}}$ is the standard contrastive loss.

3.1 CONTEXTUAL SIMILARITY DEFINITION

We compute the contextual similarity on a single batch, not the entire dataset. Two samples are contextually similar if they share the same neighbors, i.e. the intersection of their k -neighbor sets is large. Following this intuition, we calculate \tilde{w}_{ij} , the preliminary contextual similarity between samples i and j :

$$\mathcal{N}_{k+\epsilon}(i) = \{j \mid s_{ij} \leq s_{ip} + \epsilon \text{ where } p \text{ denotes the } k\text{-th closest neighbor of } i\} \quad (1)$$

$$\tilde{w}_{ij} = \begin{cases} |\mathcal{N}_{k+\epsilon}(i) \cap \mathcal{N}_{k+\epsilon}(j)| / |\mathcal{N}_{k+\epsilon}(i)| & , \text{ if } j \in \mathcal{N}_{k+\epsilon}(i) \\ 0 & , \text{ otherwise.} \end{cases} \quad (2)$$

We include i in $\mathcal{N}_{k+\epsilon}(i)$ (so if $k = 2$, then $\mathcal{N}_{k+0}(i)$ includes two elements: i and its closest neighbor). Then, we use query expansion and symmetrize to obtain the final contextual similarity. Query expansion (Arandjelović & Zisserman (2012)) is a standard evaluation-time trick in metric learning. It boosts retrieval performance by retrieving neighbors of a sample’s neighbors. Analogously, we adjust the contextual similarity by averaging \tilde{w}_{ij} over the set of close reciprocal neighbors $\mathcal{R}_{k/2}(i)$.

$$\mathcal{N}_{k/2+\epsilon}(i) = \{j \mid s_{ij} \leq s_{ip} + \epsilon \text{ where } p \text{ denotes the } k/2\text{-th closest neighbor of } i\} \quad (3)$$

$$\mathcal{R}_{k/2+\epsilon}(i) = \{j \mid j \in \mathcal{N}_{k/2+\epsilon}(i) \text{ and } i \in \mathcal{N}_{k/2+\epsilon}(j)\} \quad (4)$$

$$\hat{w}_{ij} = \frac{1}{|\mathcal{R}_{k/2+\epsilon}(i)|} \sum_{p \in \mathcal{R}_{k/2+\epsilon}(i)} \tilde{w}_{pj} \quad , \quad w_{ij} = \frac{1}{2}(\hat{w}_{ij} + \hat{w}_{ji}). \quad (5)$$

Note that $w_{ij} \in [0, 1]$ and only depend on the cosine similarities s_{ij} between embeddings. We want to optimize the embeddings such that w_{ij} converges to y_{ij} . The value of k is not arbitrary; it must be set to the number of samples per label in the mini-batch. For example, a standard metric learning setup is to randomly sample 32 labels and then sample 4 images per label. In this case, we set $k = 4$.

3.2 OPTIMIZATION

The definition in the previous sub-section clearly contains three discrete operations: (1) greater-than, (2) logical-and, and (3) intersection. For optimization, we will deal extensively with indicator matrices: we denote as $\mathbb{1}_{\mathcal{N}} \in \mathbb{R}^{n \times n}$ the indicator matrix where $\mathbb{1}_{\mathcal{N}}(i, j) = 1$ if $j \in \mathcal{N}(i)$ for some set \mathcal{N} . We use \odot to denote element-wise multiplication.

Greater-than This is used in Eq. 1 and 3 and is equivalent to the non-differentiable heaviside function θ . Our approach is to use the exact value of $\theta(\cdot)$ in the forward pass and a constant positive gradient α in the backward pass. A constant positive gradient is reasonable since θ is a (non-strictly) increasing function.




Figure 2: Illustration of $\mathcal{N}_{k+\epsilon}$ optimization using the simplified loss function \mathcal{L}_1 in Eq. 18. Negative samples which are in the set $\mathcal{N}_{k+\epsilon}(i)$ are pushed away from i by a constant gradient. Positive samples which are not in the set $\mathcal{N}_{k+\epsilon}(i)$ are pulled closer. This behavior is analogous to the standard contrastive loss: $\mathcal{L}_{\text{contrast}}$ applies a constant gradient to pairs violating a fixed margin, while \mathcal{L}_1 applies a constant gradient to pairs violating the flexible margin defined by the k -th neighbor. $k = 4$ in the illustration.

$$\text{Forward: } \theta(x) = 1 \text{ if } x \geq 0 ; 0 \text{ otherwise} \quad \text{Backward: } \frac{\partial \theta(x)}{\partial x} = \alpha. \quad (6)$$

In Section 4.1 we show with a toy experiment that this approach is robust despite being somewhat heuristic. In contrast, θ is traditionally approximated by a sigmoid (e.g. for AP approximation and in Gated Recurrent Networks Cho et al. (2014)):

$$\theta_\sigma(x) = \frac{1}{1 + \exp(x/\tau)}. \quad (7)$$

θ_σ trades off the quality of the approximation with the domain where gradients are non-zero. As temperature τ decreases, θ_σ approaches θ , but gradient vanishes everywhere except in a small region around the boundary. This behavior is not intuitive: it is undesirable to only have a large gradient at the boundary. Some prior work (e.g. ROADMAP) side-step this issue by using an upper-bound to the heaviside function (where the right side of the heaviside function increases linearly), which solves the gradient issue at the expense of grossly over-estimating the true objective. In our case, this is especially concerning since we will be multiplying together indicator functions. In Section 5.1, we show that our approach in Eq. 6 achieves better empirical results than a sigmoid approximation.

Logical-and A logical-and is used explicitly in Eq. 4 and implicitly in Eq. 2 and 5. There are two differentiable substitutes for logical-and: min and multiplication. Multiplication is smooth, but has no gradient at the origin. min is logically consistent on the continuous domain $[0,1]$, but the gradient is not continuous when inputs are equal. We found experimentally that multiplication is the best option, see row 5 of Table 3. A gradient of zero at the origin is desirable in the case of Eq. 4 and 5 (query expansion step). A work-around for the zero-gradient issue for Eq. 2 is presented in the next sub-section.

Intersection The number of elements in the intersection in Eq. 2 is calculated using matrix multiplication. A sample $p \in \mathcal{N}_{k+\epsilon}(i) \cap \mathcal{N}_{k+\epsilon}(j)$ if $p \in \mathcal{N}_{k+\epsilon}(i)$ and $p \in \mathcal{N}_{k+\epsilon}(j)$. Using multiplication for the logical-and, the number of elements in the intersection can be expressed compactly as a function of the indicator matrices:

$$M_+ = \mathbb{1}_{\mathcal{N}_{k+\epsilon}} \mathbb{1}_{\mathcal{N}_{k+\epsilon}}^T, \text{ where } M_+(i, j) = |\mathcal{N}_{k+\epsilon}(i) \cap \mathcal{N}_{k+\epsilon}(j)|. \quad (8)$$

We are now ready to state the loss function by combining the definition of contextual similarity with the optimization method.

3.3 LOSS

Our loss function straightforwardly combines the previous two sub-sections, with one exception for the intersection calculation in Eq. 2. Optimizing the intersection with M_+ tends to focus on pairs of neighboring samples because using multiplication for logical-and zeros out the gradient when both inputs are 0. This is problematic since the resulting embedding becomes clustered regardless of true similarity (see Figure 4). We mitigate this problem by optimizing the intersection of the complements $\mathcal{N}_{k+\epsilon}^c(i) \cap \mathcal{N}_{k+\epsilon}^c(j)$ in addition to optimizing the original intersection in Eq. 2. c denotes




Figure 3: This figure justifies the non-standard approach of optimizing \mathcal{L}_1 (Eq. 18) by overriding the gradient of θ . We generate synthetic 2-D data consisting of five concentric circles; each color represents a label. We use a standard three layer MLP with ReLU activations to output a 2D embedding. We train on increasingly noisy data and plot the resulting embeddings on clean data. We only sample 4 points per label per batch ($k = 4$). \mathcal{L}_1 minimization results in reasonable embeddings despite its simplicity.

the complement of a set. Maximizing the size of the intersection between two sets is equivalent to maximizing the size of the intersection between their complements, and vice versa. This can be trivially proven under the assumptions that the universal set is constant and that the size of both sets is constant.

Let $D \in \mathcal{R}^{n \times n}$ denote the matrix where $D(i, j) \in [0, 4]$ is the squared Euclidean distance between the normalized features of sample i and j . $D(i, j) = 2 - 2s_{ij}$. sg denotes stop gradient.

Step 1 (Neighborhood Optimization):

$$\mathbb{1}_{\mathcal{N}_{k+\epsilon}}(i, j) = \theta(-D(i, j) + \text{sg}(D(i, p)) + \epsilon) \text{ where } p \text{ denotes the } k\text{-th closest neighbor of } i. \quad (9)$$

Step 2 (Optimizing Intersection of Neighborhoods):

$$\begin{aligned} \tilde{W} &= \frac{1}{2} \left(\frac{M_+}{\text{sg}(|\mathcal{N}_{k+\epsilon}(i)|)} + \frac{M_-}{\text{sg}(|\mathcal{N}_{k+\epsilon}^c(i)|)} \right) \odot \mathbb{1}_{\mathcal{N}_{k+\epsilon}} \\ \text{where } M_+ &= \mathbb{1}_{\mathcal{N}_{k+\epsilon}} \mathbb{1}_{\mathcal{N}_{k+\epsilon}}^\top \text{ and } M_- = \mathbb{1}_{\mathcal{N}_{k+\epsilon}^c} \mathbb{1}_{\mathcal{N}_{k+\epsilon}^c}^\top. \end{aligned} \quad (10)$$

Step 3 (Query Expansion):

$$\mathbb{1}_{\mathcal{R}_{k/2+\epsilon}}(i, j) = \theta(-D(i, j) + \text{sg}(D(i, p)) + \epsilon) \text{ where } p \text{ denotes the } k/2\text{-th closest neighbor of } i \quad (11)$$

$$\mathbb{1}_{\mathcal{R}_{k/2+\epsilon}} = \mathbb{1}_{\mathcal{N}_{k/2+\epsilon}} \odot \mathbb{1}_{\mathcal{N}_{k/2+\epsilon}}^\top \quad (12)$$

$$\hat{W} = \frac{\mathbb{1}_{\mathcal{R}_{k/2+\epsilon}} \tilde{W}}{|\mathcal{R}_{k/2+\epsilon}(i)|}, \quad W = \frac{1}{2} (\hat{W} + \hat{W}^\top). \quad (13)$$

Finally, we use the MSE loss to optimize $w_{ij} := W(i, j)$ against the true similarity labels y_{ij} :

$$\mathcal{L}_{\text{context}} = \frac{1}{n^2} \sum_{i,j}^{n^2} (y_{ij} - w_{ij})^2. \quad (14)$$

$\mathcal{L}_{\text{context}}$ is the key to our framework and works reasonably on its own. However, it has two vulnerabilities: (1) similar to learning to rank losses, $\mathcal{L}_{\text{context}}$ can be small even if the distance between positive pairs is large, so long as there are no closer negative pairs (see Figure 2); (2) $\mathcal{L}_{\text{context}}$ tends to converge to a solution where the average cosine similarity between all pairs is large, suggesting that only a small portion of the available embedding space is utilized. In response to problem (1), we add the standard contrastive loss (Hadsell et al. (2006)), which explicitly optimizes the cosine similarity toward fixed margins δ_+ and δ_- . In response to problem (2), we add a non-standard but straightforward similarity regularizer, which regularizes the average cosine similarity between all

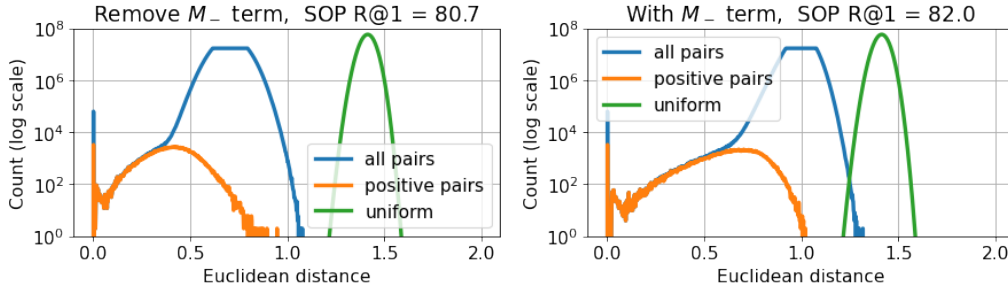


Figure 4: This figure justifies the M_- term in Eq. 10. The left plot shows the distribution of distances between pairs in normalized embedding space when we only optimize M_+ ; the right plot shows the same when we optimize both M_- and M_+ . A uniformly distributed embedding space has an average pair-wise distance of $\sqrt{2}$ because most directions are orthogonal in high dimensions. This is indicated by the green distribution. Clearly, more of the embedding space is utilized when we optimize both M_- and M_+ .

pairs toward a fixed value \tilde{s} . Our final framework (Eq. 17) minimizes a weighted combination of the three losses.

$$\mathcal{L}_{\text{contrast}} = \frac{\sum_{i,j|y_{ij}=1}(\delta_+ - s_{ij})_+}{|\{i,j|y_{ij}=1 \text{ and } \delta_+ - s_{ij} > 0\}|} + \frac{\sum_{i,j|y_{ij}=0}(s_{ij} - \delta_-)_+}{|\{i,j|y_{ij}=0 \text{ and } s_{ij} - \delta_- > 0\}|} \quad (15)$$

$$\mathcal{L}_{\text{reg}} = \left(\tilde{s} - \frac{1}{n^2} \sum_{i,j} s_{ij} \right)^2 \quad (16)$$

$$\mathcal{L}_{\text{ours}} = \lambda \mathcal{L}_{\text{context}} + \gamma \mathcal{L}_{\text{reg}} + (1 - \lambda - \gamma) \mathcal{L}_{\text{contrast}}. \quad (17)$$

Explanation of Hyperparameters α controls the magnitude of the heaviside gradient. ϵ is a margin-like parameter for the neighborhood set. δ_+ and δ_- are the positive and negative margins resp. for the contrastive loss. \tilde{s} is the desired average cosine similarity between all pairs. λ and γ control the relative weighting between the three losses.

4 ANALYSIS

The contextual similarity loss function $\mathcal{L}_{\text{context}}$ is highly non-trivial. In this section we demystify each of the three steps using a toy experiment and gradient analysis.

4.1 STEP 1: NEIGHBORHOOD OPTIMIZATION

Consider the following simplified version of $\mathcal{L}_{\text{context}}$:

$$\mathcal{L}_1 = \mathcal{L}_{\text{MSE}}(y_{ij}, \mathbb{1}_{\mathcal{N}_{k+\epsilon}}(i, j)). \quad (18)$$

This is a valid loss function that works reasonably on its own (see Figure 3). During each iteration, we sample a batch with exactly k samples from each label. Intuitively, \mathcal{L}_1 reaches a minimum of zero when $\mathcal{N}_{k+\epsilon}(i)$ includes only i and the $k-1$ other samples with the same label, i.e. all samples are correctly ranked. Using the gradient of $\theta(\cdot)$ defined in Eq. 6, we see that negative samples which are in $\mathcal{N}_{k+\epsilon}(i)$ receive a gradient of magnitude α pushing it away from i , while positive samples which are not in $\mathcal{N}_{k+\epsilon}(i)$ receive a gradient of magnitude α pushing it toward i . See Figure 2.

4.2 STEP 2: OPTIMIZING INTERSECTION OF NEIGHBORHOODS

While $\lambda \mathcal{L}_1 + (1 - \lambda) \mathcal{L}_{\text{contrastive}}$ is a reasonable loss function already, row 1 of Table 3 shows that \mathcal{L}_1 collapses without the contrastive loss ($\lambda = 1$). We propose that this is because \mathcal{L}_1 provides very sparse gradients: $\partial \mathcal{L}_1 / \partial s_{ij} = 0$ for most (i, j) pairs. Consider the following more complicated loss:

$$\mathcal{L}_2 = \mathcal{L}_{\text{MSE}}(y_{ij}, \tilde{W}(i, j)). \quad (19)$$

\tilde{W} is defined in Eq. 10. $\tilde{W}(i, j)$ has two terms multiplied together element-wise. The $\mathbb{1}_{\mathcal{N}_{k+\epsilon}}$ term was already addressed in the previous sub-section, so we focus on the first term, which optimizes the intersection between neighborhood sets M_+ . We offer a straight-forward intuition: *Maximizing*

the intersection between the neighborhood sets of two samples is equivalent to pushing one sample towards the context of the other sample and vice versa; minimizing the intersection is equivalent to pulling apart the contexts of the two samples. We show this intuition by analyzing the gradient.

For simplicity, consider $\epsilon = 0$, such that the normalization factors in Eq. 10 are equal to k . Further consider a negative pair of samples i and j ($y_{ij} = 0$), where j is wrongly ranked w.r.t. i (i.e. $j \in \mathcal{N}_k(i)$). Following the loss function equations:

$$\begin{aligned} M_+(i, j) &= \langle \mathbb{1}_{\mathcal{N}_k}(i), \mathbb{1}_{\mathcal{N}_k}(j) \rangle & M_-(i, j) &= \langle \mathbb{1}_{\mathcal{N}_k^c}(i), \mathbb{1}_{\mathcal{N}_k^c}(j) \rangle \\ \tilde{W}(i, j) &= \frac{1}{2} \left(\frac{1}{k} M_+(i, j) + \frac{1}{k} M_-(i, j) \right) \cdot \underbrace{\mathbb{1}_{\mathcal{N}_k}(i, j)}_{=1}. \end{aligned} \quad (20)$$

$\tilde{W}(i, j) \in (0, 1]$ must be a non-zero positive number. Using the equation for \mathcal{L}_2 , we see that the gradient w.r.t. $\tilde{W}(i, j)$ must be positive because $y_{ij} = 0$. We are now ready for the backward pass. For simplicity of notation, $\partial \tilde{W}(i, j) := g_{ij} > 0$ denotes the gradient of the loss w.r.t. $\tilde{W}(i, j)$.

$$\begin{aligned} \partial M_+(i, j) &= \partial M_-(i, j) = \frac{1}{2k} g_{ij} \\ \partial \mathbb{1}_{\mathcal{N}_k}(i) &= \frac{1}{2k} g_{ij} \mathbb{1}_{\mathcal{N}_k}(j), \quad \partial \mathbb{1}_{\mathcal{N}_k}(j) = \frac{1}{2k} g_{ij} \mathbb{1}_{\mathcal{N}_k}(i) \\ \partial \mathbb{1}_{\mathcal{N}_k^c}(i) &= \frac{1}{2k} g_{ij} \mathbb{1}_{\mathcal{N}_k^c}(j), \quad \partial \mathbb{1}_{\mathcal{N}_k^c}(j) = \frac{1}{2k} g_{ij} \mathbb{1}_{\mathcal{N}_k^c}(i). \end{aligned} \quad (21)$$

Use the fact that $\mathbb{1}_{\mathcal{N}_k^c} = 1 - \mathbb{1}_{\mathcal{N}_k}$ and $\partial \mathbb{1}_{\mathcal{N}_k} = -\partial \mathbb{1}_{\mathcal{N}_k^c}$ and vice versa:

$$\begin{cases} \partial \mathbb{1}_{\mathcal{N}_k}(i) &= \frac{g_{ij}}{2k} (\mathbb{1}_{\mathcal{N}_k}(j) - \mathbb{1}_{\mathcal{N}_k^c}(j)) = \frac{g_{ij}}{2k} (2\mathbb{1}_{\mathcal{N}_k}(j) - 1) \\ \partial \mathbb{1}_{\mathcal{N}_k}(j) &= \frac{g_{ij}}{2k} (\mathbb{1}_{\mathcal{N}_k}(i) - \mathbb{1}_{\mathcal{N}_k^c}(i)) = \frac{g_{ij}}{2k} (2\mathbb{1}_{\mathcal{N}_k}(i) - 1). \end{cases} \quad (22)$$

Going backward through $\theta(\cdot)$ multiplies the gradient by $-\alpha$ (negative sign accounts for optimizing distance instead of similarity):

$$\partial D(i) = -\frac{\alpha g_{ij}}{2k} (2\mathbb{1}_{\mathcal{N}_k}(j) - 1) \text{ and vice versa for } \partial D(j). \quad (23)$$

From the above equation, $\partial D(i, p) < 0$ when $\mathbb{1}_{\mathcal{N}_k}(j, p) = 1$ and $\partial D(i, p) > 0$ when $\mathbb{1}_{\mathcal{N}_k}(j, p) = 0$, for all samples p in the batch. In words, we increase the distance between i and all samples in the neighborhood of j , and we decrease the distance between i and all points outside the neighborhood of j . Recall that we assumed i and j to be a negative pair, so it makes sense to pull i away from the context of j in this fashion.

4.3 STEP 3: QUERY EXPANSION

Query expansion (QE) is an established trick for image retrieval Arandjelović & Zisserman (2012). QE expands the neighborhood set by additionally retrieving the neighbors of very close neighbors. In the unsupervised setting, QE refers to averaging the contextual similarity scores for samples in the $\mathcal{R}_{k/2+\epsilon}$ neighborhood (see Eq. 13); this leads to more robust pseudo-supervision. In our setting, we have two reasons for using QE (Eq. 11 12 and 13): (1) averaging \tilde{w}_{ij} with very close neighbors could have the same effect as label smoothing and (2) the $\mathcal{N}_{k/2+\epsilon}$ neighborhood is optimized by Eq. 12 and 13. Empirically, QE is necessary to achieve the optimal performance of our framework, since $\mathcal{L}_{\text{context}}$ achieves higher R@1 than \mathcal{L}_2 in Table 3.

5 EXPERIMENTS

Datasets We experiment on two small-scale and two large-scale datasets: Caltech-UCSD Birds (CUB-200) (Wah et al. (2011)), Stanford Cars-196 (Krause et al. (2013)), Stanford Online Products (SOP) (Oh Song et al. (2016)), and mini-iNaturalist-2021 (Van Horn et al. (2018)). CUB200 and Cars196 are smaller fine-grain classification datasets with 200 and 196 unique labels, respectively. SOP is a large-scale dataset with 120,053 product images from 22,634 classes. mini-iNaturalist-2021 is a subset of the iNaturalist 2021 species classification competition dataset, with 50 images each from 10,000 species. We are the first metric learning study to use this version of the iNaturalist dataset, so we randomly withhold 5,000 species for testing and use the other 5,000 for training. We will release our train-test split for reproducibility.

Table 1: R@k Results. We use ResNet-50 with an embedding size of 512 for all experiments. † indicates results reported by the original authors; we re-run all other baselines using the implementation by Musgrave et al. (2020b). Standard deviations are based on three trials with the same train-test split. We emphasize that our R@1 performance is at least two standard deviations better than the next best baseline on all datasets. We are at least comparable to baselines on other R@k metrics.

Method	CUB				Cars			
	R@1	R@2	R@4	R@8	R@1	R@2	R@4	R@8
Contrastive	68.5 ± 0.3	78.3 ± 0.1	86.0 ± 0.2	91.3 ± 0.1	85.4 ± 0.2	91.1 ± 0.3	94.6 ± 0.3	96.8 ± 0.1
Triplet	67.3 ± 0.2	77.9 ± 0.1	85.6 ± 0.2	91.2 ± 0.1	77.6 ± 1.3	85.4 ± 0.8	90.8 ± 0.7	94.1 ± 0.4
NtXent	65.7 ± 0.4	76.3 ± 0.2	84.3 ± 0.4	90.0 ± 0.4	79.0 ± 0.6	86.0 ± 0.3	91.0 ± 0.2	94.4 ± 0.3
MS	68.9 ± 0.5	78.5 ± 0.4	86.0 ± 0.6	91.4 ± 0.5	88.7 ± 0.4	93.0 ± 0.2	95.7 ± 0.1	97.3 ± 0.1
N-Softmax†	61.3	73.9	83.5	90.0	84.2	90.4	94.4	96.9
ProxyNCA++†	69.0 ± 0.8	79.8 ± 0.7	87.3 ± 0.7	92.7 ± 0.4	86.5 ± 0.4	92.5 ± 0.3	95.7 ± 0.2	97.7 ± 0.1
Fast-AP	63.3 ± 0.1	73.7 ± 0.4	82.2 ± 0.3	88.5 ± 0.2	81.1 ± 0.2	87.8 ± 0.4	92.2 ± 0.3	95.1 ± 0.3
Smooth-AP	66.5 ± 0.9	76.6 ± 0.5	84.8 ± 0.6	90.8 ± 0.4	74.7 ± 0.4	82.5 ± 0.7	88.0 ± 0.6	92.2 ± 0.2
ROADMAP	68.7 ± 0.5	78.3 ± 0.3	86.1 ± 0.3	91.1 ± 0.1	84.5 ± 0.5	90.3 ± 0.0	93.9 ± 0.0	96.2 ± 0.1
Ours	69.8 ± 0.2	79.8 ± 0.1	87.1 ± 0.1	92.3 ± 0.2	89.3 ± 0.0	93.7 ± 0.2	96.3 ± 0.1	97.8 ± 0.2

Method	SOP			mini-iNaturalist			
	R@1	R@10	R@100	R@1	R@4	R@16	R@32
Contrastive	82.4 ± 0.0	91.9 ± 0.0	96.0 ± 0.0	43.5 ± 0.1	62.7 ± 0.1	77.6 ± 0.1	83.2 ± 0.1
Triplet	82.0 ± 0.0	92.5 ± 0.1	96.7 ± 0.0	35.4 ± 0.1	56.5 ± 0.1	74.7 ± 0.1	81.7 ± 0.1
NtXent	79.7 ± 0.2	90.8 ± 0.0	96.1 ± 0.0	40.8 ± 0.1	61.6 ± 0.1	78.0 ± 0.0	83.9 ± 0.0
MS	81.4 ± 0.0	91.4 ± 0.0	96.1 ± 0.1	44.9 ± 0.1	63.9 ± 0.1	78.4 ± 0.1	83.9 ± 0.1
N-Softmax†	78.2	90.6	96.2	–	–	–	–
ProxyNCA++†	80.7 ± 0.5	92.0 ± 0.3	96.7 ± 0.1	–	–	–	–
Fast-AP	80.3 ± 0.1	91.0 ± 0.1	96.0 ± 0.0	35.6 ± 0.2	55.8 ± 0.1	72.8 ± 0.0	79.3 ± 0.0
Smooth-AP	82.0 ± 0.0	92.6 ± 0.0	96.9 ± 0.0	42.7 ± 0.0	63.3 ± 0.0	79.0 ± 0.0	84.7 ± 0.0
ROADMAP	83.1 ± 0.1	92.6 ± 0.0	96.6 ± 0.0	45.9 ± 0.1	65.8 ± 0.0	80.4 ± 0.1	85.7 ± 0.0
Ours	83.3 ± 0.0	92.9 ± 0.1	96.7 ± 0.0	46.2 ± 0.0	65.8 ± 0.1	80.2 ± 0.1	85.4 ± 0.1

Baselines We choose a representative set of baselines for comparison in Table 1. Refer to Sections 1 and 2 for citations. Contrastive and triplet losses are the accepted standard in the field. Multi-similarity (MS) is a popular pair-wise similarity loss. From classification-inspired methods, we compare against Normalized-Softmax (N-Softmax) (Zhai & Wu (2018)) and ProxyNCA++ (Teh et al. (2020)). From AP maximization methods, we compare against Fast-AP, Smooth-AP, and ROADMAP. For fair comparison, we re-run baselines using the same setup, with the same learning rates and schedules. We use the published optimal values for any internal hyperparameters. We were unable to match the published performances for ProxyNCA++ and N-Softmax, so we use the published R@k values for these two baselines.

Hyperparameters and Setup We use hyperparameter values that are approximately optimal across all datasets. $\lambda = 0.4$, $\gamma = 0.1$, $\alpha = 10.0$, $\epsilon = 0.05$, $k = 4$, $\delta_+ = 0.75$, $\delta_- = 0.6$, $\bar{s} = 0.25$. We follow the same training procedure as ROADMAP, but with a slightly faster learning-rate schedule for time efficiency. We use Adam with a learning-rate schedule that multiplies the learning rate by 0.3 at Epochs 15, 30, and 45. We train for 80 epochs. We report results on the model with the best test R@1 metric, as is standard in the literature. We use an initial learning rate of 8×10^{-5} on CUB, 0.00016 on Cars, 4×10^{-5} on SOP, and 8×10^{-5} on iNaturalist. We use a batch size of 256 for iNaturalist and 128 for other datasets; the larger batch size is necessary to achieve reasonable performance on iNaturalist. We use a 4 per class sampler. For CUB and Cars, we use random sampling (sample 32 classes at random, then sample 4 images per class). For SOP and iNaturalist, we use hierarchical sampling (Cakir et al. (2019)), following prior work. We use an embedding size of 512 and ResNet-50 with a linear embedding layer. Following prior work, we use layer-norm and max-pooling on the smaller CUB and Cars datasets. We always freeze batch-norm.

Performance Metrics We report Recall @ k for select k in Table 1. R@k is the percentage of test samples where at least one of the k closest neighbors have the same label. We report the average and standard deviations across three trials with the same train-test split.

Discussion Our R@1 results are at least two standard deviations better than the best baseline across all datasets. Our results are especially good on CUB and Cars, where we achieve R@1 gains of 0.8%

Table 2: Ablation Results 1. We experiment with removing one or two of the three loss terms. Overall, all three losses are necessary to achieve the best R@1 results. Reasonable results are also achieved with just \mathcal{L}_{reg} and $\mathcal{L}_{\text{contrast}}$. We emphasize that both $\mathcal{L}_{\text{context}}$ and \mathcal{L}_{reg} are our contributions.

$\mathcal{L}_{\text{context}}$	\mathcal{L}_{reg}	$\mathcal{L}_{\text{contrast}}$	λ	γ	CUB R@1	Cars R@1	SOP R@1	iNat. R@1
✓		✓	0.45	0	66.2 ± 0.5	87.9 ± 0.5	82.88	20.54
✓	✓		0.8	0.2	70.4 ± 0.6	82.8 ± 0.9	81.65	31.78
	✓	✓	0	0.17	69.6 ± 0.4	88.5 ± 0.4	83.34	44.19
✓			1	0	70.3 ± 0.4	76.8 ± 1.0	81.91	21.63
✓	✓	✓	0.4	0.1	69.8 ± 0.3	89.2 ± 0.2	83.28	46.20

Table 3: Ablation Results 2. Here, we experiment with variations of $\mathcal{L}_{\text{context}}$, both by itself ($\lambda = 1$) and regularized by a small amount of contrastive loss ($\lambda = 0.8$). We exhaustively test various ways to simplify or modify the contextual loss presented in Eq. 9 - 14. On the SOP dataset, we show that all of the modifications to $\mathcal{L}_{\text{context}}$ decrease R@1.

Ablation ($\gamma = 0$)	SOP R@1		Explanation	
	$\lambda = 0.8$	$\lambda = 1.0$		
$\lambda\mathcal{L}_1$	$+(1 - \lambda)\mathcal{L}_{\text{contrast}}$	83.13	41.99	Eq. 18 (step 1 only)
$\lambda\mathcal{L}_{1,\sigma}$	$+(1 - \lambda)\mathcal{L}_{\text{contrast}}$	79.20	75.61	Eq. 18 but using θ_σ in Eq. 7
$\lambda\mathcal{L}_2$	$+(1 - \lambda)\mathcal{L}_{\text{contrast}}$	82.39	81.05	Eq. 19 (skip step 3)
$\tilde{W} = \mathbb{1}_{\mathcal{N}_{k+\epsilon}}$ (skip step 2)		82.74	81.74	
$\lambda\mathcal{L}_{\text{context},\min}$	$+(1 - \lambda)\mathcal{L}_{\text{contrast}}$	66.57	–	Use min instead of \odot for logical-and
$\lambda\mathcal{L}_{\text{context},\sigma}$	$+(1 - \lambda)\mathcal{L}_{\text{contrast}}$	82.39	77.84	Use θ_σ for all steps
$\lambda\mathcal{L}_{\text{context},M_+}$	$+(1 - \lambda)\mathcal{L}_{\text{contrast}}$	82.31	80.72	Remove M_- term from Eq. 10
No stop gradient in Eq. 10		73.03	–	
$\lambda\mathcal{L}_{\text{context}}$	$+(1 - \lambda)\mathcal{L}_{\text{contrast}}$	83.20	82.04	Final results without \mathcal{L}_{reg}

and 0.6 %, resp. We achieve more modest R@1 gains of 0.2% and 0.3% on SOP and iNaturalist, resp. We note that these gains are significant because the standard deviation is relatively small. We also note that while ProxyNCA++ and MS are the best baselines on CUB and Cars, ROADMAP is the best baseline on SOP and iNaturalist. Our method is the best across *all* datasets, suggesting that it is more versatile.

5.1 ABLATION RESULTS

Table 2 shows R@1 on all datasets with one or two of the three losses removed. This table shows that in general, the best R@1 performance is achieved by a combination of the three losses in our framework. A combination of only \mathcal{L}_{reg} and $\mathcal{L}_{\text{contrast}}$ achieves better performance than $\mathcal{L}_{\text{contrast}}$ alone across all datasets (see first row of Table 1). The last row of Table 2 shows that adding $\mathcal{L}_{\text{context}}$ achieves 0.7% and 2% R@1 gains on Cars and iNaturalist, resp. These observations show that both of our contributions (\mathcal{L}_{reg} and $\mathcal{L}_{\text{context}}$) are necessary to achieve state-of-the-art performance.

In Table 3, we evaluate the contribution of each step in the calculation of $\mathcal{L}_{\text{context}}$. These experiments focus on dissecting $\mathcal{L}_{\text{context}}$, so we set $\gamma = 0$ (no similarity regularizer) and $\lambda = 0.8$ or 1.0. Overall, all of the modifications tested in Table 3 decrease the performance of $\mathcal{L}_{\text{context}}$ in terms of R@1. In particular: row 1 shows that including only step 1 leads to a collapsed representation when $\lambda = 1.0$; rows 3 and 4 show that steps 3 and 2 of the loss calculation are necessary; row 7 shows that the M_- term in Eq. 10 is necessary; rows 2 and 6 show that our approach to optimizing θ in Eq. 6 is better than using a sigmoid. We defer hyperparameter tuning results to the Appendix.

6 CONCLUSION

In this work we propose a novel supervised metric learning loss based on contextual similarity optimization. Our contextual loss improves the R@1 performance significantly over the current state-of-the-art on four diverse benchmarks, when regularized by the contrastive loss and a similarity regularizer. As our core technical contribution, we propose a highly non-trivial framework to optimize the contextual similarity, overcoming several non-differentiability issues. Namely, we show that the heaviside function can be effectively optimized by setting a positive gradient in the backward pass, and the intersection between sets can be optimized by indicator matrix multiplication. We carefully justify our framework by analyzing the gradient of the loss and with an exhaustive ablation study.

ETHICS STATEMENT

We note that metric learning can be applied to controversial problems such as person re-identification and face re-identification. Our work is mainly foundational, so does not contribute directly to these applications. We also limit our experimentation to the image retrieval aspect of metric learning.

REPRODUCIBILITY STATEMENT

We will release code publicly on GitHub. Our main results are reproducible by running the code. Instructions on how to run the code will be provided in a README file. We include details on hardware requirements in the Appendix.

REFERENCES

- Relja Arandjelović and Andrew Zisserman. Three things everyone should know to improve object retrieval. In *2012 IEEE conference on computer vision and pattern recognition*, pp. 2911–2918. IEEE, 2012.
- Malik Boudiaf, Jérôme Rony, Imtiaz Masud Ziko, Eric Granger, Marco Pedersoli, Pablo Piantanida, and Ismail Ben Ayed. A unifying mutual information view of metric learning: cross-entropy vs. pairwise losses. In *European conference on computer vision*, pp. 548–564. Springer, 2020.
- Andrew Brown, Weidi Xie, Vicky Kalogeiton, and Andrew Zisserman. Smooth-ap: Smoothing the path towards large-scale image retrieval. In *European Conference on Computer Vision*, pp. 677–694. Springer, 2020.
- Fatih Cakir, Kun He, Xide Xia, Brian Kulis, and Stan Sclaroff. Deep metric learning to rank. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 1861–1870, 2019.
- Kyunghyun Cho, Bart Van Merriënboer, Dzmitry Bahdanau, and Yoshua Bengio. On the properties of neural machine translation: Encoder-decoder approaches. *arXiv preprint arXiv:1409.1259*, 2014.
- Yixiao Ge, Dapeng Chen, and Hongsheng Li. Mutual mean-teaching: Pseudo label refinery for unsupervised domain adaptation on person re-identification. *arXiv preprint arXiv:2001.01526*, 2020a.
- Yixiao Ge, Feng Zhu, Dapeng Chen, Rui Zhao, and Hongsheng Li. Self-paced contrastive learning with hybrid memory for domain adaptive object re-id. In *Advances in Neural Information Processing Systems*, 2020b.
- Raia Hadsell, Sumit Chopra, and Yann LeCun. Dimensionality reduction by learning an invariant mapping. In *2006 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'06)*, volume 2, pp. 1735–1742. IEEE, 2006.
- Sungyeon Kim, Dongwon Kim, Minsu Cho, and Suha Kwak. Self-taught metric learning without labels. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 7431–7441, 2022.
- Jonathan Krause, Michael Stark, Jia Deng, and Li Fei-Fei. 3d object representations for fine-grained categorization. In *Proceedings of the IEEE international conference on computer vision workshops*, pp. 554–561, 2013.
- Yair Movshovitz-Attias, Alexander Toshev, Thomas K Leung, Sergey Ioffe, and Saurabh Singh. No fuss distance metric learning using proxies. In *Proceedings of the IEEE International Conference on Computer Vision*, pp. 360–368, 2017.
- Kevin Musgrave, Serge Belongie, and Ser-Nam Lim. A metric learning reality check. In *European Conference on Computer Vision*, pp. 681–699. Springer, 2020a.
- Kevin Musgrave, Serge Belongie, and Ser-Nam Lim. Pytorch metric learning, 2020b.

- Hyun Oh Song, Yu Xiang, Stefanie Jegelka, and Silvio Savarese. Deep metric learning via lifted structured feature embedding. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 4004–4012, 2016.
- Elias Ramzi, Nicolas Thome, Clément Rambour, Nicolas Audebert, and Xavier Bitot. Robust and decomposable average precision for image retrieval. *Advances in Neural Information Processing Systems*, 34:23569–23581, 2021.
- Eu Wern Teh, Terrance DeVries, and Graham W Taylor. Proxynca++: Revisiting and revitalizing proxy neighborhood component analysis. In *European Conference on Computer Vision*, pp. 448–464. Springer, 2020.
- Grant Van Horn, Oisin Mac Aodha, Yang Song, Yin Cui, Chen Sun, Alex Shepard, Hartwig Adam, Pietro Perona, and Serge Belongie. The inaturalist species classification and detection dataset. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 8769–8778, 2018.
- Catherine Wah, Steve Branson, Peter Welinder, Pietro Perona, and Serge Belongie. The caltech-ucsd birds-200-2011 dataset. 2011.
- Xun Wang, Xintong Han, Weilin Huang, Dengke Dong, and Matthew R Scott. Multi-similarity loss with general pair weighting for deep metric learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 5022–5030, 2019.
- Kilian Q Weinberger, John Blitzer, and Lawrence Saul. Distance metric learning for large margin nearest neighbor classification. *Advances in neural information processing systems*, 18, 2005.
- Chao-Yuan Wu, R Manmatha, Alexander J Smola, and Philipp Krahenbuhl. Sampling matters in deep embedding learning. In *Proceedings of the IEEE international conference on computer vision*, pp. 2840–2848, 2017.
- Mang Ye, Jianbing Shen, Gaojie Lin, Tao Xiang, Ling Shao, and Steven CH Hoi. Deep learning for person re-identification: A survey and outlook. *IEEE transactions on pattern analysis and machine intelligence*, 44(6):2872–2893, 2021.
- Andrew Zhai and Hao-Yu Wu. Classification is a strong baseline for deep metric learning. *arXiv preprint arXiv:1811.12649*, 2018.
- Zhun Zhong, Liang Zheng, Donglin Cao, and Shaozi Li. Re-ranking person re-identification with k-reciprocal encoding. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 1318–1327, 2017.

A APPENDIX

A.1 CODE AND ENVIRONMENT

We include code with our submission. We run experiments on 1 V100 GPU with 16 GB of memory. The CUB and Cars experiments take under one hour. The SOP and iNaturalist experiments take 4 hours and 6 hours, respectively. Some of our code is borrowed from ROADMAP. For faster experimentation, we use mixed precision floating point. Experiments take more than 12 hours on P100 GPUs, partially because mixed precision arithmetic does not appear to speed-up experiments as much on P100 GPUs compared to on V100 GPUs.

A.2 MINOR EXPERIMENTAL DETAILS

Augmentation On CUB, Cars and SOP, we use random resized crop with default parameters to crop the image to 256×256 pixels. We horizontally flip the image with 50% probability. For testing, we resize the image to 288 pixels, then center crop to 256 pixels. On iNaturalist, we random resize crop to 224×224 pixels, then flip horizontally with 50% probability. We use a smaller image size on iNaturalist so that the larger batch size can fit in the 16 GB of GPU memory. For testing, we resize the image to 256 pixels then center crop to 224 pixels.

Sampling We use a batch size of 128 on CUB, Cars and SOP. We use a batch size of 256 on iNaturalist. On all datasets, we train for 80 epochs. An epoch is defined as 15, 30, 655, and 576 batches for CUB, Cars, SOP, and iNaturalist respectively. On CUB and Cars, we use a random 4 per class sampler. On SOP and iNaturalist, we use a balanced hierarchical sampling strategy, since there are super-labels. In particular, half of each batch is randomly sampled from one super-label, and the other half is randomly sampled from another super-label. In order to sample in a balanced manner (to prevent over-emphasis of uncommon super-labels), we arrange all samples in the training dataset into half-batches with same super-labels at the beginning of each epoch. We then form batches by pairing up half-batches in a round-robin fashion.

Hyperparameters We use the same set of hyperparameters for all experiments. We present SOP R@1 results for various hyperparameter values in Figure 6. In Figure 5, we show the R@1 result on CUB and Cars for varying λ when $\gamma = 0$. We compare this result against ROADMAP with varying λ . ROADMAP is a recent state-of-the-art metric learning method, which uses a λ hyperparameter to balance between an AP loss and the contrastive loss. We show that even if λ and learning rates are tuned for ROADMAP on CUB and Cars, our method still achieves higher R@1 when the optimal λ and learning rates are chosen.

Loss Plot Figure 7 plots the contextual loss and test R@1 over the course of training on SOP, with varying λ . $\gamma = 0$. These plots show that the loss decreases over the course of training; in general, a lower contextual loss corresponds to a higher test R@1. This result suggests that the contextual loss is a reasonable objective for image retrieval.

Note on Contrastive Loss The contrastive loss has two margin parameters δ_- and δ_+ . The optimal values for δ_+ , the positive margin, is different depending on the value of γ (how much similarity regularization is used). For the “contrastive” results in row 1 of Table 1, we use $\delta_- = 0.6$ and $\delta_+ = 0.9$. For our results where $\gamma = 0.1$, we set $\delta_- = 0.6$ and $\delta_+ = 0.75$. We find this tighter positive margin to be optimal under similarity regularization. However, we note that $\delta_+ = 0.75$ is likely to be too small when only the contrastive loss is used, since positive pairs are not encouraged to be more similar than a cosine similarity of 0.75.

Note on Stop Gradient Note that we detach the normalization factors in Eq. 10. This is necessary because optimizing the size of the neighborhood is undesirable. On the contrary, we do not detach the normalization factor in Eq. 13. This is intentional even if somewhat un-intuitive. We show in Figure 9 that detaching $|\mathcal{R}_{k/2+\epsilon}(i)|$ in Eq. 13 increases R@1 on SOP when λ is high, but decreases R@1 in all other scenarios.




Figure 5: Comparison of our method to ROADMAP with varying λ . $\gamma = 0$. The orange lines show R@1 of our method with varying λ and learning rates. The purple lines show R@1 of ROADMAP with varying λ and learning rates. $\lambda = 0$ corresponds to the contrastive loss. Optimal performance of our method is always better than ROADMAP on CUB and Cars.




Figure 6: Hyperparameter tuning on SOP. We experiment with various values for ϵ , γ and δ . The R@1 values are similar for the different hyperparameter choices.




Figure 7: Plot of contextual similarity loss $\mathcal{L}_{\text{context}}$ with varying λ . Each color is a different λ . Results are on the SOP dataset. Observe that the contrastive loss implicitly minimizes the contextual loss (blue line). Also observe that in general, a lower contextual loss corresponds to a higher test R@1. This shows that our loss function is a reasonable objective. Note that the plot on the left plots $\mathcal{L}_{\text{context}}$, not the complete training loss. $\gamma = 0$.




Figure 8: Ablation results on CUB and Cars. The blue line represents test R@1 of our framework for varying λ with $\gamma = 0$. The other colors represent various modifications or simplifications of the loss function (reference Table 3 and Section 4 for full description). We perform each ablation experiment with varying λ since the optimal λ is not constant across all experiments. Clearly, the modified versions of $\mathcal{L}_{context}$ are sub-optimal.




Figure 9: Investigating the effect of detaching $|\mathcal{R}_{k/2+\epsilon}(i)|$ in Eq. 13. The blue line shows R@1 of our framework with varying λ . $\gamma = 0$. The orange line shows R@1 after detaching $|\mathcal{R}_{k/2+\epsilon}(i)|$. Overall, detaching this normalization factor is undesirable.