

A Non-isotropic Probabilistic Take on Proxy-based Deep Metric Learning

Michael Kirchhof^{1,*} , Karsten Roth^{1,*} , Zeynep Akata¹ , and Enkelejda Kasneci¹ 

¹University of Tübingen, Germany. (*) equal contribution

Abstract. Proxy-based Deep Metric Learning (DML) learns deep representations by embedding images close to their class representatives (*proxies*), commonly with respect to the angle between them. However, this disregards the embedding norm, which can carry additional beneficial context such as class- or image-intrinsic uncertainty. In addition, proxy-based DML struggles to learn class-internal structures. To address both issues at once, we introduce non-isotropic probabilistic proxy-based DML. We model images as directional von Mises-Fisher (vMF) distributions on the hypersphere that can reflect image-intrinsic uncertainties. Further, we derive non-isotropic von Mises-Fisher (nivMF) distributions for class proxies to better represent complex class-specific variances. To measure the proxy-to-image distance between these models, we develop and investigate multiple distribution-to-point and distribution-to-distribution metrics. Each framework choice is motivated by a set of ablational studies, which showcase beneficial properties of our probabilistic approach to proxy-based DML, such as uncertainty-awareness, better behaved gradients during training, and overall improved generalization performance. The latter is especially reflected in the competitive performance on the standard DML benchmarks, where our approach compares favourably, suggesting that existing proxy-based DML can significantly benefit from a more probabilistic treatment. Code is available at github.com/ExplainableML/Probabilistic_Deep_Metric_Learning.

Keywords: Deep Metric Learning, Von Mises-Fisher, Non-Isotropy, Probabilistic Embeddings, Uncertainty

1 Introduction

Understanding and encoding visual similarity is a key concept that drives applications ranging from image (video) retrieval [70,60,65,27,3] to clustering [1] and face re-identification [20,56,33,9]. Most commonly, approaches leverage Deep Metric Learning (DML) [56,60,40,70,52] to reformulate visual similarity learning into a surrogate, contrastive representation learning problem: Here, a deep network is tasked to embed images such that a simple predefined distance metric over pairs of embeddings represents their actual semantic relations. Similar contrastive learning is used for representation learning tasks s.a. supervised image

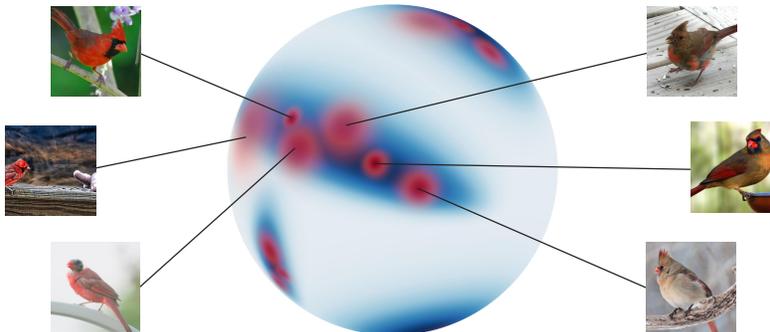


Fig. 1: Class proxy distributions (blue) and image distributions (red) embedded on the 3D unit sphere. The central proxy has a non-isotropic variance, so it can represent the high variance in body color between male (left) and female (right) cardinals and the low variance in their beak shape (top to bottom). Ambiguous images (e.g. middle left) have higher variance than images that clearly show class-discriminating features (top left, middle right). Best viewed in color.

classification [25] or self-supervised learning [18,5]. Common DML approaches are formulated as ranking tasks over data tuples (e.g. pairs [15], triplets [56] or quadruplets [6]) of similar and dissimilar samples. Unfortunately, the complexity of sampling such tuples grows exponentially with the tuples size [70]. This has motivated recent advances in DML to focus on *proxy-based* approaches, where the similar samples are summarized into learnable proxy representations [40,47] against which the sample embeddings are contrasted.

While this allows for fast convergence and reliable generalization, drawbacks may arise both in the treatment of proxies and samples: Firstly, the deterministic treatment of sample representations does not offer any degrees of freedom to address ambiguities and uncertainty (e.g., an image of a bird covered by branches). Secondly, isotropic distance scores between proxy-sample pairs (e.g., cosine similarity) provide only limited tools for the network to derive the similarity of samples within a class, as the distance to each proxy alone is insufficient to resolve relative sample placements around a proxy. This hinders class-specific variance and substructures to be successfully accounted for, which have been shown to notably benefit downstream generalization performance [52,37].

To address these issues, we propose a *probabilistic* interpretation of proxy-based DML. Driven by the fact that modern DML consistently operates on hyperspherical (i.e., normalized) representations [70,52], we derive hyperspherical von-Mises Fisher (vMF) distributions for each sample. A sample embedding’s direction controls the placement on the hypersphere, and therefore its semantic content, and its norm parametrizes the certainty of the distribution. In conjunction, we also treat class proxies probabilistically, but through *non-isotropic* vMF distributions. This enforces the distributional prior over each class proxy to explicitly account for different, non-isotropic distributions, capturing more complex

class-specific sample distributions (c.f. Figure 1). As this moves the DML training from point-based to distributional comparisons, we merge both components into a sound setup by motivating distribution-to-distribution matching metrics based on probabilistic product kernels. Our full framework is supported through an extensive set of derivations and experimental ablations that showcase and support how the extension to probabilistic proxy-based DML offers significant improvements, with competitive performance across the standard DML benchmarks – CUB200-2011 [64], CARS196 [30], and Stanford Online Products [42] – even when compared to much more complex training methods.

Overall, our contributions can be summarized as: **(1)** We propose and derive a novel probabilistic interpretation of proxy-based DML to account for sample and class ambiguities by reformulating the standard proxy-based metric learning approach to a distributional one on the hypersphere. **(2)** We extend the vMF model to a non-isotropical one for each class proxy to better incorporate and address intra-class substructures for better generalization. **(3)** We introduce various distribution-to-distribution metrics for DML and contrast them to traditional point-to-point metrics. **(4)** We support our proposed framework through various derivational and experimental ablations showcasing how a distributional treatment can positively impact the learned representation spaces. **(5)** Finally, we benchmark against standard DML approaches and provide further significant experimental support for our probabilistic approach to proxy-based DML.

2 Related Work

Deep Metric Learning comprises several conceptually different approaches. Firstly, one can define ranking tasks over data tuples such as pairs [15,70], triplets [56], quadruplets [6] or higher-order variants [42,60,67]. An underlying network then learns to solve each tuple presented by learning a representation space in which distances between embeddings correctly reflect their respective semantics/labelling. However, as the sizes of presented tuples increase, so does the tuple space each ranking task is sampled from, resulting in notable redundancy and impacted convergence behaviour [56,70,52]. As a result, a secondary branch evolved focusing on heuristics which target ranking tuples fulfilling a set of pre-defined [56,70,67,71] or learned [16,50] criteria. In a similar vein, DML research has also tried to address the sampling complexity issue through the replacement of tuple components with learned concept representations denoted as *proxies*, with some approaches leveraging proxies in a classification-style setting [73,9] or in a ranking fashion, where each sample is contrasted against a respective proxy [40,63,26,47]. Finally, benefits have also been found in orthogonal extensions and fundamental improvements to the general DML training pipeline, through various different approaches such as the usage of adversarial training [10], synthetic samples [32,75], higher-order or curvilinear metric learning [21,4], feature mining for ranking [49,38,37] or proxy-based [54] approaches, a breakdown of the overall metric space into subspaces [55,43,44], orthogonal modalities [53] or knowledge distillation [51]. Our proposed probabilistic proxy-based DML falls into this line

of work, but is orthogonal to these other approaches, as these extensions can be applied in a method-agnostic fashion. In particular, we extend proxy-based DML by specifically accounting for sample and class ambiguity through a distributional treatment of samples and proxies, and by utilizing non-isotropic proxy distributions to encourage more complex intra-class distributions around each proxy, which has been shown to be beneficial for generalization [52,39,54].

Probabilistic Embeddings. Various approaches to DML can already be framed from a more probabilistic standpoint, where softmax-based approaches on the basis of cosine similarities [9,63,73] can be seen as analytical class posteriors if each class assumes a von Mises-Fisher (vMF) distribution [45,74]. While these methods implicitly model classes as vMFs, probabilistic embedding approaches further model each sample as a distribution in the embedding space [58,31,57]. This allows the model to express uncertainty when images are ambiguous. Recent works argue that this ambiguity is captured in the image embedding’s norm [57,48,31]: [57] argues that the embedding of an image that shows many class-discriminative features of one class consists of several vectors that all point in the same direction, resulting in a higher norm. On this basis, [57,31] pioneered the use of embedding direction and norm to model each image as a vMF distribution, in particular for supervised classification. Utilizing vMF distributions, we are the first to introduce a full probabilistic proxy-based DML framework, yielding distribution-to-distribution metrics. Additionally, we propose a non-isotropic vMF for proxy distributions, which allows us to represent richer class structures in the embedding space beneficial to generalization [52,37].

3 Non-isotropic Probabilistic Proxy-based DML

3.1 A Probabilistic Interpretation of Proxy-based DML

In this section, we extend the common DML framework to a probabilistic one. Fundamentally, DML aims to find embedding functions $e : \mathcal{X} \rightarrow \mathcal{E}$ from image $\mathcal{X} \subset \mathbb{R}^{H \times W \times 3}$ to M -dimensional metric embedding spaces $\mathcal{E} \subset \mathbb{R}^M$ such that a distance function $d : \mathcal{E} \times \mathcal{E} \rightarrow \mathbb{R}$ between embeddings $z_1 = e(x_1)$ and $z_2 = e(x_2)$ of images $x_1, x_2 \in \mathcal{X}$ reflects the semantic relation between them. The embedding space \mathcal{E} is chosen to be the M -dimensional unit hypersphere $\mathcal{E} = \mathcal{S}^{M-1}$, i.e. $\|z\| = 1$. While an euclidean \mathcal{E} might appear more natural, recent works in DML [70,67,52,51,26] and other contrastive learning domains like self-supervised learning [66,5,18,8] have seen significant benefits in a directional treatment through normalization of embeddings to the unit hypersphere. This can in parts be attributed to better scaling with increased embedding dimensions [68] and semantic information being mostly directionally encoded [48]. To learn the respective embedding space \mathcal{E} , DML commonly employs ranking objectives over sample tuples. Based on the class assignments for each sample, an embedding network is tasked to minimize distances between same-class samples while maximizing them when classes differ. More recently, proxy-based approaches [40,63,26,47] directly model the class assignments by introducing class representatives during training – the proxies $p \in \mathcal{S}^{M-1}$. These are contrasted against

the sample embeddings $e(x) = z$ using a NCA-like [14] formulation (ProxyNCA, [40]), which was slightly modified by [63] as a softmax-loss

$$\mathcal{L}_{\text{NCA}++} = \log \frac{\exp(-d(\mathbf{p}^*, z)/t)}{\sum_{c=1}^C \exp(-d(\mathbf{p}_c, z)/t)}. \quad (1)$$

Here, \mathbf{p}^* denotes the ground-truth proxy associated with z , t a temperature, and d a distance metric, most commonly the negative cosine similarity $d = -s$ with $s(\mathbf{p}_c, z) = (\mathbf{p}_c z) / (\|\mathbf{p}_c\| \|z\|)$. This implies a problematic assumption: Since only angles between samples and proxies are leveraged, class-specific distribution variances around each proxy cannot be accounted for. Second, the deterministic underlying network e induces a Dirac delta distribution over sample representations [59]. This treats all the input data the same regardless of the level of ambiguity, not accounting for sample-specific uncertainties.

Therefore, we suggest to represent samples and proxies as random variables Z and P with densities ζ and ρ on \mathcal{S}^{M-1} , which allows both samples and proxies to carry uncertainty context to address sample ambiguity while encouraging to account for more complex class distributions. This converts the above loss to

$$\mathcal{L} = \log \frac{\exp(-d(\rho^*, \zeta)/t)}{\sum_{c=1}^C \exp(-d(\rho_c, \zeta)/t)}. \quad (2)$$

Below in §3.2, we discuss how precisely ρ and ζ are parametrized, and in §3.3, we find a $d(\cdot, \cdot)$ suitable for distribution-to-distribution matching .

3.2 Probabilistic Sample and Proxy Representations

Sample Representations. A common distribution on \mathcal{S}^{M-1} is the von Mises-Fisher (vMF) distribution [13,35,79]. It parametrizes the sample distribution ζ by a direction vector $\mu_z \in \mathcal{S}^{M-1}$ that points towards the mode of the distribution and a concentration parameter $\kappa_z \in \mathbb{R}_{\geq 0}$ that controls the spread around the mode, where a higher κ_z yields a sharper distribution. The density ζ of a vMF-distributed sample $Z \sim \text{vMF}(\mu_z, \kappa_z)$ at a point $\tilde{z} \in \mathcal{S}^{M-1}$ is

$$\zeta(\tilde{z}) = C_M(\kappa_z) \exp(\kappa_z s(\tilde{z}, \mu_z)). \quad (3)$$

C_M is the normalizing function which we approximate in high-dimensions (see Supp. A). The advantage of the vMF is a duality to the un-normalized image embeddings $z = e(x) \in \mathbb{R}^M$: The natural parameter of the vMF is $\nu_z = \kappa_z \mu_z \in \mathbb{R}^M$, such that if we set $\mu_z = \frac{z}{\|z\|}$ and $\kappa_z = \|z\|$, the embedding norm gives the vMF concentration without needing to explicitly predict it (as necessary for normal distribution [7,58]). This is further motivated by recent findings indicating that CNNs encode the amount of visible class discriminative features in the norm of the embedding (e.g. [57]). We validate this assumption in §4.4.

Proxy Representations. It is possible to analogously treat the proxy distributions ρ as vMF distributions with parameters $\nu_\rho = \kappa_\rho \mu_\rho$. However, a limiting factor owed to the simplicity of the vMF is its isotropy: The vMF is equivariant

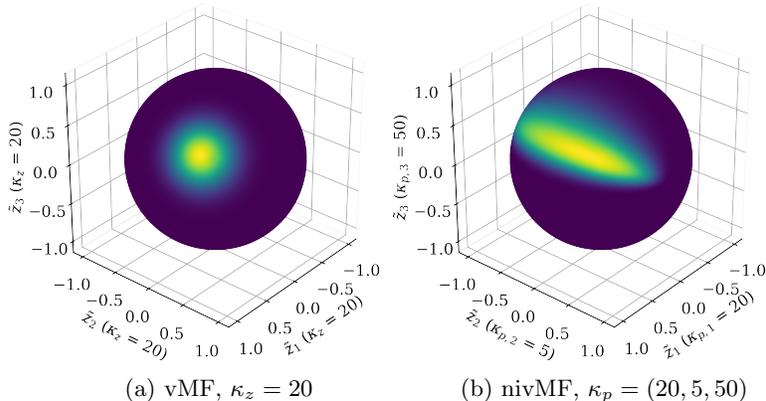


Fig. 2: Densities of (a) vMF and (b) non-isotropic vMF distributions on \mathcal{S}^2 . The density is proportional to the color gradient from violet (zero) to yellow (high).

in all directions as shown in Figure 2a. Proxies, however, need to account for more complex class distributions, i.e., non-isotropic ones (c.f. Fig. 2b). Generalized families of vMF distributions, such as Fisher-Bingham or Kent distributions [35,36,24], are able to capture non-isotropy. However, they use covariance matrices with a quadratic number of parameters and constraints on their eigenvectors. This complicates their training via gradient descent, especially in high dimensions. Hence, we propose a low-parameter vMF extension called non-isotropic von Mises-Fisher distribution (nivMF). Just like the vMF, the M -dimensional nivMF of a proxy p is parametrized by a direction $\mu_p \in \mathcal{S}^{M-1}$, but its concentration is described by a concentration *matrix* $K_p \in \mathbb{R}^{(M \times M)}$. To reduce its parameters, we assume $K_p = \text{diag}(\kappa_p) = \text{diag}(\kappa_{p,1}, \dots, \kappa_{p,M})$ to be a diagonal matrix where $\kappa_{p,m} > 0, m = 1, \dots, M$, gives the concentration per dimension. They are treated as learnable parameters (see Supp.G). Then, we define the density ρ of a nivMF distributed proxy $P \sim \text{nivMF}(\mu_p, K_p)$ at a point $\tilde{z} \in \mathcal{S}^{M-1}$ as

$$\rho = f_P(\tilde{z}) := C_M(\|K_p \mu_p\|) D(K_p) \exp(\|K_p \mu_p\| s(K_p \tilde{z}, K_p \mu_p)). \quad (4)$$

with vMF normalizer C_M , and D approximating an additional normalizing constant (see Supp. B). Intuitively, the nivMF is obtained from a vMF by a change-of-variable transformation: The unit sphere is stretched into an ellipsoid with axis lengths $\kappa_m, m = 1, \dots, M$, before the angle to the mode μ_p is measured. Thus, distances of \tilde{z} to μ_p along dimensions with high concentrations are emphasized and distances along dimensions with low concentrations are weighted less. In effect, the M -dim K_p is projected onto the $(M-1)$ -dim tangential plane of μ_p and controls the density’s spherical shape (see Fig. 2b). The remaining concentration projected on the μ_p -axis, i.e., $\|K_p \mu_p\|$, controls the density’s peakedness, analogously to the κ parameter from a standard vMF. Thus, when $K_p = cI_M$ is the identity matrix scaled by some $c > 0$, the nivMF simplifies to a vMF (up to a constant due to an approximation, see Supp. B).

3.3 Comparing Distributions Instead of Points

As proxies and images are no longer modeled as points but as distributions, we present several distribution-to-distribution metrics (in the sense of distance functions d in DML – formally, they are no metrics as they don’t fulfill the triangle inequality) and contrast them to traditional distribution-to-point metrics.

Distribution-to-Distribution Metrics. Probability product kernels (PPK) [22] are a family of metrics to compare two distributions ρ and ζ by the product of their densities. One member of this family is the expected likelihood kernel (or mutual likelihood score [58]). Although there is no analytical solution for nivMFs, we can derive a Monte-Carlo approximation

$$d_{\text{EL-nivMF}}(\rho, \zeta) := -\log \left(\int_{\mathcal{E}} \rho(a) d\zeta(a) \right) \approx -\log \left(\frac{1}{N} \sum_{\substack{i=1, \dots, N \\ z_i \sim \zeta}} \rho(z_i) \right), \quad (5)$$

where N is the number of samples. Similar to [57], we empirically found that a low number of samples ($N = 5$) is sufficient. We use [8] to sample from ζ .

The expected likelihood kernel is advantageous since it is easily Monte-Carlo approximated, but there are other distribution-to-distribution metrics we would like to survey. Hence, we derive them under a vMF assumption for ρ , where they have analytical solutions (see Supp. C). Namely, these are an analogous expected likelihood kernel $d_{\text{EL-vMF}}$, a related PPK kernel $d_{\text{B-vMF}}$, and a Kullback-Leibler distance $d_{\text{KL-vMF}}$. All three implicitly use the norm of the image embeddings in their calculations to respect the ambiguity, but differ in performance (see §4.3).

Distribution-to-point Metrics. Classical metrics like the cosine distance of the loss in Equation 1 implicitly assume a distribution for each proxy and evaluate its log-likelihood at each sample. Hence, we will refer to them as distribution-to-point metrics. E.g., the cosine metric used in Equation 1 is equivalent to the log-likelihood of the normalized sample embedding under vMF-distributed proxies with equal concentration values [17], i.e., $d_{\text{Cos}}(\rho, \zeta) := -s(\mu_p, \mu_z) = -\log(\rho(\mu_z))$. Another common example is the L2-distance $d_{\text{L2}}(\rho, \zeta) := (\nu_p - \nu_z)^2 = -\log(\rho(\nu_z))$ which is obtained by an equivariance normal distribution assumption for ρ . We analogously define $d_{\text{nivMF}}(\rho, \zeta) := -\log(\rho(\mu_z))$ under a nivMF assumption for ρ to benchmark it against the $d_{\text{EL-nivMF}}$ distance.

3.4 Probabilistic Proxy-based Deep Metric Learning

Utilizing distributional proxies ρ , distributional sample presentations ζ and the Monte-Carlo approximated Expected Likelihood Kernel $d_{\text{EL-nivMF}}(\rho, \zeta)$ we can fill in Equation 2 and define the probabilistic extension to proxy-based DML, precisely of the basic ProxyNCA ([63]), as

$$\mathcal{L}_{\text{NCA}++}^{\text{EL-nivMF}} = \log \frac{\exp(-d_{\text{EL-nivMF}}(\rho^*, \zeta)/t)}{\sum_{c=1}^C \exp(-d_{\text{EL-nivMF}}(\rho_c, \zeta)/t)}. \quad (6)$$

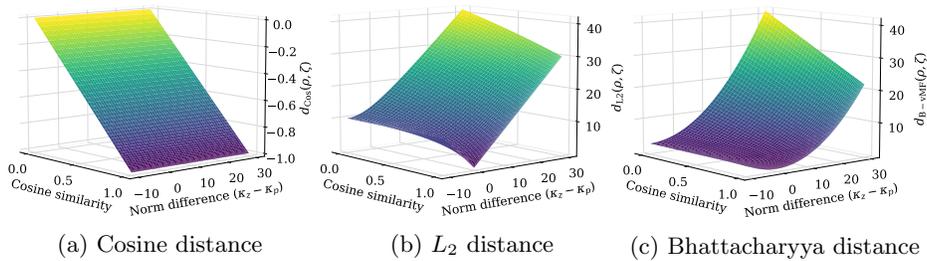


Fig. 3: Distances of a sample embedding to a vMF-distributed proxy with norm $\kappa_p = 10$. (a) and (b) treat the sample as a point and (c) as a vMF distribution.

While this can be used as standalone loss, it can also probabilistically enhance other proxy-based objectives $\mathcal{L}_{\text{Proxy-DML}}$, such as ProxyAnchor [26]. For easy usage in practice, we thus also propose using it as a regularizer via

$$\mathcal{L}_{\text{joint}}^{\text{NCA++}} = \mathcal{L}_{\text{NCA++}}^{\text{EL-nivMF}}(\rho, \zeta) + \omega \cdot \mathcal{L}_{\text{Proxy-DML}}(\mu_\rho, \mu_\zeta) \quad (7)$$

with regularization scale ω . Crucially, μ_ρ and μ_ζ of the proxy and sample distributions are shared parameters with the non-probabilistic objective’s proxies. This ensures alignment between the two learned representations spaces. The scaling ω balances the orthogonal benefits of the two approaches: An increasing ω highlights the non-probabilistic objective that encourages a better global alignment of distribution modes, and a decreasing ω yields a continuously more distributional treatment. For the remainder of this work, we use **EL-nivMF** for the standalone probabilistic extension of ProxyNCA (Eq. 6), and **PANC+EL-nivMF** for the probabilistically regularized ProxyAnchor (Eq. 7).

3.5 How Uncertainty-awareness Impacts Training

Before the experimental evaluation, we provide an insight into *how* incorporating uncertainty into the training benefits it. For this, we take a closer look at the norms of sample embeddings that, by duality, yield the concentration κ_z of ζ .

Uncertainty as Sample-wise Temperature. Figure 3 displays two distribution-to-point and one distribution-to-distribution metric with regard to the difference in norms and directions. We use the isotropic $d_{\text{B-vMF}}$ as a representative for distribution-to-distribution metrics since it has an analytical solution. While d_{Cos} ignores the difference in norms, d_{L_2} and the similar, yet smoother, $d_{\text{B-vMF}}$ incorporate it as a sample-wise temperature: The larger the norm of the sample gets, the steeper the metrics rise with increasing cosine distance. Thus, when comparing a sample to several proxies of roughly the same norm, their distances to the sample will be more uniform when the sample embedding norm is low and become more contrasted when it is high. In other words, ambiguous images produce more similar logits across all proxies and thus flatter class posterior distributions whereas highly certain images produce sharp posteriors.

Uncertainty as Gradients Scale. κ_z has another influence on the training: Differentiating the losses $\mathcal{L}_{NCA++}^{\text{Cos}}$ and $\mathcal{L}_{NCA++}^{\text{L2}}$, obtained when using the norm-agnostic d_{Cos} or the norm-aware d_{L2} as distance functions in Eq. 1, w.r.t. the cosine similarity between μ_z and μ_p (as in [26]) reveals (see Supp. F)

$$\frac{\delta \mathcal{L}_{NCA++}^{\text{Cos}}}{\delta \cos(\mu_p, \mu_z)} = \begin{cases} \frac{1}{t} \left(-1 + \frac{\exp(-d_{\text{Cos}}(\rho^*, \zeta)/t)}{\sum_{c=1}^C \exp(-d_{\text{Cos}}(\rho_c, \zeta)/t)} \right) & \text{if } p = p^* \\ \frac{1}{t} \frac{\exp(-d_{\text{Cos}}(\rho^*, \zeta)/t)}{\sum_{c=1}^C \exp(-d_{\text{Cos}}(\rho_c, \zeta)/t)} & \text{else} \end{cases} \quad (8)$$

$$\frac{\delta \mathcal{L}_{NCA++}^{\text{L2}}}{\delta \cos(\mu_p, \mu_z)} = \begin{cases} \frac{2\kappa_p \kappa_z}{t} \left(-1 + \frac{\exp(-d_{\text{L2}}(\rho^*, \zeta)/t)}{\sum_{c=1}^C \exp(-d_{\text{L2}}(\rho_c, \zeta)/t)} \right) & \text{if } p = p^* \\ \frac{2\kappa_p \kappa_z}{t} \frac{\exp(-d_{\text{L2}}(\rho^*, \zeta)/t)}{\sum_{c=1}^C \exp(-d_{\text{L2}}(\rho_c, \zeta)/t)} & \text{else} \end{cases}, \quad (9)$$

where p^* denotes the ground-truth class. Besides the sample-wise temperature in d_{L2} , the gradients differ in that the gradient of $\mathcal{L}_{NCA++}^{\text{L2}}$ scales proportionally to κ_z . This means that in batch-wise gradient descent, samples with a high embedding norm are pulled towards ground-truth proxies and pushed away from others stronger than samples with low norm. In other words, the impact of an image on the structuring process of the embedding space depends on its ambiguity. This holds similarly for the distribution-to-distribution metrics, but is harder to derive than for d_{L2} . This analysis unveils that using the Euclidean d_{L2} distance is adequate during training albeit switching to the hyperspherical d_{Cos} at retrieval-time, as it can be seen as a simple approximation to the uncertainty-aware training of hyperspherical distribution-to-distribution metrics.

4 Experiments

We now detail the experiments (§4.1) that benchmark our method (§4.2), before surveying different distr.-to-distr. metrics (§4.3) and the role of the norm (§4.4).

4.1 Experimental Details

Implementations. All experiments use `PyTorch` [46]. We follow standard DML protocols by leveraging ImageNet-pretrained ResNet50 [19] and Inception-V1 networks with Batch-Normalization [62] as encoders. Their weights are taken from `torchvision` [34] and `timm` [69]. To further ensure standardized training, we built upon the code and standardized DML protocols proposed in [52], using the Adam optimizer [28], a learning rate of 10^{-5} and weight decay of $4 \cdot 10^{-3}$. In the more open state-of-the-art comparison (Table 2), we additionally use step-wise learning rate scheduling. To ensure comparability and access to fast similarity search methods, all test-time retrieval uses cosine distances. To sample from vMF-distributions, we make use of [8] and [respective implementations](#). Further details on our method and hyperparameters are provided in Supp. H. All experiments were run on NVIDIA 2080Ti GPUs with 12GB VRAM.

Datasets. We benchmark on three standard datasets: CUB200-2011 [64] (has a 100/100 split of train and test bird classes with 11,788 images in total), CARS196 [30] (contains a 98/98 split of car classes and 16,185 images), and Stanford Online Products (SOP) [42] (covers 22,634 product categories and 120,053 images).

Table 1: We re-run various strong benchmarks in the *standardized comparison* setting of [52]. We find strong improvements both when enhancing simple ProxyNCA towards probabilistic DML (**EL-nivMF**) and when using our approach as a regularizer on top of more versatile approaches (*PANC* + **EL-nivMF**).

| BENCHMARKS→ | CUB200-2011 | | CARS196 | | SOP | |
|---|-------------|------------|------------|------------|------------|------------|
| APPROACHES ↓ | R@1 | mAP@1000 | R@1 | mAP@1000 | R@1 | mAP@1000 |
| Sample-based Baselines. | | | | | | |
| Margin [70] | 62.9 ± 0.4 | 32.7 ± 0.3 | 80.1 ± 0.2 | 32.7 ± 0.4 | 78.4 ± 0.1 | 46.8 ± 0.1 |
| Multisimilarity [67] | 62.8 ± 0.2 | 31.1 ± 0.3 | 81.6 ± 0.3 | 31.7 ± 0.1 | 76.0 ± 0.1 | 43.3 ± 0.1 |
| Standard versus Probabilistic. | | | | | | |
| ProxyNCA [40,63] | 63.2 ± 0.2 | 33.4 ± 0.1 | 78.8 ± 0.2 | 31.9 ± 0.2 | 76.2 ± 0.1 | 43.0 ± 0.1 |
| EL-nivMF | 64.8 ± 0.4 | 34.3 ± 0.3 | 82.1 ± 0.3 | 33.4 ± 0.2 | 76.6 ± 0.2 | 43.3 ± 0.1 |
| Probabilistic DML as Regularization. | | | | | | |
| ProxyAnchor (<i>PANC</i> , [26]) | 64.4 ± 0.3 | 33.2 ± 0.3 | 82.4 ± 0.4 | 34.2 ± 0.3 | 78.0 ± 0.1 | 45.5 ± 0.1 |
| <i>PANC</i> + EL-nivMF | 66.5 ± 0.3 | 35.3 ± 0.1 | 83.6 ± 0.2 | 35.1 ± 0.1 | 78.2 ± 0.1 | 45.6 ± 0.1 |

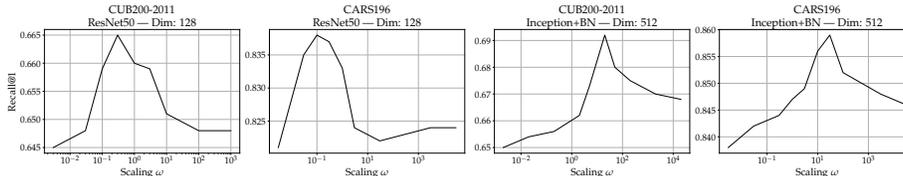


Fig. 4: *Probabilistic regularization as a function of the scaling factor ω* . We find a notable benefit when accounting for both orthogonal enhancements, i.e., the more probabilistic treatment (*decreasing ω*) and the better global alignment of the proxy distribution modes (*increasing ω*).

4.2 Quantitative evaluation of probabilistic proxy-based DML

Standardized comparison. We first follow protocols proposed in [52], which suggest comparisons under equal pipeline and implementation settings (and no learning rate scheduling) to determine the true benefits of a proposed method, unbiased by external covariates. Particularly, we thus compare the standard ProxyNCA (see Eq. 1) against our proposed **EL-nivMF** extension of ProxyNCA that includes sample and proxy distributions with distribution-to-distribution metrics during training. We further apply **EL-nivMF** as a probabilistic regularizer on top of the strong, but hyperparameter-heavy ProxyAnchor objective. Here, we only optimize the scaling ω . Finally, we rerun the two strongest sample-based methods used in [52]. In all cases, Table 1 shows significant improvements in performance and outperforms the sample-based methods. First, converting from standard to probabilistic proxy-based DML (ProxyNCA \rightarrow **EL-nivMF**) increases R@1 on CUB200-2011 by 1.6pp, 3.3pp on Cars196 and 0.4pp on SOP. This highlights the benefits of accounting for uncertainty and explicitly encouraging non-isotropic intra-class variance. However, due to the large number of proxies and low number of samples per class, on SOP benefits are limited when

Table 2: *Comparison to Literature*, separated by backbones and embedding dimensions. **Bold** denotes best results for a respective Backbone/Dim. subset, **bold** the overall best. Results show that our probabilistically regularized ProxyAnchor method matches or beats previous, in parts notably more complex state-of-the-art methods.

| BENCHMARKS → | | CUB200 [64] | | | CARS196 [30] | | | SOP [42] | | | |
|---------------------------|--------------------|-------------|-------------|-------------|--------------|-------------|-------------|-------------|-------------|-------------|-------------|
| METHODS ↓ | Venue | Arch/Dim. | R@1 | R@2 | NMI | R@1 | R@2 | NMI | R@1 | R@10 | NMI |
| Margin [70] | <i>ICCV '17</i> | R50/128 | 63.6 | 74.4 | 69.0 | 79.6 | 86.5 | 69.1 | 72.7 | 86.2 | 90.7 |
| Div&Conq [55] | <i>CVPR '19</i> | R50/128 | 65.9 | 76.6 | 69.6 | 84.6 | 90.7 | 70.3 | 75.9 | 88.4 | 90.2 |
| MIC [49] | <i>ICCV '19</i> | R50/128 | 66.1 | 76.8 | 69.7 | 82.6 | 89.1 | 68.4 | 77.2 | 89.4 | 90.0 |
| PADS [50] | <i>CVPR '20</i> | R50/128 | 67.3 | 78.0 | 69.9 | 83.5 | 89.7 | 68.8 | 76.5 | 89.0 | 89.9 |
| RankMI [23] | <i>CVPR '20</i> | R50/128 | 66.7 | 77.2 | 71.3 | 83.3 | 89.8 | 69.4 | 74.3 | 87.9 | 90.5 |
| <i>PANC + EL-niVMF</i> | - | R50/128 | 67.0 | 77.6 | 70.0 | 84.0 | 90.0 | 69.5 | 78.6 | 90.5 | 90.1 |
| NormSoft [73] | <i>BMVC '19</i> | R50/512 | 61.3 | 73.9 | - | 84.2 | 90.4 | - | 78.2 | 90.6 | - |
| EPSHN [72] | <i>WACV '20</i> | R50/512 | 64.9 | 75.3 | - | 82.7 | 89.3 | - | 78.3 | 90.7 | - |
| Circle [61] | <i>CVPR '20</i> | R50/512 | 66.7 | 77.2 | - | 83.4 | 89.7 | - | 78.3 | 90.5 | - |
| DiVA [37] | <i>ECCV '20</i> | R50/512 | 69.2 | 79.3 | 71.4 | 87.6 | 92.9 | 72.2 | 79.6 | 91.2 | 90.6 |
| DCML-MDW [76] | <i>CVPR '21</i> | R50/512 | 68.4 | 77.9 | 71.8 | 85.2 | 91.8 | 73.9 | 79.8 | 90.8 | 90.8 |
| <i>PANC + EL-niVMF</i> | - | R50/512 | 69.3 | 79.3 | 72.1 | 86.2 | 91.9 | 70.3 | 79.4 | 90.7 | 90.6 |
| Group [12] | <i>ECCV '20</i> | IBN/512 | 65.5 | 77.0 | 69.0 | 85.6 | 91.2 | 72.7 | 75.1 | 87.5 | 90.8 |
| DR-MS [11] | <i>TAI '20</i> | IBN/512 | 66.1 | 77.0 | - | 85.0 | 90.5 | - | - | - | - |
| ProxyGML [78] | <i>NeurIPS '20</i> | IBN/512 | 66.6 | 77.6 | 69.8 | 85.5 | 91.8 | 72.4 | 78.0 | 90.6 | 90.2 |
| DRML [77] | <i>ICCV '21</i> | IBN/512 | 68.7 | 78.6 | 69.3 | 86.9 | 92.1 | 72.1 | 71.5 | 85.2 | 88.1 |
| <i>PANC + MemVir</i> [29] | <i>ICCV '21</i> | IBN/512 | 69.0 | 79.2 | - | 86.7 | 92.0 | - | 79.7 | 91.0 | - |
| <i>PANC + EL-niVMF</i> | - | IBN/512 | 69.5 | 80.0 | 71.0 | 86.4 | 92.0 | 71.3 | 79.2 | 90.4 | 90.2 |

compared to datasets such as CUB200-2011 and CARS196, as the estimation of our proxy distributions becomes noticeably noisier. When using **EL-nivMF** as probabilistic regularization, we find boosts of over $2.1pp$ and $1.2pp$ on CUB200-2011 and CARS196, respectively, with expected smaller improvements of $0.2pp$ on SOP. Generally however, the consistent improvements, whether as a standalone objective or as a regularization method, highlight the versatility of a probabilistic take on DML, and offer a strong proof-of-concept for future DML research to built upon.

Impact of different scaling factors ω . Figure 4 showcases the generalization performance as a function of the scaling weight ω (see Eq. 7). Higher ω denotes a more non-probabilistic treatment to the point of ignoring the distributional aspects and returning to the auxiliary ProxyAnchor loss [26]. Lower ω indicates a higher emphasis on distributional treatment of proxies (and samples). Across benchmarks and backbones, the best performance is reached with an ω that is neither high nor 0. Thus, the results highlight that our probabilistic proxy-based DML helps the better global realignment of each proxy distribution mode via ProxyAnchor, and vice-versa. Overall, R@1 increases up to $4pp$ at the most suitable scaling choice. This optimum is reached robustly in a large area around the peak (note the logarithmic x-axes).

Comparison against SOTA. After these strictly standardized comparisons, we now compare the combination of ProxyAnchor and **EL-nivMF**, which

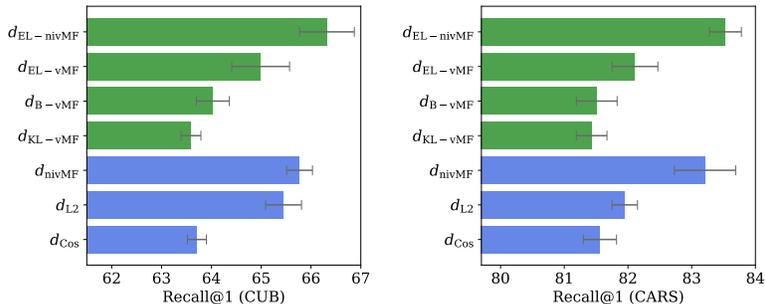


Fig. 5: Distance-to-point (blue) vs. distance-to-distance (green) metrics on CUB and CARS. Bars show average R@1 with standard deviation.

performed best in the previous study to the larger DML literature. The hyperparameters and pipeline components (e.g., learning rate, weight decay) differ between the approaches, and so the comparison should be taken with a grain of salt [52,41,57], but we still separate by the backbones and embedding dimensionalities, which are identified as the largest factors of variation [52]. Accounting for that, we find competitive performance on all benchmarks (c.f. Table 2), even when compared against other, much more complex state-of-the-art methods relying on multitask learning (DiVA [37], MIC [49]) or reinforcement learning (PADS [50]). This makes our probabilistic take on proxy-based DML a generally attractive approach to DML, with further potential improvements down the line by implementing the probabilistic perspective into these orthogonal extensions.

Computational overhead. We do note that training with **EL-nivMF** requires the differentiable drawing of samples from vMF-distributions (see Eq. 5 and [8]). This can increase the overall training time, but we found 2-5 samples to already be suitable, limiting the impact on overall walltime to $< 25\%$ against pure ProxyNCA. This is in line with other extensions of ProxyNCA (s.a. [21,49,55,50,37]). The retrieval walltime remains unaffected as cosine-similarity is deployed. As an alternative for rapid training, we provide further probabilistic distribution-to-distribution distances (d_{EL-vMF} , d_{B-vMF} , d_{KL-vMF}) along with analytical solutions (Supp. C), so that no sampling is required and computational overhead is negligible. We study them in the next section.

4.3 Quantitative Comparison of Metrics

Sections 3.3 and 3.2 provided numerous modeling choices for distributions and distance metrics that can be plugged into the probabilistic DML framework in Eq. 2. This section investigates these possibilities, ultimately motivating the particular choice of $d_{EL-nivMF}$, and also compares to more traditional distribution-to-point metrics. To ensure fair comparisons, we return to the standardized benchmark protocol of [52] using a 512-dimensional ResNet-50. All hyperparameters are fixed, except for the initial proxy norm and temperature, which are tuned via grid search on a validation set.

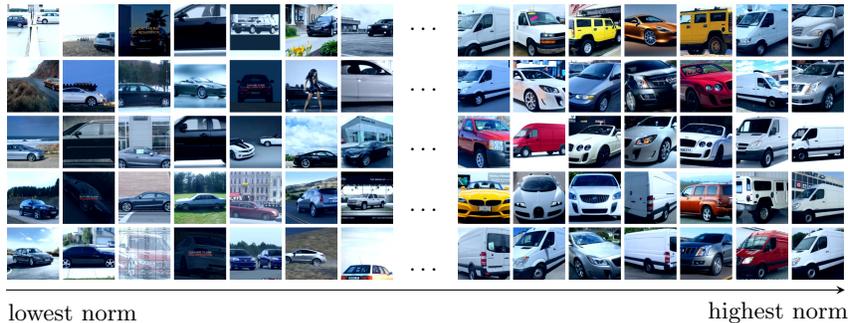


Fig. 6: CARS train images with lowest (left) to highest (right) embedding norms.

Figure 5 shows the R@1 of all three distribution-to-point and four distribution-to-distribution metrics on CUB and CARS. Comparing the distribution-to-point metrics, d_{L2} outperforms d_{Cos} on both datasets, but is dominated by d_{nivMF} . The non-isotropic approach also performs best within the distribution-to-distribution metrics. Within the three isotropic distribution-to-distribution metrics, d_{KL-vMF} shows the worst performance, with a small gap to the Bhattacharyya and a larger gap to the expected likelihood PPKs. This stands in line with preliminary findings of [7]. The latter performs within one standard deviation of d_{L2} . Altogether, we find that adding non-isotropy to the standard d_{Cos} (i.e., using d_{nivMF}) increases the R@1 by 2.1pp on CUB and 1.7pp on CARS. Further considering the image norm (i.e., $d_{EL-nivMF}$) adds another 0.6pp on CUB and 0.3pp on CARS.

The enhancement by non-isotropic modeling can be seen as inductive bias towards better resolution of intra-class variances and substructures (see Supp. L), which drives generalization performance [52,37,32,77,72]. The strong performance of d_{L2} is surprising as many current approaches use a d_{Cos} -based loss [9,26,63]. The crux is that d_{L2} in our setting still uses the cosine distance at retrieval-time, similar to, e.g., [2]. Using d_{L2} also as the retrieval metric would reduce the R@1 by up to $-5.34pp$ across all metrics and datasets, with the highest reduction appearing on the d_{L2} -trained model itself (see Supp. J). This supports the usage of the norm only during training, discussed in §3.5, where d_{L2} shares the uncertainty-awareness of distribution-to-distribution metrics, explaining the small gap between d_{L2} and d_{EL-vMF} . Thus, ultimately, we conjecture that it doesn’t matter whether an approach is motivated from a distribution-to-distribution or distribution-to-point perspective, as long as it considers the ambiguity of images (and proxies) during training.

4.4 Embedding Norms Encode Uncertainty

In the previous section, we found that considering the norms of embeddings during training leads to a higher performance. In this section, we qualitatively support that the learned norms actually correspond to a sample-wise ambiguity.

For this, we study the **EL-nivMF** model on CARS. Figure 6 shows the images with the lowest and highest embedding norm in the training set. In many

samples with low norm, characteristic parts of the cars are cropped out by the data augmentation (this also happens in the test set, hindering perfect accuracy). Others are overlaid or portray multiple distracting objects. In high-norm images, illumination and camera angle facilitate the detection of class-discriminative features. A competing hypothesis could be that high-norm images comprise mostly car classes with more distinctive designs. However, the differences between low and high norm images also hold within classes, see Supp. K and M. These findings are in line with [57,48,31] and support the hypothesis that the image norm indicates image certainties, motivated by being the sum of visible class-discriminative parts [57]. This justifies the $\kappa_z = \|z\|$ duality underlying the vMF assumption and is consistent with our analysis of uncertainty-aware training in §3.5.

5 Conclusion

This work proposes non-isotropic probabilistic proxy-based deep metric learning (DML) through uncertainty-aware training and non-isotropic proxy-distributions. Uncertainty-aware training is achieved by treating sample embeddings not as deterministic points but as directional distributions parametrized by embedding directions and, beyond popular DML approaches, norms. This allows semantic ambiguities to be decoupled from the directional semantic context, which mathematically manifests itself in sample-wise temperature scaling and certainty-weighted gradients. Additionally, our non-isotropic von Mises-Fisher distribution for proxies better models intra-class uncertainty, which introduces a low-parameter inductive prior for better generalizing embedding spaces. We support our approach through various ablation studies, which showcase that our proposed framework can operate both as a standalone objective and a probabilistic regularizer on top of existing proxy-based objectives. In both cases, we further found strong performances on the standard DML benchmarks, in parts matching or beating existing state-of-the-art methods. Our findings strongly indicate that a probabilistic treatment of proxy-based DML offers simple, orthogonal enhancements to existing DML methods and enables better generalization.

Limitations. We find that for applications with only few samples per class, the ability to estimate the non-isotropic proxy densities is limited (c.f. performance on SOP). For future work in such sparse settings, returning to the proposed isotropic distribution-to-distribution metrics or introducing across-class priors for the covariance matrices might serve as alternatives.

Acknowledgements. This work has been partially funded by the ERC (853489 - DEXIM) and DFG (2064/1 – Project number 390727645) under Germany’s Excellence Strategy. Michael Kirchhof and Karsten Roth thank the International Max Planck Research School for Intelligent Systems (IMPRS-IS) for support. Karsten Roth further acknowledges his membership in the European Laboratory for Learning and Intelligent Systems (ELLIS) PhD program.

References

1. Bouchacourt, D., Tomioka, R., Nowozin, S.: Multi-level variational autoencoder: Learning disentangled representations from grouped observations. In: Thirty-Second AAAI Conference on Artificial Intelligence (AAAI) (2018) [1](#)
2. Boudiaf, M., Rony, J., Ziko, I.M., Granger, E., Pedersoli, M., Piantanida, P., Ayed, I.B.: A unifying mutual information view of metric learning: cross-entropy vs. pairwise losses. In: Proceedings of the European Conference on Computer Vision (ECCV) (2020) [13](#)
3. Brattoli, B., Tighe, J., Zhdanov, F., Perona, P., Chalupka, K.: Rethinking zero-shot video classification: End-to-end training for realistic applications. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2020) [1](#)
4. Chen, S., Luo, L., Yang, J., Gong, C., Li, J., Huang, H.: Curvilinear distance metric learning. In: Advances in Neural Information Processing Systems 32, pp. 4223–4232. Curran Associates, Inc. (2019), <http://papers.nips.cc/paper/8675-curvilinear-distance-metric-learning.pdf> [3](#)
5. Chen, T., Kornblith, S., Norouzi, M., Hinton, G.: A simple framework for contrastive learning of visual representations. In: Proceedings of the 37th International Conference on Machine Learning (ICML) (2020) [2](#), [4](#)
6. Chen, W., Chen, X., Zhang, J., Huang, K.: Beyond triplet loss: a deep quadruplet network for person re-identification. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2017) [2](#), [3](#)
7. Chun, S., Oh, S.J., De Rezende, R.S., Kalantidis, Y., Larlus, D.: Probabilistic embeddings for cross-modal retrieval. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2021) [5](#), [13](#)
8. Davidson, T.R., Falorsi, L., De Cao, N., Kipf, T., Tomczak, J.M.: Hyperspherical variational auto-encoders. 34th Conference on Uncertainty in Artificial Intelligence (UAI) (2018) [4](#), [7](#), [9](#), [12](#)
9. Deng, J., Guo, J., Xue, N., Zafeiriou, S.: ArcFace: Additive angular margin loss for deep face recognition. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2019) [1](#), [3](#), [4](#), [13](#)
10. Duan, Y., Zheng, W., Lin, X., Lu, J., Zhou, J.: Deep adversarial metric learning. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2018) [3](#)
11. Dutta, U.K., Harandi, M., Sekhar, C.C.: Unsupervised deep metric learning via orthogonality based probabilistic loss. *IEEE Transactions on Artificial Intelligence* **1**(1) (2020) [11](#)
12. Elezi, I., Vascon, S., Torcinovich, A., Pelillo, M., Leal-Taixé, L.: The group loss for deep metric learning. In: Proceedings of the European Conference on Computer Vision (ECCV) (2020) [11](#)
13. Fisher, R.A.: Dispersion on a sphere. *Proceedings of the Royal Society of London. Series A. Mathematical and Physical Sciences* **217**(1130) (1953) [5](#)
14. Goldberger, J., Hinton, G.E., Roweis, S., Salakhutdinov, R.R.: Neighbourhood components analysis. In: Advances in Neural Information Processing Systems (NeurIPS) (2004) [5](#)
15. Hadsell, R., Chopra, S., LeCun, Y.: Dimensionality reduction by learning an invariant mapping. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (2006) [2](#), [3](#)

16. Harwood, B., Kumar, B., Carneiro, G., Reid, I., Drummond, T., et al.: Smart mining for deep metric learning. In: Proceedings of the IEEE International Conference on Computer Vision (ICCV) (2017) [3](#)
17. Hasnat, M.A., Bohné, J., Milgram, J., Gentric, S., Chen, L.: von Mises-Fisher mixture model-based deep learning: Application to face verification. arXiv preprint arXiv:1706.04264 (2017) [7](#)
18. He, K., Fan, H., Wu, Y., Xie, S., Girshick, R.: Momentum contrast for unsupervised visual representation learning. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2020) [2](#), [4](#)
19. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2016) [9](#)
20. Hu, J., Lu, J., Tan, Y.: Discriminative deep metric learning for face verification in the wild. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2014) [1](#)
21. Jacob, P., Picard, D., Histace, A., Klein, E.: Metric learning with horde: High-order regularizer for deep embeddings. In: The IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2019) [3](#), [12](#)
22. Jebara, T., Kondor, R.: Bhattacharyya and expected likelihood kernels. In: Learning Theory and Kernel Machines (2003) [7](#)
23. Kemertas, M., Pishdad, L., Derpanis, K.G., Fazly, A.: RankMI: A mutual information maximizing ranking loss. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2020) [11](#)
24. Kent, J.T.: The Fisher-Bingham distribution on the sphere. Journal of the Royal Statistical Society: Series B (Methodological) **44**(1) (1982) [6](#)
25. Khosla, P., Teterwak, P., Wang, C., Sarna, A., Tian, Y., Isola, P., Maschinot, A., Liu, C., Krishnan, D.: Supervised contrastive learning. Advances in Neural Information Processing Systems (NeurIPS) (2020) [2](#)
26. Kim, S., Kim, D., Cho, M., Kwak, S.: Proxy anchor loss for deep metric learning. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2020) [3](#), [4](#), [8](#), [9](#), [10](#), [11](#), [13](#)
27. Kim, S., Kim, D., Cho, M., Kwak, S.: Embedding transfer with label relaxation for improved metric learning. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2021) [1](#)
28. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. In: Bengio, Y., LeCun, Y. (eds.) 3rd International Conference on Learning Representations (ICLR) (2015) [9](#)
29. Ko, B., Gu, G., Kim, H.G.: Learning with memory-based virtual classes for deep metric learning. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV) (2021) [11](#)
30. Krause, J., Stark, M., Deng, J., Fei-Fei, L.: 3d object representations for fine-grained categorization. In: Proceedings of the IEEE International Conference on Computer Vision Workshops (CVPR) (2013) [3](#), [9](#), [11](#)
31. Li, S., Xu, J., Xu, X., Shen, P., Li, S., Hooi, B.: Spherical confidence learning for face recognition. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2021) [4](#), [14](#)
32. Lin, X., Duan, Y., Dong, Q., Lu, J., Zhou, J.: Deep variational metric learning. In: Proceedings of the European Conference on Computer Vision (ECCV) (2018) [3](#), [13](#)

33. Liu, W., Wen, Y., Yu, Z., Li, M., Raj, B., Song, L.: Sphereface: Deep hypersphere embedding for face recognition. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2017) [1](#)
34. Marcel, S., Rodriguez, Y.: Torchvision the machine-vision package of torch. *MM '10, Association for Computing Machinery* (2010) [9](#)
35. Mardia, K.V., Jupp, P.E.: *Directional statistics* (2009) [5](#), [6](#)
36. Mardia, K.V.: *Statistics of directional data*. *Journal of the Royal Statistical Society: Series B (Methodological)* **37**(3) (1975) [6](#)
37. Milbich, T., Roth, K., Bharadhwaj, H., Sinha, S., Bengio, Y., Ommer, B., Cohen, J.P.: Diva: Diverse visual feature aggregation for deep metric learning. In: Vedaldi, A., Bischof, H., Brox, T., Frahm, J.M. (eds.) *Proceedings of the European Conference on Computer Vision (ECCV)* (2020) [2](#), [3](#), [4](#), [11](#), [12](#), [13](#)
38. Milbich, T., Roth, K., Brattoli, B., Ommer, B.: Sharing matters for generalization in deep metric learning. *IEEE Trans. Pattern Anal. Mach. Intell.* **44**(1), 416–427 (jan 2022). <https://doi.org/10.1109/TPAMI.2020.3009620>, <https://doi.org/10.1109/TPAMI.2020.3009620> [3](#)
39. Milbich, T., Roth, K., Sinha, S., Schmidt, L., Ghassemi, M., Ommer, B.: Characterizing generalization under out-of-distribution shifts in deep metric learning. In: Ranzato, M., Beygelzimer, A., Dauphin, Y., Liang, P., Vaughan, J.W. (eds.) *Advances in Neural Information Processing Systems*. vol. 34, pp. 25006–25018. Curran Associates, Inc. (2021), <https://proceedings.neurips.cc/paper/2021/file/d1f255a373a3cef72e03aa9d980c7eca-Paper.pdf> [4](#)
40. Movshovitz-Attias, Y., Toshev, A., Leung, T.K., Ioffe, S., Singh, S.: No fuss distance metric learning using proxies. In: *Proceedings of the IEEE International Conference on Computer Vision (ICCV)* (2017) [1](#), [2](#), [3](#), [4](#), [5](#), [10](#)
41. Musgrave, K., Belongie, S., Lim, S.N.: A metric learning reality check. In: Vedaldi, A., Bischof, H., Brox, T., Frahm, J.M. (eds.) *Proceedings of the European Conference on Computer Vision (ECCV)* (2020) [12](#)
42. Oh Song, H., Xiang, Y., Jegelka, S., Savarese, S.: Deep metric learning via lifted structured feature embedding. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2016) [3](#), [9](#), [11](#)
43. Opitz, M., Waltner, G., Possegger, H., Bischof, H.: Bier-boosting independent embeddings robustly. In: *Proceedings of the IEEE International Conference on Computer Vision (ICCV)* (2017) [3](#)
44. Opitz, M., Waltner, G., Possegger, H., Bischof, H.: Deep metric learning with BIER: Boosting independent embeddings robustly. *IEEE transactions on pattern analysis and machine intelligence* (2018) [3](#)
45. Park, J., Yi, S., Choi, Y., Cho, D.Y., Kim, J.: Discriminative few-shot learning based on directional statistics. *arXiv preprint arXiv:1906.01819* (2019) [4](#)
46. Paszke, A., Gross, S., Chintala, S., Chanan, G., Yang, E., DeVito, Z., Lin, Z., Desmaison, A., Antiga, L., Lerer, A.: Automatic differentiation in pytorch. In: *NIPS Workshop on Automatic Differentiation* (2017) [9](#)
47. Qian, Q., Shang, L., Sun, B., Hu, J., Li, H., Jin, R.: Softtriple loss: Deep metric learning without triplet sampling. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)* (2019) [2](#), [3](#), [4](#)
48. Ranjan, R., Castillo, C.D., Chellappa, R.: L2-constrained softmax loss for discriminative face verification. *arXiv preprint arXiv:1703.09507* (2017) [4](#), [14](#)
49. Roth, K., Brattoli, B., Ommer, B.: Mic: Mining interclass characteristics for improved metric learning. In: *Proceedings of the IEEE International Conference on Computer Vision (ICCV)* (2019) [3](#), [11](#), [12](#)

50. Roth, K., Milbich, T., Ommer, B.: PADS: Policy-adapted sampling for visual similarity learning. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2020) [3](#), [11](#), [12](#)
51. Roth, K., Milbich, T., Ommer, B., Cohen, J.P., Ghassemi, M.: Simultaneous similarity-based self-distillation for deep metric learning. In: Proceedings of the 38th International Conference on Machine Learning (ICML) (2021) [3](#), [4](#)
52. Roth, K., Milbich, T., Sinha, S., Gupta, P., Ommer, B., Cohen, J.P.: Revisiting training strategies and generalization performance in deep metric learning. In: Proceedings of the 37th International Conference on Machine Learning (ICML) (2020) [1](#), [2](#), [3](#), [4](#), [9](#), [10](#), [12](#), [13](#)
53. Roth, K., Vinyals, O., Akata, Z.: Integrating language guidance into vision-based deep metric learning. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 16177–16189 (June 2022) [3](#)
54. Roth, K., Vinyals, O., Akata, Z.: Non-isotropy regularization for proxy-based deep metric learning. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 7420–7430 (June 2022) [3](#), [4](#)
55. Sanakoyeu, A., Tschernezki, V., Buchler, U., Ommer, B.: Divide and conquer the embedding space for metric learning. In: The IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2019) [3](#), [11](#), [12](#)
56. Schroff, F., Kalenichenko, D., Philbin, J.: Facenet: A unified embedding for face recognition and clustering. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2015) [1](#), [2](#), [3](#)
57. Scott, T.R., Gallagher, A.C., Mozer, M.C.: von Mises-Fisher loss: An exploration of embedding geometries for supervised learning. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV) (2021) [4](#), [5](#), [7](#), [12](#), [14](#)
58. Shi, Y., Jain, A.K.: Probabilistic face embeddings. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV) (2019) [4](#), [5](#), [7](#)
59. Sinha, S., Roth, K., Goyal, A., Ghassemi, M., Akata, Z., Larochelle, H., Garg, A.: Uniform priors for data-efficient learning. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops. pp. 4017–4028 (June 2022) [5](#)
60. Sohn, K.: Improved deep metric learning with multi-class n-pair loss objective. In: Advances in Neural Information Processing Systems (NeurIPS) (2016) [1](#), [3](#)
61. Sun, Y., Cheng, C., Zhang, Y., Zhang, C., Zheng, L., Wang, Z., Wei, Y.: Circle loss: A unified perspective of pair similarity optimization. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2020) [11](#)
62. Szegedy, C., Liu, W., Jia, Y., Sermanet, P., Reed, S., Anguelov, D., Erhan, D., Vanhoucke, V., Rabinovich, A.: Going deeper with convolutions. In: Computer Vision and Pattern Recognition (CVPR) (2015) [9](#)
63. Teh, E.W., DeVries, T., Taylor, G.W.: ProxyNCA++: Revisiting and revitalizing proxy neighborhood component analysis. In: Proceedings of the European Conference on Computer Vision (ECCV) (2020) [3](#), [4](#), [5](#), [7](#), [10](#), [13](#)
64. Wah, C., Branson, S., Welinder, P., Perona, P., Belongie, S.: The Caltech-UCSD birds-200-2011 dataset. Tech. Rep. CNS-TR-2011-001, California Institute of Technology (2011) [3](#), [9](#), [11](#)
65. Wang, J., Zhou, F., Wen, S., Liu, X., Lin, Y.: Deep metric learning with angular loss. In: Proceedings of the IEEE International Conference on Computer Vision (ICCV) (2017) [1](#)

66. Wang, T., Isola, P.: Understanding contrastive representation learning through alignment and uniformity on the hypersphere. In: Proceedings of the 37th International Conference on Machine Learning (ICML) (2020) [4](#)
67. Wang, X., Han, X., Huang, W., Dong, D., Scott, M.R.: Multi-similarity loss with general pair weighting for deep metric learning. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2019) [3](#), [4](#), [10](#)
68. Weisstein, E.W.: Hypersphere (2002) [4](#)
69. Wightman, R.: Pytorch image models. <https://github.com/rwightman/pytorch-image-models> (2019) [9](#)
70. Wu, C.Y., Manmatha, R., Smola, A.J., Krahenbuhl, P.: Sampling matters in deep embedding learning. In: Proceedings of the IEEE International Conference on Computer Vision (ICCV) (2017) [1](#), [2](#), [3](#), [4](#), [10](#), [11](#)
71. Xuan, H., Stylianou, A., Pless, R.: Improved embeddings with easy positive triplet mining. In: Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV) (March 2020) [3](#)
72. Xuan, H., Stylianou, A., Pless, R.: Improved embeddings with easy positive triplet mining. In: Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV) (March 2020) [11](#), [13](#)
73. Zhai, A., Wu, H.: Making classification competitive for deep metric learning. arXiv Preprint arXiv:1811.12649 (2018) [3](#), [4](#), [11](#)
74. Zhe, X., Chen, S., Yan, H.: Directional statistics-based deep metric learning for image classification and retrieval. *Pattern Recognition* **93** (2018) [4](#)
75. Zheng, W., Chen, Z., Lu, J., Zhou, J.: Hardness-aware deep metric learning. The IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2019) [3](#)
76. Zheng, W., Wang, C., Lu, J., Zhou, J.: Deep compositional metric learning. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2021) [11](#)
77. Zheng, W., Zhang, B., Lu, J., Zhou, J.: Deep relational metric learning. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV) (2021) [11](#), [13](#)
78. Zhu, Y., Yang, M., Deng, C., Liu, W.: Fewer is more: A deep graph metric learning perspective using fewer proxies. In: Larochelle, H., Ranzato, M., Hadsell, R., Balcan, M.F., Lin, H. (eds.) *Advances in Neural Information Processing Systems (NeurIPS)* (2020) [11](#)
79. Zimmermann, R.S., Sharma, Y., Schneider, S., Bethge, M., Brendel, W.: Contrastive learning inverts the data generating process. In: Proceedings of the 38th International Conference on Machine Learning (ICML) (2021) [5](#)