

Bayesian Metric Learning for Uncertainty Quantification in Image Retrieval

Warburg, Frederik; Miani, Marco; Brack, Silas; Hauberg, Søren

Published in: Proceedings of the 37th Conference on Neural Information Processing Systems

Publication date: 2024

Document Version Early version, also known as pre-print

Link back to DTU Orbit

Citation (APA):

Warburg, F., Miani, M., Brack, S., & Hauberg, S. (in press), Bayesian Metric Learning for Uncertainty Quantification in Image Retrieval. In *Proceedings of the 37th Conference on Neural Information Processing Systems* (Vol. 36). Neural Information Processing Systems Foundation.

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

• Users may download and print one copy of any publication from the public portal for the purpose of private study or research.

- · You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

Bayesian Metric Learning for Uncertainty Quantification in Image Retrieval

Frederik Warburg^{*1} Marco Miani^{*1} Silas Brack¹ Søren Hauberg¹

Abstract

We propose the first Bayesian encoder for metric learning. Rather than relying on neural amortization as done in prior works, we learn a distribution over the network weights with the Laplace Approximation. We actualize this by first proving that the contrastive loss is a valid log-posterior. We then propose three methods that ensure a positive definite Hessian. Lastly, we present a novel decomposition of the Generalized Gauss-Newton approximation. Empirically, we show that our Laplacian Metric Learner (LAM) estimates wellcalibrated uncertainties, reliably detects out-ofdistribution examples, and yields state-of-the-art predictive performance.

1. Introduction

Metric learning seeks data representations where similar observations are near and dissimilar ones are far. This construction elegantly allows for building retrieval systems with simple nearest-neighbor search. Such systems easily cope with a large number of classes, and new classes can organically be added without retraining. While these retrieval systems show impressive performance, they quickly, and with no raised alarms, deteriorate with out-of-distribution data (Shi & Jain, 2019). In particular, in safety-critical applications, the lack of uncertainty estimation is a concern as retrieval errors may propagate unnoticed through the system, resulting in erroneous and possibly dangerous decisions.

We present the Laplacian Metric Learner (LAM) to estimate reliable uncertainties of image embeddings as demonstrated in Fig. 1. LAM is the first Bayesian method proposed for metric learning. We learn a distribution over the network weights (weight posterior) from which we obtain a stochastic representation by embedding an image through sampled neural networks. The Bayesian formulation has multiple benefits, namely (1) robustness to out-of-distribution examples, (2) calibrated in-distribution uncertainties, and (3) a slight improvement in predictive performance.



Figure 1: **Reliable stochastic embeddings.** Current stateof-the-art, PFE (Shi & Jain, 2019), and LAM (ours) learn stochastic representations. LAM estimates reliable uncertainties of the latent representation that intuitively follow the amount of blur, noise, or occlusion in the input image.

Our method extends the Laplace Approximation (MacKay, 1992) for metric learning. We present a probabilistic interpretation of the contrastive loss (Hadsell et al., 2006) which justifies that it can be interpreted as an unnormalized negative log-posterior. We then propose three solutions to ensure a positive definite Hessian for the contrastive loss and present two approaches to compute the Generalized Gauss-Newton (Foresee & Hagan, 1997) approximation for ℓ_2 -normalized networks. Finally, we boost our method with the online training procedure from Miani et al. (2022) and achieve state-of-the-art performance.

We are not the first to consider uncertainty quantification in image retrieval. Seminal works (Shi & Jain, 2019; Oh et al., 2018) have addressed the lack of uncertainties in retrieval with *amortized inference* (Gershman & Goodman, 2014), where a neural network predicts a stochastic embedding. The issues with this approach are that (1) it requires strong assumptions on the distribution of the embedding, (2) the networks are often brittle and difficult to optimize, and (3) out-of-distribution detection relies on the network's capacity to extrapolate uncertainties. As neural networks extrapolate poorly (Xu et al., 2021), the resulting *predicted* uncertainties

^{*}Equal contribution ¹Technical University of Denmark. Correspondence to: Frederik Warburg <frwa@dtu.dk>.

Code at https://github.com/FrederikWarburg/
bayesian-metric-learning

are unreliable for out-of-distribution data (Detlefsen et al., 2019) and are thus, in practice, of limited value.

In contrast, our method does not assume any distribution on the stochastic embeddings, is simple to optimize, and does not rely on a neural network to extrapolate uncertainties. Instead, our weight posterior is derived from the curvature of the loss landscape and the uncertainties of the latent embeddings deduced (rather than learned) with sampling. We show through rigorous experiments that this leads to reliable outof-distribution performance and calibrated uncertainties in both controlled toy experiments and challenging real-world applications such as bird, face, and place recognition.

2. Related Work

Metric learning attempts to map data to an embedding space, where similar data are close together and dissimilar data are far apart. This is especially useful for retrieval tasks with many classes and few observations per class such as place recognition (Warburg et al., 2020) and face recognition (Schroff et al., 2015) or for tasks where classes are not well-defined, such as food tastes (Wilber et al., 2022).

There exist many metric losses that optimize for a wellbehaved embedding space. We refer to the excellent survey by Musgrave et al. (2020) for an overview. We here focus on the *contrastive loss* (Hadsell et al., 2006)

$$\mathcal{L}_{con}(\theta) = \frac{1}{2} \| f_{\theta}(x_a) - f_{\theta}(x_p) \|^2 + \frac{1}{2} \max\left(0, m - \| f_{\theta}(x_a) - f_{\theta}(x_n) \|^2\right),$$
(1)

which has shown state-of-the-art performance (Musgrave et al., 2020) and is one of the most commonly used metric losses. Here, f_{θ} is a neural network parametrized by θ which maps from the observation space to the embedding space. The loss consists of two terms, one that attracts observations from the same class (*anchor* x_a and *positive* x_p), and one that repels observations from different classes (*anchor* x_a and *negative* x_n). The margin *m* ensures that negatives are repelled sufficiently far. We will later present a probabilistic extension of the contrastive loss that allows us to learn stochastic, rather than deterministic, features in the embedding space.

Uncertainty in deep learning is currently studied across many domains to mitigate fatal accidents and allow for human intervention when neural networks make erroneous predictions. Current methods can be divided into methods that apply amortized optimization to train a neural network to predict the parameters of the output distribution, and methods that do not. The amortized methods, best known from the variational autoencoder (VAE) (Kingma & Welling, 2013; Rezende et al., 2014), seem attractive at first as they can directly estimate the output distribution (without requiring sampling), but they suffer from mode collapse and are sensitive to out-of-distribution data due to the poor extrapolation capabilities of neural networks (Nalisnick et al., 2018; Detlefsen et al., 2019). 'Bayes by Backprop' (Blundell et al., 2015) learns a distribution over parameters variationally but is often deemed too brittle for practical applications. Alternatives to amortized methods includes deep ensembles (Lakshminarayanan et al., 2017), stochastic weight averaging (SWAG) (Maddox et al., 2019), Monte-Carlo dropout (Gal & Ghahramani, 2016) and Laplace Approximation (LA) (Laplace, 1774; MacKay, 1992) which all approximate the generally intractable weight posterior $p(\theta|D)$ of a neural network. We propose the first principled method to approximate the weight posterior in metric learning.

Laplace approximations (LA) can be applied for every loss function \mathcal{L} that can be interpreted as an unnormalized log-posterior by performing a second-order Taylor expansion around a chosen weight vector θ^* such that

$$\mathcal{L}(\theta) \approx \mathcal{L}^* + (\theta - \theta^*)^\top \nabla \mathcal{L}^* + \frac{1}{2} (\theta - \theta^*)^\top \nabla^2 \mathcal{L}^* (\theta - \theta^*),$$
⁽²⁾

where \mathcal{L}^* is the loss evaluated in θ^* . Imposing the unnormalized log-posterior to be a second-order polynomial is equivalent to assuming the posterior to be Gaussian. If θ^* is a MAP estimate, the first-order term vanishes, and the second-order term can be interpreted as a precision matrix, the inverse of the covariance. Assuming θ^* is a MAP estimate, this second-order term is negative semi-definite for common (convex) supervised losses, such as the mean-squared error and cross-entropy. Recently, Daxberger et al. (2021) demonstrated that post-hoc LA is scalable and produces well-behaved uncertainties for classification and regression. The Laplacian Autoencoder (LAE) (Miani et al., 2022) improves on the post-hoc LA with an online Monte Carlo EM training procedure to learn a well-behaved posterior. It demonstrates state-of-the-art uncertainty quantification for unsupervised representation learning. We first extend LA to the contrastive loss and achieve state-of-the-art performance with the online EM training procedure.

Uncertainty in metric learning is not a new idea (Vilnis & McCallum, 2014), but the large majority of recent methods apply amortized inference to predict distributions in the embedding space (Warburg et al., 2021; Chang et al., 2020; Chun et al., 2021; Oh et al., 2018; Shi & Jain, 2019; Song & Soleymani, 2019; Sun et al., 2020), making them sensitive to mode collapse and out-of-distribution data. Taha et al. (2019b;a) explore deep ensembles and Monte-Carlo dropout as alternatives, but these methods suffer from increased training time, poor empirical performance, and limited Bayesian interpretation (Daxberger et al., 2021). We explore LA in metric learning and attain state-of-the-art performance.



Figure 2: **Model overview**. We learn a distribution over parameters, such that we embed an image through sampled encoders f_{θ} to points z_i (red dots) in a latent space \mathcal{Z} . We reduce these latent samples to a single measure of uncertainty by estimating the parameters of a von Mises-Fisher distribution.

3. Laplacian Metric Learning

To perform Bayesian retrieval, we propose to estimate the weight posterior of the embedding network f_{θ} such that we can sample data embeddings to propagate uncertainty through the decision process. The embedding network is parametrized by $\theta \in \Theta$ and trained with the contrastive loss. The network maps an image $x \in \mathcal{X} := \mathbb{R}^{HWC}$ to an embedding $z \in \mathcal{Z}$, which is restricted to be on a Z-dimensional unit sphere $\mathcal{Z} := \mathcal{S}^Z$. This normalization to the unit sphere is commonly done in retrieval to obtain faster retrieval and a slight performance boost (Arandjelović et al., 2016; Radenović et al., 2018). Fig. 2 illustrate our Bayesian mapping from image to latent space.

To obtain an approximate posterior over the weights θ we rely on the Laplace approximation (LA). We first motivate the post-hoc LA as this is the simplest. We then extend this approach to online LA which marginalizes the Laplace approximation during training to improve model fit. In Appendix D, we prove that the contrastive loss is a valid unnormalized log-posterior on the compact spherical space. The proof draws inspiration from electrostatics, and the main idea is to define a PDF for the attracting and repelling terms. We then show that the logarithm of the product of these PDFs is equivalent (up to some constant) to the contrastive loss (details are postponed to Appendix D). Since the contrastive loss is a valid log-posterior, we can proceed with LA.

A post-hoc Laplace approximation is found by first training a standard deterministic network through gradient steps with the contrastive loss to find the *maximum a posteriori* (MAP) parameters θ^* . Since we are in a local optimum, the first-order term in the second-order Taylor expansion (Eq. 2) vanishes, and we can define the parameter distribution as

$$p(\theta|\mathcal{D}) = \mathcal{N}\left(\theta \middle| \theta^*, \left(\nabla_{\theta}^2 \mathcal{L}_{\text{con}}\left(\theta^*; \mathcal{D}\right) + \sigma_{\text{prior}}^{-2} \mathbb{I}\right)^{-1}\right).$$
(3)

The advantage of post-hoc LA is that the training procedure does not change, and already trained neural networks can be made Bayesian. In practice, however, stochastic gradientbased training does not locate isolated minima of the loss landscape, but rather ends up exploring regions near local minima. The Hessian (and hence the posterior covariance) can change significantly during this exploration, and the post-hoc LA can become unstable.

Online Laplace approximations (Miani et al., 2022) avoids this instability by marginalizing the LA during training with Monte Carlo EM. This helps the training recover a solution θ^* where the Hessian reflects the loss landscape. Specifically, at each step t during training we keep in memory a Gaussian distribution on the parameters $q^t(\theta) = \mathcal{N}(\theta | \theta_t, H_{\theta_t}^{-1})$. The parameters are updated through an expected gradient step

$$\theta_{t+1} = \theta_t + \lambda \mathbb{E}_{\theta \sim q^t} [\nabla_\theta \mathcal{L}_{\text{con}}(\theta; \mathcal{D})]$$
(4)

and a discounted Laplace update

$$H_{\theta_{t+1}} = (1 - \alpha)H_{\theta_t} + \nabla^2_{\theta}\mathcal{L}_{\text{con}}(\theta; \mathcal{D}), \qquad (5)$$

where α describes an exponential moving average, similar to momentum-like training. The initialization follows the isotropic prior $q^0(\theta) = \mathcal{N}(\theta|0, \sigma_{\text{prior}}^2 \mathbb{I})$.

In practice, the Hessian scales quadratically in memory wrt. the number of parameters. To mitigate this, we approximate this Hessian by its diagonal (LeCun et al., 1989; Denker & LeCun, 1990).

Hessian of the contrastive loss. Both post-hoc and online LA require the Hessian of the contrastive loss $\nabla_{\theta}^{2} \mathcal{L}_{con}(\theta; \mathcal{D})$. The Hessian is commonly approximated with the Generalized Gauss-Newton (GGN) approximation (Foresee & Hagan, 1997; Daxberger et al., 2021; Dangel et al., 2020; Detlefsen et al., 2021). The GGN decomposes the loss into $\mathcal{L} = g \circ f$, where g is usually chosen as the loss function and f the model function, and only f is linearized.

However, in our case, this decomposition is non-trivial. Recall that the last layer of our network is an ℓ_2 normalization layer, which projects embeddings onto a hyper-sphere. This normalization layer can either be viewed as part of the model f (linearized normalization layer) or part of the loss g (nonlinearized normalization layer). We show in Appendix F.1 that the former can be interpreted as using the *Euclidean* distance and in Appendix F.2 that the latter as using the *Arccos* distance for the contrastive loss (Eq. 1). We highlight that these share the zero- and first-order terms for normalized embeddings but, due to the GGN linearization, not the second-order derivatives. The Euclidean interpretation leads to simpler derivatives and interpretations, and we will therefore use it for our derivations. We emphasize that the Arccos is theoretically a more accurate approximation, because the



Figure 3: Hessian approximations. To ensure a positive definite Hessian approximation we propose three approximations. In (a) only the positives p contribute to the Hessian as the negatives n are ignored. In (b) we consider one point at a time, e.g., only the anchor a contributes. In (c) we consider all interactions.

 ℓ_2 -layer is not linearized, and we provide derivations in Appendix F.

The GGN matrix for contrastive loss with the *Euclidean* interpretation is given by

$$\nabla_{\theta}^{2} \mathcal{L}_{con}(\theta; \mathcal{I}) = \sum_{ij \in \mathcal{I}} H_{\theta}^{ij} = \sum_{ij \in \mathcal{I}_{p}} H_{\theta}^{ij} + \sum_{ij \in \mathcal{I}_{n}} H_{\theta}^{ij} \quad (6)$$

$$\stackrel{\text{GGN}}{\approx} \sum_{ij \in \mathcal{I}_{p}} J_{\theta}^{ij^{\top}} \underbrace{\begin{pmatrix} 1 & -1 \\ -1 & 1 \end{pmatrix}}_{:=H_{p}} J_{\theta}^{ij} + \sum_{ij \in \mathcal{I}_{n}} J_{\theta}^{ij^{\top}} \underbrace{\begin{pmatrix} -1 & 1 \\ 1 & -1 \end{pmatrix}}_{:=H_{n}} J_{\theta}^{ij},$$

where $J_{\theta}^{ij} = (J_{\theta}f_{\theta}(x_i)^{\top}, J_{\theta}f_{\theta}(x_j)^{\top})^{\top}$, with J_{θ} being the Jacobian wrt. the parameters, and where H_p and H_n are the Hessian of the contrastive loss wrt. the model output for positive and negative pairs. Notice that the first sum runs over positive pairs and the second sum runs over negative pairs within the margin. Negative pairs outside the margin do not contribute to the Hessian, and can therefore be ignored to reduce the computational load (Appendix B).

The eigenvalues of the Hessian wrt. to the output are (0, 2)and (-2, 0) for the positive H_p and negative H_n terms, so we are not guaranteed to have a positive definite Hessian, H_{θ} . To avoid covariances with negative eigenvalues, we propose three solutions to ensure a positive definite Hessian. Proofs are in Appendix F.3.

Ensuring positive definiteness of the Hessian. We do not want to be restricted in the choice of the prior, so we must ensure that $\nabla_{\theta}^2 \mathcal{L}_{con}(\theta^*; \mathcal{D})$ is positive definite itself. Differently from the standard convex losses, this is not ensured by the GGN approximation (Immer et al., 2021). Our main insight is that we can ensure a positive definite Hessian H_{θ} by only manipulating the Hessians H_p and H_n in Eq. 6.

1. Positive: The repelling term is ignored, such that only positive pairs contribute to the Hessian.

$$H_p = \begin{pmatrix} 1 & -1 \\ -1 & 1 \end{pmatrix}, \qquad H_n = \begin{pmatrix} 0 & 0 \\ 0 & 0 \end{pmatrix}$$
(7)

2. Fixed: The cross derivatives are ignored.

$$H_p = \begin{pmatrix} 1 & 0\\ 0 & 1 \end{pmatrix}, \qquad H_n = \begin{pmatrix} -1 & 0\\ 0 & -1 \end{pmatrix}$$
(8)

3. Full: Nothing is ignored, but rather positive definiteness is ensured with ReLU, $\max(0, \nabla^2_{\theta} \mathcal{L}_{con}(\theta; \mathcal{D}))$, on the Hessian of the loss wrt. the parameters.

$$H_p = \begin{pmatrix} 1 & -1 \\ -1 & 1 \end{pmatrix}, \qquad H_n = \begin{pmatrix} -1 & 1 \\ 1 & -1 \end{pmatrix}$$
(9)

The *positive* approximation is inspired by Shi & Jain (2019). which also only uses positive pairs to train their uncertainty module. The gradient arrows in Fig. 3 (a) illustrate that negative pairs are neglected when computing the Hessian of the contrastive loss. The *fixed* approximation considers one data point at the time, assuming the other one is fixed. Thus, given a pair of data points, this can be interpreted as first moving one data point, and then the second (rather than both at the same time). We formalize this in Appendix F.3. Fig. 3 (b) illustrate this idea when all points except a are fixed. Lastly, we propose the *full* Hessian of the contrastive loss (Fig. 3 (c)) and ensure a positive definiteness by computing the ReLU of the Hessian. We note that this approximation assumes a diagonal Hessian. We experimentally determine, which of these three methods to ensure a positive definiteness leads to better performance (Section 4).

Intuition of Hessian approximations. With Fig. 4 we try to provide some intuition on these three approximations and their impact in a clustered versus cluttered latent space. Positive points will increase the precision (magnitude of the Hessian) while negative points inside the margin will decrease the precision (as indicated with the arrows in Fig. 4). Negative points outside the margin will not affect the gradient or the Hessian. This is the desired behavior as a perfect clustered latent space (Fig. 4 b) will have high precision (low uncertainty), whereas in a cluttered latent space (Fig. 4 a) will have many negatives within the margin, which will decrease the precision (large uncertainty). Therefore, the positive approximation (Eq. 7) will provide the lowest uncertainty, as the negatives are excluded. In practice, we cannot compute the Hessian for all pairs, as this scales quadratically in time and memory with the dataset size. Furthermore, most pairs, namely the negatives outside the margin, do not contribute to the Hessian, so it is wasteful to compute their Hessian. Therefore, we make use of a common trick in metric learning, namely hard negative mining (Musgrave et al., 2020).

Hard negative mining leads to biased Hessian approximation. Metric losses operate locally, hence non-zero losses will only occur for negative examples that are close to the anchor in the latent space. Presenting the model for



Figure 4: Intuition of H_p and H_n . Illustration of a cluttered latent space, where observations from different classes are close, and a clustered latent space with distinct clusters with observations from the same class. Intuitively, negative examples \star inside the margin decreases the precision (higher variance) with H_n , and positive points • will increase the precision (lower variance) with H_p .

randomly sampled data will in practice rarely result in negative pairs that are close, which leads to extremely long training times. Therefore, *hard negative mining* is often used to construct batches that have an over-representation of hard negative pairs (negative examples that are close to the anchor). We use similar mining when computing the Hessian, which leads to a biased estimate of the Hessian. We thus introduce a scaling parameter $0 \le w_n \le 1$ to obtain

$$\hat{H}_{\theta} = (1 - w_n) J_{\theta}^{\top} H_p J_{\theta} + w_n J_{\theta}^{\top} H_n J_{\theta}, \qquad (10)$$

which corrects the biased estimate (see Appendix B.1 for proof). Our positive approximation (Eq. 7) can be seen as the extreme case with $w_n = 0$, whereas our full approximation Eq. 9 ($w_n = 0.5$) corresponds to unbiased sampling.

Having an estimate of the Hessian in place and three methods to ensure it is positively definite, we can perform both post-hoc and online LA to obtain a distribution over the weights. We can sample from this weight posterior and embed an input image via each sampled network to obtain multiple samples in the latent space (see Fig. 2). We reduce these samples to a single measure of uncertainty by estimating the parameters of a von Mises-Fisher distribution.

Why the von Mises-Fisher distribution? The uncertainty deduced from LA is usually computed by the variance of the samples (Daxberger et al., 2021). This assumes that the samples follow a Gaussian distribution. However, for ℓ_2 -normalized networks, we assume that all the probability mass lies on a Z-dimensional hyper-sphere. The von Mises-Fisher distribution describes such distribution, and it is parametrized with a directional mean μ and a scalar concentration parameter κ , which can be interpreted as the inverse of an isotropic covariance $\kappa = 1/\sigma^2$, i.e., small κ means high uncertainty and large κ means low uncertainty.

There exist several methods to estimate κ . We opt for the simplest and most computationally efficient (Sra, 2012) (see Appendix C). Prior works (Shi & Jain, 2019; Taha et al., 2019b;a) have treated an estimated latent distribution as

a Gaussian, although all probability mass lies on the unit sphere. This is insufficient, as samples from the Gaussian distribution will not lie on the unit sphere. In the experimental work, we correct this by projecting the samples onto the unit sphere.

4. Experiments

We benchmark our method against strong probabilistic retrieval models. Probabilistic Face Embeddings (PFE) (Shi & Jain, 2019) and Hedge Image Embedding (HIB) (Oh et al., 2018) perform amortized inference and thus estimate the mean and variance of latent observation. We also compare against MC Dropout (Gal & Ghahramani, 2016) and Deep Ensemble (Lakshminarayanan et al., 2017), two approximate Bayesian methods, which have successfully been applied in image retrieval (Taha et al., 2019b;a).

We compare the models' *predictive performance* with the recall (recall@k) and mean average precision (mAP@k) among the k nearest neighbors (Warburg et al., 2021; Musgrave et al., 2020; Arandjelović et al., 2016). We evaluate the models' abilities to *interpolate* and *extrapolate* uncertainties by measuring the Area Under the Sparsification Curve (AUSC) and Expected Calibration Error (ECE) on in-distribution (ID) data, and the Area Under Receiver Operator Curve (AUROC) and Area Under Precision-Recall Curve (AUPRC) on out-of-distribution (OOD) data. We provide more details on these metrics in Appendix G.3.

We extend StochMan (Detlefsen et al., 2021) with the Hessian backpropagation for the contrastive loss. The training code is implemented in PyTorch (Paszke et al., 2017) and is available on GitHub¹. Appendix G provides more details on the experimental setup.

Experimental Summary. We begin with a short summary of our experimental results. Across five datasets, three different network architectures, and three different sizes of the latent space (ranging from 3 to 2048), we find that LAM has well-calibrated uncertainties, reliably detects OOD examples, and achieves state-of-the-art predictive performance. Fig. 5a shows that the uncertainties from online LAM reliably identify OoD examples. Online LAM outperforms other Bayesian methods, such as post-hoc LAM and MC dropout, on this task, which in turn clearly improves upon amortized methods that rely on a neural network to extrapolate uncertainties. Fig. 5b shows that LAM consistently matches or outperforms existing image retrieval methods in terms of predictive performance. We find that the fixed Hessian approximation with the Arccos distance performs the best, especially on higher dimensional data.

¹Code: https://github.com/FrederikWarburg/ bayesian-metric-learning

Bayesian Metric Learning for Uncertainty Quantification in Image Retrieval



Figure 5: Summary of experimental results. LAM consistently outperforms existing methods on OOD detection, as measured by AUROC, and matches or surpasses in predictive performance measured by mAP@k. Error bar shows \pm one standard deviation measured across five runs.

Ablation: Positive Definiteness of the Hessian. We experimentally study which method to ensure a positive definite Hessian has the best performance measured in both predictive performance (mAP@5) and uncertainty quantification (AUROC, AUSC). We found that all these methods perform similarly on simple datasets and low dimensional hyper-spheres, but the fixed approximation with Arccos distance performs better on more challenging datasets and higher dimensional hyper-spheres. We present results on one of these more challenging datasets, namely the LFW (Huang et al., 2007) face recognition dataset with the CUB200 (Wah et al., 2011) bird dataset as an OOD dataset. We use a ResNet50 (He et al., 2016) with a GeM pooling layer (Radenović et al., 2018) and a 2048 dimensional embedding and diagonal, last-layer LA (Daxberger et al., 2021).

Table 1 shows the performance for post-hoc and online LA with fixed, positive, or full Hessian approximation using either Euclidean or Arccos distance. Across all metrics, the online LA with Arccos distance and the fixed Hessian approximation performs similarly or the best. We proceed to benchmark this method against several strong probabilistic baselines on closed-set retrieval and a more challenging open-set retrieval.

Closed-Set Retrieval. OOD capabilities are critical for identifying distributional shifts, outliers, and irregular user inputs, which might hinder the propagation of erroneous decisions in an automated system. We evaluate OOD performance on the commonly used benchmarks (Nalisnick et al., 2018), where we use (1) FashionMNIST (Xiao et al., 2017) as ID and MNIST (LeCun et al., 1998) as OOD, and (2) CI-FAR10 (Krizhevsky, 2009) as ID and SVHN (Netzer et al., 2011) as OOD. We use, respectively, a standard 2- or 3-layer relu convolutional network followed by a single linear layer on which we compute LA with a diagonal Hessian.

FashionMNIST (ID) vs MNIST (OOD). Table 2 shows that both PFE and post-hoc LAM have a similar predictive performance to the deterministic model. This is not surprising,

Table 1: Ablation on Hessian approximation and GGN decomposition. Online LA with the fixed approximation and Arccos distance performs best. Error bars show one std. deviation across five runs.

		mAP@ $5\uparrow$	AUROC \uparrow	AUSC \uparrow
్ల	Euclidean fix	0.70 ± 0.0	0.57 ± 0.25	0.44 ± 0.01
	Euclidean pos	0.70 ± 0.0	0.58 ± 0.23	0.45 ± 0.01
-hc	Euclidean full	0.70 ± 0.0	0.56 ± 0.26	0.44 ± 0.01
ost	Arccos fix	0.69 ± 0.0	0.53 ± 0.20	0.46 ± 0.02
Ч	Arccos pos	0.70 ± 0.0	0.29 ± 0.11	0.48 ± 0.01
	Arccos full	0.69 ± 0.0	0.55 ± 0.18	0.45 ± 0.01
Online	Euclidean fix	0.63 ± 0.01	0.77 ± 0.04	0.31 ± 0.02
	Euclidean pos	0.70 ± 0.0	0.38 ± 0.10	0.47 ± 0.01
	Euclidean full	0.67 ± 0.01	0.59 ± 0.04	0.42 ± 0.01
	Arccos fix	$\textbf{0.71} \pm \textbf{0.0}$	$\textbf{0.78} \pm \textbf{0.18}$	0.50 ± 0.03
	Arccos pos	0.70 ± 0.0	0.23 ± 0.03	0.46 ± 0.00
	Arccos full	$\textbf{0.71} \pm \textbf{0.0}$	0.70 ± 0.12	$\textbf{0.51} \pm \textbf{0.01}$

as both methods are initialized with the deterministic parameters, and then uncertainties are learned (PFE) or deduced (post-hoc LAM) with frozen weights. The awareness of uncertainties during training, grants the online LAM slightly higher predictive performance.

PFE uses amortized inference to predict variances. This works reasonably within distribution but does not work well for OoD detection. This is because a neural network is trusted to extrapolate far away from the data distribution. Table 2 shows that MC dropout, LAM (post-hoc), and LAM (online) assign high uncertainty to observations outside the training distribution. Fig. 7 shows that both post-hoc and online LAM are near perfectly calibrated, which leads to very low ECE measure (Table 2).

Cifar10 (ID) vs SVHN (OOD) is a slightly harder setting. Table 2 yields similar conclusions as before; Bayesian approaches such as MC dropout, LAM (post-hoc), and LAM (online) are much better at detecting OoD examples than neural amortized methods such as PFE. Online LAM has a similar predictive performance to state-of-the-art while having better ID (lower ECE and higher AUSC) and OoD

Table 2: Closed-set results. LAM matches or outperforms existing methods in terms of predictive performance. It produces
reliable uncertainties ID and OOD on two standard datasets FashionMNIST and CIFAR10. Confidence intervals show one
standard deviation computed across five runs.

		Image retrieval			00	D	CALIBRATION			
		mAP@1↑	mAP@5 \uparrow	mAP@10 \uparrow	AUROC ↑	AUPRC \uparrow	AUSC ↑	$\text{ECE}\downarrow$		
FashionMNIST	Deterministic Deep Ensemble PFE HIB MC dropput	$\begin{array}{c} 0.78 \pm 0.01 \\ 0.69 \\ 0.78 \pm 0.00 \\ 0.69 \pm 0.08 \\ 0.76 \pm 0.02 \end{array}$	$\begin{array}{c} 0.73 \pm 0.01 \\ 0.62 \\ 0.74 \pm 0.00 \\ 0.63 \pm 0.09 \\ 0.71 \pm 0.02 \end{array}$	$\begin{array}{c} 0.72 \pm 0.01 \\ 0.59 \\ 0.72 \pm 0.00 \\ 0.61 \pm 0.09 \\ 0.70 \pm 0.02 \end{array}$	$\begin{array}{c} - \\ 0.41 \\ 0.53 \pm 0.03 \\ 0.60 \pm 0.12 \\ 0.02 + 0.02 \end{array}$	$\begin{array}{c}\\ 0.46\\ 0.46\pm 0.01\\ 0.60\pm 0.11\\ 0.02\pm 0.02\\ \end{array}$	$\begin{array}{c}\\ 0.61\\ 0.65\pm 0.01\\ 0.65\pm 0.08\\ 0.04\pm 0.06\end{array}$	$\begin{array}{c}\\ 0.04\\ 0.26 \pm 0.02\\ 0.54 \pm 0.08\\ 0.02 + 0.04 \end{array}$		
	LAM (post-hoc) LAM (online)	$\begin{array}{c} 0.76 \pm 0.03 \\ 0.78 \pm 0.00 \\ \textbf{0.81} \pm \textbf{0.00} \end{array}$	0.71 ± 0.03 0.74 ± 0.00 0.77 ± 0.01	0.70 ± 0.03 0.72 ± 0.00 0.76 ± 0.01	$\begin{array}{c} 0.93 \pm 0.03 \\ 0.96 \pm 0.02 \\ \textbf{0.98} \pm \textbf{0.01} \end{array}$	0.93 ± 0.03 0.96 ± 0.02 0.98 ± 0.01	$\begin{array}{c} 0.84 \pm 0.06 \\ 0.86 \pm 0.01 \\ \textbf{0.89} \pm \textbf{0.01} \end{array}$	$\begin{array}{c} 0.03 \pm 0.04 \\ 0.03 \pm 0.00 \\ \textbf{0.02} \pm \textbf{0.00} \end{array}$		
Cifar10	Deterministic Deep Ensemble MC dropout HIB PFE LAM (post-hoc) LAM (online)	$\begin{array}{c} \textbf{0.66} \pm \textbf{0.00} \\ \textbf{0.66} \\ 0.46 \pm 0.01 \\ 0.11 \pm 0.01 \\ \textbf{0.66} \pm \textbf{0.00} \\ \textbf{0.66} \pm \textbf{0.00} \\ \textbf{0.66} \pm \textbf{0.01} \end{array}$	$\begin{array}{c} 0.59 \pm 0.00 \\ \textbf{0.61} \\ 0.37 \pm 0.01 \\ 0.07 \pm 0.00 \\ 0.60 \pm 0.00 \\ 0.60 \pm 0.00 \\ 0.60 \pm 0.00 \end{array}$	$\begin{array}{c} 0.58 \pm 0.00 \\ \textbf{0.59} \\ 0.34 \pm 0.01 \\ 0.05 \pm 0.00 \\ 0.58 \pm 0.00 \\ 0.58 \pm 0.00 \\ 0.57 \pm 0.01 \end{array}$	$ \begin{array}{c} \\ 0.42 \\ 0.60 \pm 0.03 \\ 0.44 \pm 0.17 \\ 0.21 \pm 0.02 \\ 0.50 \pm 0.11 \\ \textbf{0.78} \pm \textbf{0.04} \end{array} $		$\begin{array}{c}\\ 0.72\\ 0.61 \pm 0.01\\ 0.29 \pm 0.03\\ 0.56 \pm 0.01\\ 0.81 \pm 0.01\\ \textbf{0.83} \pm \textbf{0.01} \end{array}$			





Figure 6: **Calibration Curves.** LAM is near perfectly calibrated on FashionMNIST and CIFAR10.

(higher AUROC and AUPRC) performance. Fig. 7 shows the calibration plot for CIFAR10. For this dataset, online LAM has a near-perfect calibration curve. Fig. 6 shows the ROC curves for CIFAR10 and highlights that online LAM is better at distinguishing ID and OOD examples.

Open-Set Retrieval. A key advantage of metric learning methods is that they easily cope with a large number of classes and new classes can be added seamlessly. We therefore evaluate LAM's performance on challenging open-set retrieval, where none of the classes in the test set are available during training. We first test with CUB200 (Wah et al., 2011) as ID and CAR196 (Krause et al., 2013) as OoD similarly to Warburg et al. (2021), and second, test with LFW (Huang et al., 2007) as ID and CUB200 as OoD. We use a ResNet50 (He et al., 2016) with a GeM pooling layer (Radenović et al., 2018) and a 2048 dimensional embedding and diagonal, last-layer LA (Daxberger et al., 2021).

CUB200 (ID) vs CARS196 (OoD). The CUB-200-2011 dataset (Wah et al., 2011) has 200 bird species captured

Figure 7: **Receive Operator Curves** LAM assign high uncertainty to OOD observations.

from different perspectives and in different environments. We follow the zero-shot train/test split (Musgrave et al., 2020). In this zero-shot setting, the trained models have not seen any of the bird species in the test set, and the learned features must generalize well across species. Table 3 shows that LAM matches or surpasses the predictive performance of all other methods. LAM (post-hoc) achieves state-of-the-art predictive performance, while LAM (online) matches the predictive performance of the deterministic trained model while achieving state-of-the-art AUROC and AUPRC for OoD detection.

LFW (ID) vs CUB200 (OOD). Face recognition is another challenging metric learning task with many applications in security and surveillance. The goal is to retrieve images of the same person as in the query image. Table 3 shows that online LAM outperforms existing methods both in terms of predictive performance and uncertainty quantification. Fig. 8 shows that PFE assigns higher uncertainty to images from the ID dataset (faces) than those from the OOD dataset (birds). In contrast, both online and post-hoc LAM better associate high variance to OOD examples, while PFE predicts

		IN	AGE RETRIEV	AL	00	ID	
		mAP@1↑	mAP@5↑	mAP@10↑	AUROC ↑	AUPRC ↑	AUSC ↑
	Deterministic	0.62 ± 0.01	0.48 ± 0.01	0.42 ± 0.01	—		
_	Deep Ensemble	0.21	0.11	0.07	0.47	0.55	0.21
500	PFE	0.62 ± 0.01	0.5 ± 0.01	0.43 ± 0.01	0.44 ± 0.16	0.5 ± 0.08	0.61 ± 0.02
CUB3	HIB	0.33 ± 0.04	0.19 ± 0.02	0.14 ± 0.02	0.54 ± 0.12	0.61 ± 0.1	0.31 ± 0.07
	MC dropout	0.61 ± 0.00	0.48 ± 0.00	0.42 ± 0.00	0.73 ± 0.08	0.68 ± 0.07	0.63 ± 0.01
	LAM (post-hoc)	$\textbf{0.65} \pm \textbf{0.01}$	$\textbf{0.52} \pm \textbf{0.01}$	$\textbf{0.45} \pm \textbf{0.01}$	0.56 ± 0.16	0.61 ± 0.11	$\textbf{0.66} \pm \textbf{0.03}$
	LAM (online)	0.61 ± 0.00	0.48 ± 0.00	0.42 ± 0.00	$\textbf{0.80} \pm \textbf{0.03}$	$\textbf{0.75} \pm \textbf{0.03}$	0.63 ± 0.01
LFW	Deterministic	0.44 ± 0.00	0.68 ± 0.00	0.65 ± 0.00	_	_	
	Deep Ensemble	0.36	0.57	0.54	0.52	0.64	0.33
	PFE	0.44 ± 0.00	0.68 ± 0.00	0.65 ± 0.00	0.03 ± 0.02	0.41 ± 0.0	0.49 ± 0.01
	MC dropout	0.42 ± 0.00	0.65 ± 0.01	0.63 ± 0.01	0.03 ± 0.01	0.41 ± 0.0	0.46 ± 0.01
	LAM (post-hoc)	0.44 ± 0.01	0.68 ± 0.01	0.65 ± 0.00	0.65 ± 0.14	0.72 ± 0.11	0.45 ± 0.03
	LAM (online)	$\textbf{0.46} \pm \textbf{0.00}$	$\textbf{0.71} \pm \textbf{0.00}$	$\textbf{0.69} \pm \textbf{0.00}$	$\textbf{0.71} \pm \textbf{0.22}$	$\textbf{0.78} \pm \textbf{0.17}$	$\textbf{0.50} \pm \textbf{0.02}$

Table 3: **Open-set results.** LAM matches or outperforms existing methods in terms of predictive performance and produces state-of-the-art uncertainty quantification for challenging zero-shot metric learning datasets LFW and CUB200. Confidence intervals show one standard deviation computed across five runs.



Figure 8: **Images with lowest and highest variance** for PFE, post-hoc LAM, and online LAM across LFW (ID) and CUB200 (OOD) datasets. LAM associates high uncertainty to OOD examples, whereas PFE predicts higher uncertainties for ID images. Shows the best-performing PFE and online LAM across five runs.

high variance to ID examples. Furthermore, online LAM seems to assign the highest variance to images in which the background is complex and thus camouflages the birds.

4.1. Visual Place recognition

Lastly, we evaluate LAM on the challenging task of visual place recognition, which has applications that span from human trafficking investigation (Stylianou et al., 2019) to the long-term operation of autonomous robots (Davison et al., 2007). Our focus is the latter, where the goal is to retrieve images taken within a radius of 25 meters from a query image. The high number of unique places and varying visual appearance of each location – including weather, dynamic, structural, view-point, seasonal, and day/night changes – makes visual place recognition a very challenging metric learning problem. Reliable uncertainties and reliable out-

of-distribution behavior are important to avoid incorrect loop-closure, which can deteriorate the autonomous robots' location estimate. We evaluate on MSLS (Warburg et al., 2020), which is the largest and most diverse place recognition dataset currently available comprised of 1.6M images from 30 cities spanning six continents. We use the standard train/test split, training on 24 cities and testing on six other cities. We use the same model as in open-set retrieval.

Table 4 shows that online LAM yields state-of-the-art uncertainties for visual place recognition measured with AUSC, while matching the predictive performance of the alternative probabilistic and deterministic methods on both the MSLS validation and the challenge set. Fig. 10 shows the sparsification curves on the challenge set. Both online and posthoc LAM have monotonically increasing sparsification curves, implying that when we remove the most uncertain observations, the predictive performance increase. This il-

	Validation Set						Challenge Set					
	R@1↑	R@5↑	R@10↑	M@5↑	M@10 \uparrow	AUSC↑	R@1↑	R@5↑	R@10↑	M@5↑	M@10 \uparrow	AUSC↑
Deterministic	0.77	0.88	0.90	0.61	0.56		0.58	0.74	0.78	0.45	0.43	
MC Dropout	0.75	0.87	0.87	0.59	0.54	0.77	0.55	0.71	0.76	0.43	0.41	0.57
PFE	0.77	0.88	0.90	0.61	0.56	0.73	0.58	0.74	0.78	0.45	0.44	0.57
LAM (post-hoc)	0.76	0.86	0.89	0.60	0.55	0.74	0.58	0.74	0.78	0.45	0.44	0.59
LAM (online)	0.76	0.87	0.90	0.60	0.56	0.77	0.57	0.74	0.78	0.45	0.43	0.63

Table 4: **Results on MSLS.** LAM yields state-of-the-art uncertainties and matches the predictive performance of deterministic trained models. We evaluate on both the validation set and the official challenge set (Warburg et al., 2020).



Figure 9: **Images with lowest and highest variance** for PFE, post-hoc LAM, and online LAM across MSLS validation set. LAM reliably associates high uncertainty to images that are blurry, are captured facing the pavement, or contain vegetation. These images do not contain features that are descriptive of a specific place, making them especially challenging to geographically locate.



Figure 10: **Sparsification curve.** Online and posthoc LAM's sparsification curves monotonically increase, illustrating that they reliably associate higher uncertainty to harder observations.

lustrates that LAM produces reliable uncertainties for this challenging open-set retrieval task. Fig. 9 shows the queries associated with the highest and lowest uncertainty. LAM predicts high uncertainty to images with are blurry, captured facing into the pavement, or contain mostly vegetation. These images do not have features that are descriptive of a specific place, making them hard to geographically locate.

Limitations. Similar to other Bayesian methods, LAM

relies on n samples to obtain uncertainties. This makes inferences n times slower. Computing the Hessian at every step during online LAM also makes training time slower. To combat long training times and high memory usage, we use a last-layer LAM and thus only estimate and sample for a weight posterior of the last layer. The last-layer LAM training time is 3 hours for online LAM vs 2.3 hours for deterministic contrastive loss on LFW, and 30 minutes vs 15 minutes loss on CUB200 on an NVIDIA RTX A5000.

5. Conclusion

In this paper, we have introduced a Bayesian encoder for metric learning, the Laplacian Metric Learner (LAM), which uses the Laplace approximation. We prove that the contrastive loss is indeed a valid unnormalized log-posterior, and develop three Hessian approximations, which ensures a positive definite covariance matrix. We propose a novel decomposition of the Generalized Gauss-Newton approximation that improves Hessian approximations of ℓ_2 -normalized networks. Empirically, we demonstrate that LAM consistently produces well-calibrated uncertainties, reliably detects out-of-distribution examples, and achieves state-of-theart predictive performance on both closed-set and challenging open-set image retrieval tasks. Acknowledgement. This work was supported by research grants (15334, 42062) from VILLUM FONDEN. This project received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement 757360). The work was partly funded by the Novo Nordisk Foundation through the Center for Basic Machine Learning Research in Life Science (NNF20OC0062606). The authors acknowledge the Pioneer Centre for AI, DNRF grant P1.

References

- Arandjelović, R., Gronat, P., Torii, A., Pajdla, T., and Sivic, J. Netvlad: CNN architecture for weakly supervised place recognition. In *Proceedings of the IEEE Conference* on Computer Vision and Pattern Recognition, pp. 5297– 5307, 2016.
- Blundell, C., Cornebise, J., Kavukcuoglu, K., and Wierstra, D. Weight uncertainty in neural networks, 2015.
- Chang, J., Lan, Z., Cheng, C., and Wei, Y. Data uncertainty learning in face recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 5710–5719, 2020.
- Christensen, P. E., Warburg, F., Jia, M., and Belongie, S. Searching for structure in unfalsifiable claims, 2022.
- Chun, S., Oh, S. J., De Rezende, R. S., Kalantidis, Y., and Larlus, D. Probabilistic embeddings for cross-modal retrieval. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 8415–8424, 2021.
- Dangel, F., Kunstner, F., and Hennig, P. BackPACK: Packing more into backprop. In *International Conference on Learning Representations (ICLR)*, 2020.
- Davison, A. J., Reid, I. D., Molton, N. D., and Stasse, O. Monoslam: Real-time single camera slam. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 29 (6):1052–1067, 2007. doi: 10.1109/TPAMI.2007.1049.
- Daxberger, E., Kristiadi, A., Immer, A., Eschenhagen, R., Bauer, M., and Hennig, P. Laplace redux–effortless Bayesian deep learning. In *NeurIPS*, 2021.
- Denker, J. and LeCun, Y. Transforming neural-net output levels to probability distributions. *Advances in Neural Information Processing Systems*, 1990.
- Detlefsen, N. S., Jørgensen, M., and Hauberg, S. Reliable training and estimation of variance networks. In *Advances in Neural Information Processing Systems*, 2019.
- Detlefsen, N. S., Pouplin, A., Feldager, C. W., Geng, C., Kalatzis, D., Hauschultz, H., Duque,

M. G., Warburg, F., Miani, M., and Hauberg, S. StochMan, 2021. URL https://github.com/MachineLearningLifeScience/stochman/.

- Foresee, F. D. and Hagan, M. T. Gauss-Newton approximation to Bayesian learning. In *Proceedings of International Conference on Neural Networks (ICNN'97)*, volume 3, pp. 1930–1935. IEEE, 1997.
- Gal, Y. and Ghahramani, Z. Dropout as a Bayesian approximation: Representing model uncertainty in deep learning. In *international conference on machine learning*, pp. 1050–1059. PMLR, 2016.
- Gershman, S. and Goodman, N. Amortized inference in probabilistic reasoning. In *Proceedings of the annual meeting of the cognitive science society*, volume 36, 2014.
- Hadsell, R., Chopra, S., and LeCun, Y. 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.
- He, K., Zhang, X., Ren, S., and Sun, J. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 770–778, 2016.
- Huang, G. B., Ramesh, M., Berg, T., and Learned-Miller, E. Labeled faces in the wild: A database for studying face recognition in unconstrained environments. Technical Report 07-49, University of Massachusetts, Amherst, October 2007.
- Immer, A., Korzepa, M., and Bauer, M. Improving predictions of Bayesian neural nets via local linearization. In *International Conference on Artificial Intelligence and Statistics*, pp. 703–711. PMLR, 2021.
- Kenney, J. F. and Keeping, E. Mathematics of statistics, vol. ii, 1951.
- Kingma, D. P. and Welling, M. Auto-encoding variational Bayes. *arXiv preprint arXiv:1312.6114*, 2013.
- Krause, J., Stark, M., Deng, J., and Fei-Fei, L. 3d object representations for fine-grained categorization. In 2013 IEEE International Conference on Computer Vision Workshops, pp. 554–561, 2013.
- Krizhevsky, A. Learning multiple layers of features from tiny images, 2009.
- Lakshminarayanan, B., Pritzel, A., and Blundell, C. Simple and scalable predictive uncertainty estimation using deep ensembles. *Advances in Neural Information Processing Systems*, 30, 2017.

- Laplace, P. S. Mémoire sur la probabilité des causes par les événements. *Mémoire de l'Académie Royale des Sciences*, 1774.
- LeCun, Y., Denker, J., and Solla, S. Optimal brain damage. Advances in Neural Information Processing Systems, 1989.
- LeCun, Y., Bottou, L., Bengio, Y., and Haffner, P. Gradientbased learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998.
- MacKay, D. J. C. A practical Bayesian framework for backpropagation networks. *Neural Computation*, 4(3): 448–472, 1992.
- Maddox, W. J., Izmailov, P., Garipov, T., Vetrov, D. P., and Wilson, A. G. A simple baseline for Bayesian uncertainty in deep learning. *Advances in Neural Information Processing Systems*, 32, 2019.
- Miani, M., Warburg, F., Moreno-Muñoz, P., Detlefsen, N. S., and Hauberg, S. Laplacian autoencoders for learning stochastic representations. In Advances in Neural Information Processing Systems, 2022.
- Musgrave, K., Belongie, S., and Lim, S.-N. A metric learning reality check. In *European Conference on Computer Vision*, pp. 681–699. Springer, 2020.
- Nalisnick, E., Matsukawa, A., Teh, Y. W., Gorur, D., and Lakshminarayanan, B. Do deep generative models know what they don't know? *arXiv preprint arXiv:1810.09136*, 2018.
- Netzer, Y., Wang, T., Coates, A., Bissacco, A., Wu, B., and Ng, A. Y. Reading digits in natural images with unsupervised feature learning, 2011.
- Oh, S. J., Murphy, K., Pan, J., Roth, J., Schroff, F., and Gallagher, A. Modeling uncertainty with hedged instance embedding. arXiv preprint arXiv:1810.00319, 2018.
- Paszke, A., Gross, S., Chintala, S., Chanan, G., Yang, E., DeVito, Z., Lin, Z., Desmaison, A., Antiga, L., and Lerer, A. Automatic differentiation in PyTorch. In *NIPS-W*, 2017.
- Radenović, F., Tolias, G., and Chum, O. Fine-tuning CNN image retrieval with no human annotation. *IEEE transactions on pattern analysis and machine intelligence*, 41(7): 1655–1668, 2018.
- Rezende, D. J., Mohamed, S., and Wierstra, D. Stochastic backpropagation and approximate inference in deep generative models. In *International conference on machine learning*, pp. 1278–1286. PMLR, 2014.

- Schroff, F., Kalenichenko, D., and Philbin, J. Facenet: A unified embedding for face recognition and clustering. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 815–823, 2015.
- Shi, Y. and Jain, A. K. Probabilistic face embeddings. In Proceedings of the IEEE/CVF International Conference on Computer Vision, pp. 6902–6911, 2019.
- Song, Y. and Soleymani, M. Polysemous visual-semantic embedding for cross-modal retrieval. In *Proceedings* of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 1979–1988, 2019.
- Sra, S. A short note on parameter approximation for von Mises-Fisher distributions: and a fast implementation of i s (x). *Computational Statistics*, 27(1):177–190, 2012.
- Stylianou, A., Xuan, H., Shende, M., Brandt, J., Souvenir, R., and Pless, R. Hotels-50k: A global hotel recognition dataset. *CoRR*, 2019.
- Sun, J. J., Zhao, J., Chen, L.-C., Schroff, F., Adam, H., and Liu, T. View-invariant probabilistic embedding for human pose. In *European Conference on Computer Vision*, pp. 53–70. Springer, 2020.
- Taha, A., Chen, Y.-T., Misu, T., Shrivastava, A., and Davis, L. Unsupervised data uncertainty learning in visual retrieval systems. arXiv preprint arXiv:1902.02586, 2019a.
- Taha, A., Chen, Y.-T., Yang, X., Misu, T., and Davis, L. Exploring uncertainty in conditional multi-modal retrieval systems. *arXiv preprint arXiv:1901.07702*, 2019b.
- The Week. Is facial recognition technology racist?, 2018. URL https://www.theweek.co.uk/95383/is-facial-recognition-racist.
- Vilnis, L. and McCallum, A. Word representations via Gaussian embedding. arXiv preprint arXiv:1412.6623, 2014.
- Wah, C., Branson, S., Welinder, P., Perona, P., and Belongie, S. Caltech-UCSD birds 200, 2011.
- Wang, M. and Deng, W. Deep face recognition: A survey. *Neurocomputing*, 429:215–244, March 2021.
- Warburg, F., Hauberg, S., Lopez-Antequera, M., Gargallo, P., Kuang, Y., and Civera, J. Mapillary street-level sequences: A dataset for lifelong place recognition. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), June 2020.
- Warburg, F., Jørgensen, M., Civera, J., and Hauberg, S. Bayesian triplet loss: Uncertainty quantification in image retrieval. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 12158–12168, 2021.

- Wilber, M., Kwak, I. S., Kriegman, D., and Belongie, S. Learning concept embeddings with combined humanmachine expertise. In *Proceedings of the IEEE International Conference on Computer Vision*, pp. 981–989, 2015.
- Xiao, H., Rasul, K., and Vollgraf, R. Fashion-MNIST: a novel image dataset for benchmarking machine learning algorithms, 2017.
- Xu, K., Zhang, M., Li, J., Du, S. S., ichi Kawarabayashi, K., and Jegelka, S. How neural networks extrapolate: From feedforward to graph neural networks, 2021.

A. Notation

Let \mathcal{X} be the data space, \mathcal{Z} be the latent space, Θ be the parameter space. For a given neural network architecture f and a parameter vector $\theta \in \Theta$, f_{θ} is a function

$$f_{\theta}: \mathcal{X} \to \mathcal{Z} \tag{11}$$

The discrete set $C = \{c_i\}_{i=1,...,C}$ of classes induces a partition of the points in the data space. A dataset is a collection of *data point*-class pairs $(x, c) \in \mathcal{X} \times C$.

$$\mathcal{D} = \{(x_i, c_i)\}_{i=1,\dots,D}$$
(12)

and it is itself partitioned by the classes into C sets

$$\mathcal{D}_c = \{ (x', c') \in \mathcal{D} | c' = c \} \subset \mathcal{D} \qquad \forall c \in \mathcal{C}.$$
(13)

We assume these sets always to contain at least one element $\mathcal{D}_c \neq \emptyset$ and we use the notation $|\mathcal{D}_c|$ to refer to their cardinalities. In the metric learning setting, instead of enforcing properties of a single data point, the goal is to enforce relations between data points. Thus we will often make use of pairs $\mathfrak{p}_{ij} = ((x_i, c_i), (x_j, c_j)) \in \mathcal{D}^2$, specifically we use the terms

$$\mathfrak{p}_{ij} \text{ is positive pair if } c_i = c_j \tag{14}$$

$$\mathfrak{p}_{ij} \text{ is negative pair if } c_i \neq c_j. \tag{15}$$

To ease the later notation we will consider the trivial pairs composed by a data point with itself. We then define the positive set and negative set, and compute their cardinalities, respectively

$$\mathcal{D}_{\text{pos}}^2 := \{ \mathfrak{p} \in \mathcal{D}^2 \text{ such that } p \text{ is positive} \} \qquad |\mathcal{D}_{\text{pos}}^2| = \sum_{c \in \mathcal{C}} |\mathcal{D}_c|^2 \tag{16}$$

$$\mathcal{D}_{\text{neg}}^2 := \{ \mathfrak{p} \in \mathcal{D}^2 \text{ such that } p \text{ is negative} \} \qquad |\mathcal{D}_{\text{neg}}^2| = \sum_{\substack{c_i, c_j \in \mathcal{C}^2 \\ \text{s.t. } c_i \neq c_j}} |\mathcal{D}_{c_i}| |\mathcal{D}_{c_j}| \qquad (17)$$

A common trick in metric learning is to introduce a margin $m \in \mathbb{R}^+$. The margin induces a further partition of the \mathcal{D}_{neg}^2 set into pairs with close or far embeddings. For any pair $\mathfrak{p}_{ij} = ((x_i, c_i), (x_j, c_j))$ we describe the pair as being

inside the margin if
$$||f_{\theta}(x_i) - f_{\theta}(x_j)|| \le m$$
 (18)

outside the margin if
$$||f_{\theta}(x_i) - f_{\theta}(x_j)|| > m$$
 (19)

irrespective of the classes c_i and c_j . We can make use of this definition to consider only the pairs that have a non-zero loss. From this logic, it is convenient to define the set $\mathcal{D}_{\text{neg inside}}^2 := \{\mathfrak{p} \in \mathcal{D}_{\text{neg}}^2 \text{ inside the margin}\}.$

The *target*, or label, is the value that encodes the information we want to learn. In classical settings, we have one scalar for each data point: a class for classification, a value for regression. In our setting, we consider a target $y_{ij} \in \mathbb{R}$ for every pair $\mathfrak{p}_{ij} \in \mathcal{D}^2$. With this notation in mind, we can continue to define the contrastive loss.

B. Contrastive Loss

Recall the definition of the contrastive loss (Hadsell et al., 2006)

$$\mathcal{L}_{\rm con}(\theta) = \frac{1}{2} \|f_{\theta}(x_a) - f_{\theta}(x_p)\|^2 + \frac{1}{2} \max\left(0, m - \|f_{\theta}(x_a) - f_{\theta}(x_n)\|^2\right),\tag{20}$$

where f is a network with parameters θ and x_p is a data point the same class as the anchor x_a and different from the negative data point x_n . The loss over the whole dataset is then informally defined as $\mathcal{L}_{con}(\theta; \mathcal{D}) = \mathbb{E}_{\mathcal{D}}[\mathcal{L}_{con}(\theta)]$, where the expectation is taken over tuples (x_a, x_p, x_n) satisfying the positive and negative constrains. This definition is intuitive and compact, but not formal enough to show that the contrastive loss is in fact an unnormalized log posterior. The main issue is the explicit notation for positive x_p and negative x_n , which at first glance seems innocent, but for later derivations becomes cumbersome. Instead, we express the loss in an equivalent but more verbose way.

In order to understand that we are doing nothing more than a change in notation, it is first convenient to express Eq. 20 as

$$\mathcal{L}_{\rm con}(\theta) = \begin{cases} \frac{1}{2} \|f_{\theta}(x_i) - f_{\theta}(x_j)\|^2 & \text{if the pair } \mathfrak{p}_{ij} \text{ is positive} \\ 0 & \text{if the pair } \mathfrak{p}_{ij} \text{ is negative and outside the margin} \\ m - \frac{1}{2} \|f_{\theta}(x_i) - f_{\theta}(x_j)\|^2 & \text{if the pair } \mathfrak{p}_{ij} \text{ is negative and inside the margin} \end{cases}$$
(21)

which reveals the three cases of the *per-observation* contrastive loss. Notice that, up to a neglectable additive constant, all cases shares the same form of being scalar multiples of a distance between the two embeddings $z_i = f_{\theta}(x_i)$ and $z_j = f_{\theta}(x_j)$. This allows us to combine the three scenarios into a single-case *per-observation* contrastive loss, parametrized by a scalar $y \in \mathbb{R}$, as

$$\mathcal{L}_y: \quad \mathcal{Z}^2 \longrightarrow \mathbb{R}$$
$$z_i, z_j \longmapsto \frac{1}{2} y \| z_i - z_j \|^2$$
(22)

and we later entrust the distinction between the three scenarios, as in Eq. 21, to the scalar y. Although it is important to define this loss for a general target $y \in \mathbb{R}$ (we will interpret it as a Von-Mises-Fisher concentration parameter κ in Appendix D), practically we will make use of the specific instances $y = y_{ij} \in \mathbb{R}$ with the target values defined for every data indexes i and j as

$$y_{ij} := \begin{cases} \frac{1}{|\mathcal{D}_{pos}^2|} & \text{if the pair } \mathfrak{p}_{ij} \text{ is positive} \\ 0 & \text{if the pair } \mathfrak{p}_{ij} \text{ is negative and outside the margin} \\ -\frac{1}{|\mathcal{D}_{neel}^2|} & \text{if the pair } \mathfrak{p}_{ij} \text{ is negative and inside the margin} \end{cases}$$
(23)

We define the contrastive loss over the entire dataset $\mathcal{L}(;\mathcal{D}): \Theta \to \mathbb{R}$ as a sum over all pairs of the *per-observation* contrastive loss (22) with targets (23)

$$\mathcal{L}(\theta; \mathcal{D}) := \sum_{\mathfrak{p}_{ij} \in \mathcal{D}^2} \mathcal{L}_{y_{ij}}(\underbrace{f_{\theta}(x_i)}_{z_i}, \underbrace{f_{\theta}(x_j)}_{z_j}).$$
(24)

Notice that we slightly overload the notation $\mathcal{L}(\theta)$ and $\mathcal{L}_y(\theta; \mathcal{D})$ for the dataset loss and the per-observation-loss, respectively.

Expanded expression. In order to better highlight the equivalence, it may be useful to express the loss explicitly using all the previous definitions.

$$\mathcal{L}(\theta; \mathcal{D}) = \sum_{\mathfrak{p}_{ij} \in \mathcal{D}^2} \mathcal{L}_{y_{ij}}(f_{\theta}(x_i), f_{\theta}(x_j))$$
(25)

$$= \sum_{\mathfrak{p}_{ij} \in \mathcal{D}^2} \frac{1}{2} y_{ij} \| f_{\theta}(x_i) - f_{\theta}(x_j) \|^2$$
(26)

$$= \sum_{\mathfrak{p}_{ij} \in \mathcal{D}_{\text{pos}}^{2}} \frac{1}{2} \frac{1}{|\mathcal{D}_{\text{pos}}^{2}|} \|f_{\theta}(x_{i}) - f_{\theta}(x_{j})\|^{2} + \sum_{\mathfrak{p}_{ij} \in \mathcal{D}_{\text{neg inside}}^{2}} -\frac{1}{2} \frac{1}{|\mathcal{D}_{\text{neg}}^{2}|} \|f_{\theta}(x_{i}) - f_{\theta}(x_{j})\|^{2}$$
(27)

$$= \frac{1}{|\mathcal{D}_{\text{pos}}^2|} \sum_{\mathfrak{p}_{ij} \in \mathcal{D}_{\text{pos}}^2} \frac{1}{2} \|f_{\theta}(x_i) - f_{\theta}(x_j)\|^2 - \frac{1}{|\mathcal{D}_{\text{neg}}^2|} \sum_{\mathfrak{p}_{ij} \in \mathcal{D}_{\text{neg inside}}^2} \frac{1}{2} \|f_{\theta}(x_i) - f_{\theta}(x_j)\|^2$$
(28)

The scaling factor $\frac{1}{|\mathcal{D}_{pos}^2|}$ for positives and $\frac{1}{|\mathcal{D}_{neg}^2|}$ for negatives, together with the sum over pairs, leads to a per-type average. In this sense, we can informally say that $\mathcal{L}(\theta) = \mathbb{E}[\mathcal{L}_{con}(\theta)]$ under a distribution that assigns equal probabilities of a pair being positive or negative. We highlight that scaling is key in order to ensure positive definiteness in Proposition F.3.

B.1. Minibatch

In the previous Section, we defined the contrastive loss for the entire dataset (24). However, in practice it is approximated with minibatching. This approximation induces a bias that we can correct by simply adjusting the target y_{ij} as in Eq. 30.

Minibatching recap. When dealing with a huge number of identically distributed data, it is common to assume that a big enough arbitrary subset will follow the same distribution and thus have the same properties, specifically we assume the expected value to be similar $\mathbb{E}_{\mathcal{S}}[f(s)] \approx \mathbb{E}_{\mathcal{S}'}[f(s)]$. The idea of minibatching relies on this to approximate a sum over a set \mathcal{S} with a *scaled* sum over a subset $\mathcal{S}' \subseteq \mathcal{S}$,

$$\sum_{s \in \mathcal{S}} f(s) \approx \frac{|\mathcal{S}|}{|\mathcal{S}'|} \sum_{s \in \mathcal{S}'} f(s)$$
⁽²⁹⁾

where S = D and S' is the set of data points in the minibatch.

Minibatching contrastive. Conversely, in the constrastive setting, we need to consider a subset of the pairs (rather than single observations), i.e. $S = D^2$. In practice this subset will not be representative of the positive and negative ratio, thus we need a different scaling to account for that. This process can be viewed as taking two independent minibatches at the same time, one representative of the positives and one of the negatives. This intuition is formalized by the following Definition and Proposition.

Definition. Consider a minibatch set of pairs $\mathcal{B} \subseteq \mathcal{D}^2$ and its partition $\mathcal{B} = \mathcal{B}_{pos} \cup \mathcal{B}_{neg}$ in positives $\mathcal{B}_{pos} \subseteq \mathcal{D}_{pos}^2$ and negatives $\mathcal{B}_{neg} \subseteq \mathcal{D}_{neg}^2$, then we can define a new scaled target as

$$y_{ij}^{\mathcal{B}} := \begin{cases} \frac{|\mathcal{B}|}{|\mathcal{D}^2||\mathcal{B}_{\text{pos}}|} & \text{if the pair } \mathfrak{p}_{ij} \text{ is positive} \\ 0 & \text{if the pair } \mathfrak{p}_{ij} \text{ is negative and outside the margin} \\ -\frac{|\mathcal{B}|}{|\mathcal{D}^2||\mathcal{B}_{\text{neg}}|} & \text{if the pair } \mathfrak{p}_{ij} \text{ is negative and inside the margin} \end{cases}$$
(30)

and so we can properly approximate the loss by minibatching positives and negatives independently

Proposition B.1. Assume that the positives in the batch are representative of the positives in the whole dataset, and similarly for the negatives, i.e. $\mathbb{E}_{\mathcal{D}_{pos}^2}[\mathcal{L}_y] \approx \mathbb{E}_{\mathcal{B}_{pos}}[\mathcal{L}_y]$ and $\mathbb{E}_{\mathcal{D}_{neg}^2}[\mathcal{L}_y] \approx \mathbb{E}_{\mathcal{B}_{neg}}[\mathcal{L}_y]$. Then the loss, as defined in Eq. 24, can be approximated by using the target in Eq. 30 with

$$\mathcal{L}(\theta; \mathcal{D}) \approx \frac{|\mathcal{D}^2|}{|\mathcal{B}|} \sum_{\mathfrak{p}_{ij} \in \mathcal{B}} \mathcal{L}_{y_{ij}^{\mathcal{B}}}(f_{\theta}(x_i), f_{\theta}(x_j))$$
(31)

Proof. The equality is proven by applying the logic of Eq. 29 two times independently, once for the positive pairs with $\mathcal{B}_{pos} = \mathcal{S}' \subseteq \mathcal{S} = \mathcal{D}_{pos}^2$ and once for the negatives with $\mathcal{B}_{neg} = \mathcal{S}' \subseteq \mathcal{S} = \mathcal{D}_{neg}^2$ and then rearranging the terms

$$\mathcal{L}(\theta; \mathcal{D}) = \sum_{\mathfrak{p}_{ij} \in \mathcal{D}^2} \mathcal{L}_{y_{ij}}(f_{\theta}(x_i), f_{\theta}(x_j))$$
(32)

$$= \sum_{\mathfrak{p}_{ij} \in \mathcal{D}^2_{\text{pos}}} \mathcal{L}_{y_{ij}}(f_{\theta}(x_i), f_{\theta}(x_j)) + \sum_{\mathfrak{p}_{ij} \in \mathcal{D}^2_{\text{neg}}} \mathcal{L}_{y_{ij}}(f_{\theta}(x_i), f_{\theta}(x_j))$$
(33)

$$\approx \frac{|\mathcal{D}_{\text{pos}}^2|}{|\mathcal{B}_{\text{pos}}|} \sum_{\mathfrak{p}_{ij} \in \mathcal{B}_{\text{pos}}} \mathcal{L}_{y_{ij}}(f_{\theta}(x_i), f_{\theta}(x_j)) + \frac{|\mathcal{D}_{\text{neg}}^2|}{|\mathcal{B}_{\text{neg}}|} \sum_{\mathfrak{p}_{ij} \in \mathcal{B}_{\text{neg}}} \mathcal{L}_{y_{ij}}(f_{\theta}(x_i), f_{\theta}(x_j))$$
(34)

$$=\frac{|\mathcal{D}^2|}{|\mathcal{B}|}\left(\frac{|\mathcal{D}_{\text{pos}}^2|}{|\mathcal{D}^2|}\frac{|\mathcal{B}|}{|\mathcal{B}_{\text{pos}}|}\sum_{\mathfrak{p}_{ij}\in\mathcal{B}_{\text{pos}}}\mathcal{L}_{y_{ij}}(f_{\theta}(x_i),f_{\theta}(x_j))+\frac{|\mathcal{D}_{\text{neg}}^2|}{|\mathcal{D}^2|}\frac{|\mathcal{B}|}{|\mathcal{B}_{\text{neg}}|}\sum_{\mathfrak{p}_{ij}\in\mathcal{B}_{\text{neg}}}\mathcal{L}_{y_{ij}}(f_{\theta}(x_i),f_{\theta}(x_j))\right)$$
(35)

$$= \frac{|\mathcal{D}^2|}{|\mathcal{B}|} \left(\frac{\operatorname{pos}_{\mathcal{D}^2}^{\%}}{\operatorname{pos}_{\mathcal{B}}^{\%}} \sum_{\mathfrak{p}_{ij} \in \mathcal{B}_{\operatorname{pos}}} \mathcal{L}_{y_{ij}}(f_{\theta}(x_i), f_{\theta}(x_j)) + \frac{\operatorname{neg}_{\mathcal{D}^2}^{\%}}{\operatorname{neg}_{\mathcal{B}}^{\%}} \sum_{\mathfrak{p}_{ij} \in \mathcal{B}_{\operatorname{neg}}} \mathcal{L}_{y_{ij}}(f_{\theta}(x_i), f_{\theta}(x_j)) \right)$$
(36)

$$= \frac{|\mathcal{D}^2|}{|\mathcal{B}|} \left(\sum_{\mathfrak{p}_{ij} \in \mathcal{B}_{\text{pos}}} \mathcal{L}_{y_{ij}^{\mathcal{B}}}(f_{\theta}(x_i), f_{\theta}(x_j)) + \sum_{\mathfrak{p}_{ij} \in \mathcal{B}_{\text{neg}}} \mathcal{L}_{y_{ij}^{\mathcal{B}}}(f_{\theta}(x_i), f_{\theta}(x_j)) \right)$$
(37)

$$=\frac{|\mathcal{D}^2|}{|\mathcal{B}|}\sum_{\mathfrak{p}_{ij}\in\mathcal{B}}\mathcal{L}_{y_{ij}^{\mathcal{B}}}(f_{\theta}(x_i), f_{\theta}(x_j))$$
(38)

where $pos^{\%}$ and $neg^{\%}$ are the percentage of respectively positives and negatives in a given set, indicated in the subscript. \Box

We highlight that this scaling is linear, and thus is reflected in both first and second-order derivatives. This will later become important for scaling the hessian.

C. Normalization layer and Von Mises–Fisher

It is common in metric learning to add a normalization layer at the end of the neural network architecture. This, besides improving performances, has an interesting geometric interpretation. Moreover, happens to be fundamental in order to interpret the loss in a probabilistic way, and consequently to apply the Laplace approximation in a meaningful way.

Adding an l^2 -normalization layer is a practical way to enforce that all outputs lie on the unit sphere. In other words, we are assuming our latent manifold to be

$$\mathcal{Z} := S^Z \subset \mathbb{R}^{Z+1}. \tag{39}$$

We can generalize the concept of Gaussian to the unit sphere. Start from a normal distribution with isotropic covariance $\sigma^2 \mathbb{I}$ and mean μ . Conditioning on ||z|| = 1 leads to a distribution on S^Z , the so-called von Mises–Fisher distribution \mathcal{N}^S . More succinctly, the restriction of any isotropic multivariate normal density to the unit hypersphere, gives a von Mises-Fisher density, up to normalization.

$$\mathcal{N}^{S}(z|\mu,\kappa) \sim \mathcal{P}(z) = c_{\kappa} e^{-\kappa \frac{\|z-\mu\|^{2}}{2}}$$
(40)

where, importantly, the normalization constant $c_{\kappa} \in \mathbb{R}$ only depends on κ and not on μ . The von Mises-Fisher distribution is parametrized with a directional mean μ and a scalar concentration parameter κ , which can be interpreted as the inverse of an isotropic covariance $\kappa = \frac{1}{\sigma^2}$.

Being both μ and z of the unitary norm, the norm $||z - \mu||^2$ can be more efficiently computed through the scalar product $\langle z, \mu \rangle$, that is the reason it is also called Arccos distance. Moreover, we highlight that on the unit sphere, the equivalence

$$||z - \mu||^2 = 2 - \langle z, \mu \rangle = 4 - ||z + \mu||^2$$
(41)

holds. This will be key in proving the equivalence of the probabilistic setting from Eq. 54 to Eq. 55.

Parameters estimators. Having access to *N* samples drawn from a Gaussian distribution with unknown mean and variance, it is common to compute an empirical estimate of such parameters. Various estimators exists, each satisfying different properties like being unbiased (Kenney & Keeping, 1951).

In a similar fashion, for Von-Mises-Fisher distributions we make use of two empirical estimators Sra (2012) of the parameters μ and $\kappa = \frac{1}{\sigma^2}$. We compute the empirical mean direction

$$\bar{\mu} = \frac{\mu}{\bar{R}}, \qquad \bar{R} = \|\mu\|, \qquad \mu = \frac{1}{N} \sum_{i=0}^{N} z_i,$$
(42)

from these samples and the approximate concentration parameter

$$\bar{\kappa} = \frac{\bar{R}(D - \bar{R}^2)}{1 - \bar{R}^2}.$$
(43)

D. Probabilistic view

In the metric learning framework, a dataset \mathcal{D} plays two very different roles at the same time. This happens in the same spirit as, in electrostatics, a charged particle at the same time *generates* an electric field and *interact* with the existing electric field. This subtle difference, although being neglectable in the classic non-Bayesian setting, is conceptually important in the probabilistic setting. Thus, for now, we assume to be given a *generating* dataset \mathcal{D}_G and an *interacting* dataset \mathcal{D}_I . Only after the derivations, we will set them to be equal.

For some fixed value $\kappa > 0$, each data point $x \in \mathcal{X}$ induces two Von Mises-Fisher distributions on the latent space

$$\mathcal{P}^{\to\leftarrow}(z|x,\theta) \sim \mathcal{N}^S(z|f_\theta(x),\kappa) \tag{44}$$

$$\mathcal{P}^{\leftarrow \rightarrow}(z|x,\theta) \sim \mathcal{N}^S(z| - f_\theta(x), \kappa) \tag{45}$$

one centered in the embedding point $f_{\theta}(x)$ and one centered in the antipodal point $-f_{\theta}(x)$. It is convenient to explicit the densities of these two distributions, according to Eq. 40

$$\mathcal{P}^{\rightarrow\leftarrow}(z|x,\theta) = c_{\kappa} \exp\left(-\kappa \frac{\|z - f_{\theta}(x)\|^2}{2}\right)$$
(46)

$$\mathcal{P}^{\leftarrow \rightarrow}(z|x,\theta) = c_{\kappa} \exp\left(-\kappa \frac{\|z + f_{\theta}(x)\|^2}{2}\right).$$
(47)



Figure 11: Representation of the Von Mises-Fisher densities \mathcal{P}^{\leftarrow} and \mathcal{P}^{\leftarrow} in the one-dimensional unit circle $\mathcal{Z} = \mathcal{S}^1$. Intuitively, the distribution corresponding to an attractive force is supported near the embedding $z = f_{\theta}(x)$, while the distribution corresponding to a repulsive force is supported far from the embedding, i.e. close to the antipodal $z = -f_{\theta}(x)$. The higher the value of κ is, the higher the precision, the lower the variance and the narrower the support.

The generating dataset \mathcal{D}_G induces a distribution $\mathcal{P}^c_{\mathcal{Z}} \in \Delta(\mathcal{Z})$ on the latent space for each class $c \in \mathcal{C}$ defined by

$$\mathcal{P}_{\mathcal{Z}}^{c}(z|\mathcal{D}_{G},\theta) = \frac{c_{\kappa}}{|\mathcal{D}_{G}|} \prod_{\substack{(x_{j},c_{j})\in\mathcal{D}_{G}\\\text{s.t. }c_{j}=c}} \mathcal{P}^{\rightarrow\leftarrow}(z|x_{j},\theta) \prod_{\substack{(x_{j},c_{j})\in\mathcal{D}_{G}\\\text{s.t. }c_{j}\neq c}} \mathcal{P}^{\leftarrow\rightarrow}(z|x_{j},\theta) \tag{48}$$

that is a product of one *attractive* term \mathcal{P}^{\leftarrow} for all anchor points of the same class, times a *repulsive* term \mathcal{P}^{\leftarrow} for all anchor points of different classes. In the same spirit as the product of Gaussian densities is a Gaussian density, we can show that $\mathcal{P}_{\mathcal{Z}}^c$ is a Von Mises-Fisher distribution itself, the precision of which is the sum of the precisions. This implies that the normalization constant only depends on κ for this distribution as well.

Having access to a likelihood in the latent space, the interacting dataset likelihood is then classically defined as the product

$$\mathcal{P}(\mathcal{D}_I | \mathcal{D}_G, \theta) = \prod_{(x_i, c_i) \in \mathcal{D}_I} \mathcal{P}_{\mathcal{Z}}^{c_i}(z | \mathcal{D}_G, \theta) \big|_{z = f_\theta(x_i)}$$
(49)

Using the definition in Eq. 48 and rearranging the terms, we can rewrite this likelihood in extended form as

$$\mathcal{P}(\mathcal{D}_{I}|\mathcal{D}_{G},\theta) = \bar{c}_{\kappa} \prod_{\substack{\mathfrak{p}_{ij} \in \mathcal{D}_{I} \times \mathcal{D}_{G} \\ \text{s.t. } \mathfrak{p}_{ij} \text{ is positive}}} \mathcal{P}^{\rightarrow \leftarrow}(z|x_{j},\theta) \Big|_{z=f_{\theta}(x_{i})} \prod_{\substack{\mathfrak{p}_{ij} \in \mathcal{D}_{I} \times \mathcal{D}_{G} \\ \text{s.t. } \mathfrak{p}_{ij} \text{ is negative}}} \mathcal{P}^{\leftarrow \rightarrow}(z|x_{j},\theta) \Big|_{z=f_{\theta}(x_{i})}$$
(50)

Setting the generating and interacting dataset to be the same, we can define the constrastive learning dataset likelihood as

$$\mathcal{P}(\mathcal{D}|\theta) := \mathcal{P}(\mathcal{D}_I|\mathcal{D}_G, \theta) \big|_{\mathcal{D}_I = \mathcal{D}_G = \mathcal{D}}$$
(51)

D.1. Equivalence of the two settings

The contrastive term for a single pair is equivalent to the negative log-likelihood of a von Mises-Fisher distribution, up to additive constants. Specifically, a positive target $y = \kappa > 0$ is related to the attractive term $\mathcal{P}^{\rightarrow \leftarrow}$

$$\mathcal{L}_{\kappa}(f_{\theta}(x_e), f_{\theta}(x_a)) = \frac{1}{2}\kappa \|f_{\theta}(x_e) - f_{\theta}(x_a)\|^2$$
(52)

$$= \log(c_{\kappa}) - \log \mathcal{P}^{\to \leftarrow}(f_{\theta}(x_e)|x_a, \theta)$$
(53)

while a negative target $y = -\kappa < 0$ is related to the repulsive term $\mathcal{P}^{\leftarrow \rightarrow}$

$$\mathcal{L}_{-\kappa}(f_{\theta}(x_e), f_{\theta}(x_a)) = -\frac{1}{2}\kappa \|f_{\theta}(x_e) - f_{\theta}(x_a)\|^2$$
(54)

$$= -2\kappa + \frac{1}{2}\kappa \|f_{\theta}(x_e) + f_{\theta}(x_a)\|^2$$
(55)

$$= -2\kappa + \log(c_{\kappa}) - \log \mathcal{P}^{\leftarrow \rightarrow}(f_{\theta}(x_e)|x_a, \theta).$$
(56)

This equivalence is reflected in the equivalence between the contrastive loss and the dataset negative log-likelihood, up to an additive constant. Expanding the definition of dataset likelihood we have

$$\mathcal{P}(\mathcal{D}|\theta) = \prod_{(x_i,c_i)\in\mathcal{D}} \mathcal{P}_{\mathcal{Z}}^{c_i}(f_{\theta}(x_i)|\mathcal{D},\theta)$$
(57)

$$= \bar{c}_{\kappa} \prod_{\substack{\mathfrak{p}_{ij} \in \mathcal{D}^{2} \\ \text{s.t. } \mathfrak{p}_{ij} \text{ is positive}}} \mathcal{P}^{\rightarrow \leftarrow}(z|x_{j},\theta) \big|_{z=f_{\theta}(x_{i})} \prod_{\substack{\mathfrak{p}_{ij} \in \mathcal{D}^{2} \\ \text{s.t. } \mathfrak{p}_{ij} \text{ is negative}}} \mathcal{P}^{\leftarrow \rightarrow}(z|x_{j},\theta) \big|_{z=f_{\theta}(x_{i})}$$
(58)

$$= \bar{c}_{\kappa} c_{\kappa}^{|\mathcal{D}|} \prod_{\substack{\mathfrak{p}_{ij} \in \mathcal{D}^{2} \\ \text{s.t. } \mathfrak{p}_{ij} \text{ is positive}}} \exp\left(-\kappa \frac{\|f_{\theta}(x_{i}) - f_{\theta}(x_{j})\|^{2}}{2}\right) \prod_{\substack{\mathfrak{p}_{ij} \in \mathcal{D}^{2} \\ \text{s.t. } \mathfrak{p}_{ij} \text{ is negative}}} \exp\left(-\kappa \frac{\|f_{\theta}(x_{i}) + f_{\theta}(x_{j})\|^{2}}{2}\right)$$
(59)

Considering the log-likelihood

$$\log \mathcal{P}(\mathcal{D}|\theta) = \sum_{(x_i, c_i) \in \mathcal{D}} \log \mathcal{P}_{\mathcal{Z}}^{c_i}(f_{\theta}(x_i)|\mathcal{D}, \theta)$$
(60)

$$= \log(\bar{c}_{\kappa}) + \sum_{\substack{\mathfrak{p}_{ij} \in \mathcal{D}^2\\ \text{s.t. } \mathfrak{p}_{ij} \text{ is positive}}} \log \mathcal{P}^{\to\leftarrow}(z|x_j,\theta) \big|_{z=f_{\theta}(x_i)} + \sum_{\substack{\mathfrak{p}_{ij} \in \mathcal{D}^2\\ \text{s.t. } \mathfrak{p}_{ij} \text{ is negative}}} \log \mathcal{P}^{\leftarrow\to}(z|x_j,\theta) \big|_{z=f_{\theta}(x_i)}$$
(61)

$$= \operatorname{const}(\kappa) - \sum_{\mathbf{n}_{i} \in \mathcal{D}^{2}} \mathcal{L}_{\kappa}(f_{\theta}(x_{i}), f_{\theta}(x_{j})) - \sum_{\mathbf{n}_{i} \in \mathcal{D}^{2}} \mathcal{L}_{-\kappa}(f_{\theta}(x_{i}), f_{\theta}(x_{j}))$$
(62)

s.t.
$$\mathfrak{p}_{ij} \subset \mathcal{P}$$

s.t. \mathfrak{p}_{ij} is positive s.t. \mathfrak{p}_{ij} is negative (62)

$$= \operatorname{const}(\kappa) - \mathcal{L}(\theta; \mathcal{D}) \tag{63}$$

thus the loss is equal, up to an additive constant, to the negative log-likelihood and we can proceed with the Bayesian interpretation. We highlight that $const(\kappa)$ is not dependent on θ , and so is neglected by the derivatives

$$\nabla_{\theta} \log \mathcal{P}(\mathcal{D}|\theta) = -\nabla_{\theta} \mathcal{L}(\theta; \mathcal{D})$$
(64)

$$\nabla_{\theta}^{2} \log \mathcal{P}(\mathcal{D}|\theta) = -\nabla_{\theta}^{2} \mathcal{L}(\theta; \mathcal{D}).$$
(65)

This shows that contrastive loss minimization is indeed a likelihood maximization. We highlight that this proof assumes margin m > 2, i.e. all negatives are inside the margin (because on the manifold S^Z the maximum distance is 2). Following the exact same structure, a proof can be done with arbitrary m. The only difference would be in the definition Eq. 47, instead of a Von Mises-Fisher it would be a discontinuous density and the normalization constants would depend on both κ and m.

E. Bayesian Metric learning

Having defined the dataset likelihood $\mathcal{P}(\mathcal{D}|\theta)$ conditioned to the parameter θ , the parameter likelihood $\mathcal{P}(\theta|\mathcal{D})$ conditioned to the dataset is defined according to Bayes rule as

$$\mathcal{P}(\theta|\mathcal{D}) = \frac{\mathcal{P}(\mathcal{D}|\theta)\mathcal{P}(\theta)}{\mathcal{P}(\mathcal{D})}.$$
(66)

The aim of metric learning is to maximise such likelihood with respect to the parameter θ , specifically, it aims at finding the Maximum A Posteriori

$$\theta^{\text{MAP}} \in \operatorname*{arg\,max}_{\theta \in \Theta} \mathcal{P}(\theta|\mathcal{D}). \tag{67}$$

We highlight that the argmax is the same in log-scale

$$\arg\max_{\theta\in\Theta} \mathcal{P}(\theta|\mathcal{D}) = \arg\max_{\theta\in\Theta} \log \mathcal{P}(\theta|\mathcal{D}) = \arg\max_{\theta\in\Theta} \log \mathcal{P}(\mathcal{D}|\theta) + \log \mathcal{P}(\theta)$$
(68)

where the log-prior takes the form of a standard l^2 regularize (weight decay), and, as shown before, maximizing the log-likelihood log $p(\mathcal{D}|\theta)$ is equivalent to minimizing of the contrastive loss $\mathcal{L}(\theta; \mathcal{D})$.

However in the Bayesian setting, we do not seek a maximum, but we seek an expression of the whole distribution $\mathcal{P}(\theta|\mathcal{D})$, for which we use the notation $q(\theta)$. We do that by maximizing the dataset likelihood $\mathcal{P}(\mathcal{D})$ by integrating out θ with $\mathcal{P}(\mathcal{D}) = \mathbb{E}_{\theta \sim q}[\mathcal{P}(\mathcal{D}|\theta)]$. This maximization is then defined by

$$q^{*}(\theta) := \mathcal{P}^{*}(\theta|\mathcal{D}) \in \underset{\mathcal{P}(\theta|\mathcal{D}) \in \mathcal{G}(\Theta)}{\operatorname{arg\,max}} \mathbf{E}_{\theta \sim \mathcal{P}(\theta|\mathcal{D})}[\mathcal{P}(\mathcal{D}|\theta)]$$
(69)

Notice that, limited by the ability to parametrize distributions, we aim at finding such a maximum on q over only some subspace of distributions on Θ . To this end, we consider the Laplace approximation, which is one way of choosing a subspace of the parameter distribution. In Laplace post-hoc we restrict ourselves to the space of Gaussians $\mathcal{G}(\Theta) \subset \Delta(\Theta)$ centered in θ^{MAP} . This is a strong assumption since there are (1) no guarantees of $p(\theta|\mathcal{D})$ being Gaussian and (2) no guarantees of the distribution to be centered in θ^{MAP} . Online Laplace lifts the latter assumption, and only assumes the parameters to be Gaussian distributed.

With the general Bayesian framework in mind, we now proceed to derive post-hoc and online Laplace.

E.1. Laplace post-hoc

The Bayes rule

$$\mathcal{P}(\theta|\mathcal{D}) = \frac{\mathcal{P}(\mathcal{D}|\theta)\mathcal{P}(\theta)}{\mathcal{P}(\mathcal{D})}$$
(70)

implies that

$$\nabla_{\theta}^{2} \log \mathcal{P}(\theta|\mathcal{D}) = \nabla_{\theta}^{2} \log \mathcal{P}(\mathcal{D}|\theta) + \nabla_{\theta}^{2} \log \mathcal{P}(\theta).$$
(71)

Assuming an isotropic Gaussian prior $\mathcal{P}(\theta) \sim \mathcal{N}(\theta|0, \sigma_{\text{prior}}^2 \mathbb{I})$ implies $\nabla_{\theta}^2 \log \mathcal{P}(\theta) = -\sigma_{\text{prior}}^{-2} \mathbb{I}$ and we have

$$\nabla_{\theta}^{2} \log \mathcal{P}(\theta|\mathcal{D}) = \nabla_{\theta}^{2} \log \mathcal{P}(\mathcal{D}|\theta) - \sigma_{\text{prior}}^{-2} \mathbb{I}$$
(72)

$$= -\nabla_{\theta}^{2} \mathcal{L}(\theta; \mathcal{D}) - \sigma_{\text{prior}}^{-2} \mathbb{I}.$$
(73)

Thus, we have two options:

• IF $\mathcal{P}(\theta|\mathcal{D})$ is a Gaussian it holds $\forall \theta^* \in \Theta$

$$\mathcal{P}(\theta|\mathcal{D}) \sim \mathcal{N}(\theta|\mu = \theta^{\text{MAP}}, \Sigma = (\nabla_{\theta}^{2} \mathcal{L}(\theta^{*}; \mathcal{D}) + \sigma_{\text{prior}}^{-2} \mathbb{I})^{-1})$$
(74)

• ELSE we can do a second order Taylor approximation of $\log \mathcal{P}(\theta|\mathcal{D})$ around θ^{MAP} and we have the approximation

$$\mathcal{P}(\theta|\mathcal{D}) \sim \mathcal{N}(\theta|\mu = \theta^{\text{MAP}}, \Sigma = (\nabla_{\theta}^{2} \mathcal{L}(\theta^{\text{MAP}}; \mathcal{D}) + \sigma_{\text{prior}}^{-2} \mathbb{I})^{-1})$$
(75)

E.2. Laplace online

At every step t we have some Gaussian on the parameter space

$$q^{t}(\theta) \sim \mathcal{N}(\theta|\mu = \theta_{t}, \Sigma = (H_{t})^{-1})$$
(76)

Where the values θ_t and H_t are iteratively defined as

$$\theta_{t+1} = \theta_t + \lambda \nabla_\theta \mathcal{L}(\theta_t; \mathcal{D}) \qquad \theta_0 = \mu_{\text{prior}}$$
(77)

and

$$H_{t+1} = (1-\alpha)H_t + \nabla^2_{\theta}\mathcal{L}(\theta_t; \mathcal{D}) \qquad H_0 = \sigma_{\text{prior}}^{-2}\mathbb{I}$$
(78)

For some learning rate λ and memory factor α . This means that $q^0(\theta)$ is actually the prior distribution, which is updated with the first and second order derivatives of the loss. The updates can be improved by computing the derivatives not only in the single point θ_t , but rather on the expected value with θ following the distribution q^t . This leads to the update rules

$$\theta_{t+1} = \theta_t + \lambda \mathbb{E}_{\theta \sim q^t(\theta)} [\nabla_\theta \mathcal{L}(\theta; \mathcal{D})] \qquad \theta_0 = \mu_{\text{prior}}$$
(79)

and

$$H_{t+1} = (1-\alpha)H_t + \mathbb{E}_{\theta \sim q^t(\theta)}[\nabla^2_{\theta}\mathcal{L}(\theta;\mathcal{D})] \qquad H_0 = \sigma_{\text{prior}}^{-2}\mathbb{I}$$
(80)

F. Derivatives

In order to perform the Laplace based learning, we need to compute the second-order derivative of the loss with respect to the parameter θ . Let us start by fixing a target $y \in \mathbb{R}$ and two data points $x_1, x_2 \in \mathcal{X}$ and compute the second order derivative of one contrastive term $\mathcal{L}_y(f_\theta(x_1), f_\theta(x_2))$. This term can be viewed as a composition of functions, graphically represented as

To make the derivation cleaner is it useful to define, for a given function $f_{\theta} : \mathcal{X} \to \mathcal{Z}$, and auxiliary function $\mathfrak{F}_{\theta} : \mathcal{X}^2 \to \mathcal{Z}^2$ defined by $\mathfrak{F}_{\theta}(x_1, x_2) := (f_{\theta}(x_1), f_{\theta}(x_2))$. In this way the graphical representation is

$$\begin{pmatrix} x_1 \\ x_2 \end{pmatrix} \xrightarrow{\mathfrak{F}_{\theta}} \begin{pmatrix} z_1 \\ z_2 \end{pmatrix} \xrightarrow{\mathcal{L}_y} l$$
(82)

and we can directly apply the chain rule. Before doing so, it is convenient to expand some derivatives to express the Z^2 -size matrix as two Z-size submatrixes. The Jacobian of \mathfrak{F} evaluated in (x_1, x_2) , is an operator from the tangent space $T_{\theta}\Theta$ to the tangent space $T_{(z_1, z_2)}Z^2$, which can be written in block matrix form as

$$J_{\theta}\mathfrak{F}_{\theta}(x_1, x_2) = \begin{pmatrix} J_{\theta}f_{\theta}(x_1) \\ J_{\theta}f_{\theta}(x_2) \end{pmatrix}$$
(83)

and, similarly, the hessian of $\mathcal{L}_y(z_1, z_2)$ can be written in block form as

$$\nabla_{(z_1, z_2)}^2 \mathcal{L}_y(z_1, z_2) = \begin{pmatrix} \nabla_{z_1}^2 \mathcal{L}_y(z_1, z_2) & \nabla_{z_1} \nabla_{z_2} \mathcal{L}_y(z_1, z_2) \\ \nabla_{z_2} \nabla_{z_1} \mathcal{L}_y(z_1, z_2) & \nabla_{z_2}^2 \mathcal{L}_y(z_1, z_2) \end{pmatrix}$$
(84)

The Hessian of the per-observation Contrastive loss is then

$$\nabla^2_{\theta} \mathcal{L}_y(f_{\theta}(x_1), f_{\theta}(x_2)) = \nabla^2_{\theta} \mathcal{L}_y(\mathfrak{F}_{\theta}(x_1, x_2))$$
(85)

$$\stackrel{\text{GGN}}{\approx} J_{\theta} \mathfrak{F}_{\theta}(x_1, x_2)^{\top} \cdot \nabla^2_{(z_1, z_2)} \mathcal{L}_y(z_1, z_2) \cdot J_{\theta} \mathfrak{F}_{\theta}(x_1, x_2)$$
(86)

$$= \begin{pmatrix} J_{\theta}f_{\theta}(x_1) \\ J_{\theta}f_{\theta}(x_2) \end{pmatrix}^{\top} \begin{pmatrix} \nabla_{z_1}^2 \mathcal{L}_y(z_1, z_2) & \nabla_{z_1} \nabla_{z_2} \mathcal{L}_y(z_1, z_2) \\ \nabla_{z_2} \nabla_{z_1} \mathcal{L}_y(z_1, z_2) & \nabla_{z_2}^2 \mathcal{L}_y(z_1, z_2) \end{pmatrix} \begin{pmatrix} J_{\theta}f_{\theta}(x_1) \\ J_{\theta}f_{\theta}(x_2) \end{pmatrix}$$
(87)

where $(z_1, z_2) = h_{\theta}(x_1, x_2)$ is the point where we have to evaluate the derivative wrt to z. Consequently, the Hessian of the contrastive loss is

$$\nabla^2_{\theta} \mathcal{L}(\theta; \mathcal{D}) = \sum_{\mathfrak{p}_{ij} \in \mathcal{D}^2} \nabla^2_{\theta} \mathcal{L}_{y_{ij}}(f_{\theta}(x_i), f_{\theta}(x_j))$$
(88)

$$=\sum_{\mathfrak{p}_{ij}\in\mathcal{D}^2} \begin{pmatrix} J_{\theta}f_{\theta}(x_i)\\ J_{\theta}f_{\theta}(x_j) \end{pmatrix}^{\top} \begin{pmatrix} \nabla_{z_i}^2 \mathcal{L}_{y_{ij}}(z_i, z_j) & \nabla_{z_j} \nabla_{z_j} \mathcal{L}_{y_{ij}}(z_i, z_j)\\ \nabla_{z_j} \nabla_{z_i} \mathcal{L}_{y_{ij}}(z_i, z_j) & \nabla_{z_j}^2 \mathcal{L}_{y_{ij}}(z_i, z_j) \end{pmatrix} \begin{pmatrix} J_{\theta}f_{\theta}(x_i)\\ J_{\theta}f_{\theta}(x_j) \end{pmatrix}$$
(89)

We now proceed to find the derivatives of the per-observation loss \mathcal{L}_y wrt. the Neural Network outputs $z_i = f_{\theta}(x_i)$ and $z_j = f_{\theta}(x_j)$. The derivatives varies based on the specific loss term that we consider. In the following we derive the derivatives for Euclidean \mathcal{L}_y^E and Arccos \mathcal{L}_y^A cases.

Split choice. The GGN assumes access to a composition of two functions, the specific split choice affects the result. As said in Appendix C it is common to have a normalization layer at the end of the Neural Network. This can be schematized as follow

$$\begin{pmatrix} x_1 \\ x_2 \end{pmatrix} \xrightarrow{\text{NN}} \begin{pmatrix} z_1 \\ z_2 \end{pmatrix} \xrightarrow{\ell_2 \text{-norm}} \begin{pmatrix} \frac{z_1}{\|z_1\|} \\ \frac{z_2}{\|z_2\|} \end{pmatrix} \xrightarrow{\text{distance}} l$$
(90)

and this leads to (at least) two possible split choices, wheter we include the normalization layer in the network left or right function, which we proceed to study further. We highlight that these two split choices can be interpreted as different loss function or, equivalently, as different distance metric: Euclidean or Arccos.

F.1. Euclidean distance

If we consider the ℓ_2 -normalization layer as part of the Neural Network f

$$\begin{pmatrix} x_1 \\ x_2 \end{pmatrix} \xrightarrow{\text{NN}} \begin{pmatrix} z_1 \\ z_2 \end{pmatrix} \xrightarrow{\ell_2 \text{-norm}} \begin{pmatrix} \frac{z_1}{\|z_1\|} \\ \frac{z_2}{\|z_2\|} \end{pmatrix} \xrightarrow{\text{distance}} l \tag{91}$$

then the loss $\mathcal{L}_y = \mathcal{L}_y^E =$ is the *Euclidean* distance defined as

$$\mathcal{L}_{y}^{E}(z_{i}, z_{j}) := \frac{1}{2} y \|z_{i} - z_{j}\|^{2}$$
(92)

and it holds

$$\nabla_{z_1}^2 \mathcal{L}_y^E(z_1, z_2) = \nabla_{z_2}^2 \mathcal{L}_y^E(z_1, z_2) = y\mathbb{I}$$
(93)

$$\nabla_{z_1} \nabla_{z_2} \mathcal{L}_y^E(z_1, z_2) = \nabla_{z_2} \nabla_{z_1} \mathcal{L}_y^E(z_1, z_2) = -y \mathbb{I}$$
(94)

which leads to

$$\nabla_{(z_1,z_2)}^2 \mathcal{L}_y^E(z_1,z_2) = y \begin{pmatrix} \mathbb{I} & -\mathbb{I} \\ -\mathbb{I} & \mathbb{I} \end{pmatrix}.$$
(95)

We highlight that the matrix $\begin{pmatrix} \mathbb{I} & -\mathbb{I} \\ -\mathbb{I} & \mathbb{I} \end{pmatrix}$ is positive semi-definite, this will be useful in Proposition F.3. **Intuition.** The overall hessian expression can be further simplified by considering the matrix square root.

$$\nabla^2_{\theta} \mathcal{L}(\theta; \mathcal{D}) = \sum_{\mathfrak{p}_{ij} \in \mathcal{D}^2} \nabla^2_{\theta} \mathcal{L}_{y(p_{ij})}(f_{\theta}(x_i), f_{\theta}(x_j))$$
(96)

$$=\sum_{\mathfrak{p}_{ij}\in\mathcal{D}^2}y_{ij}\begin{pmatrix}J_{\theta}f_{\theta}(x_i)\\J_{\theta}f_{\theta}(x_j)\end{pmatrix}^{\top}\begin{pmatrix}\mathbb{I} & -\mathbb{I}\\-\mathbb{I} & \mathbb{I}\end{pmatrix}\begin{pmatrix}J_{\theta}f_{\theta}(x_i)\\J_{\theta}f_{\theta}(x_j)\end{pmatrix}$$
(97)

$$=\frac{1}{2}\sum_{\mathfrak{p}_{ij}\in\mathcal{D}^2}y_{ij}\begin{pmatrix}J_{\theta}f_{\theta}(x_i)\\J_{\theta}f_{\theta}(x_j)\end{pmatrix}^{\top}\begin{pmatrix}\mathbb{I} & -\mathbb{I}\\-\mathbb{I} & \mathbb{I}\end{pmatrix}\begin{pmatrix}\mathbb{I} & -\mathbb{I}\\-\mathbb{I} & \mathbb{I}\end{pmatrix}\begin{pmatrix}J_{\theta}f_{\theta}(x_i)\\J_{\theta}f_{\theta}(x_j)\end{pmatrix}$$
(98)

$$= \frac{1}{2} \sum_{\mathfrak{p}_{ij} \in \mathcal{D}^2} y_{ij} \begin{pmatrix} J_{\theta} f_{\theta}(x_i) - J_{\theta} f_{\theta}(x_j) \\ J_{\theta} f_{\theta}(x_j) - J_{\theta} f_{\theta}(x_i) \end{pmatrix}^{\top} \begin{pmatrix} J_{\theta} f_{\theta}(x_i) - J_{\theta} f_{\theta}(x_j) \\ J_{\theta} f_{\theta}(x_j) - J_{\theta} f_{\theta}(x_i) \end{pmatrix}$$
(99)

$$= \sum_{\mathbf{p}_{ij} \in \mathcal{D}^2} y_{ij} \left(J_{\theta} f_{\theta}(x_i) - J_{\theta} f_{\theta}(x_j) \right)^{\top} \left(J_{\theta} f_{\theta}(x_i) - J_{\theta} f_{\theta}(x_j) \right).$$
(100)

This expression give raise to two interpretations. (1) This formulation draws parallels with the GGN approximation of the MSE loss. (2) It can be viewed as the squared distance of the jacobians product, where the sign y_{ij} determines the sign. This is parallel with the contrastive loss

$$\mathcal{L}(\theta; \mathcal{D}) = \frac{1}{2} \sum_{\mathfrak{p}_{ij} \in \mathcal{D}^2} y_{ij} \| f_{\theta}(x_i) - f_{\theta}(x_j) \|^2$$
(101)

$$= \frac{1}{2} \sum_{\mathfrak{p}_{ij} \in \mathcal{D}^2} y_{ij} \left(f_{\theta}(x_i) - f_{\theta}(x_j) \right)^{\top} \left(f_{\theta}(x_i) - f_{\theta}(x_j) \right)$$
(102)

with the only difference being the Jacobian operator (and the 2 factor).

F.2. Arccos distance

If we consider the ℓ_2 -normalization layer as part of the loss function \mathcal{L}_y

$$\begin{pmatrix} x_1 \\ x_2 \end{pmatrix} \xrightarrow{\text{NN}} \begin{pmatrix} z_1 \\ z_2 \end{pmatrix} \underbrace{ \stackrel{\ell_2 \text{-norm}}{\underbrace{ \begin{array}{c} z_1 \\ ||z_1|| \\ ||z_2|| \end{array}}} } \underbrace{ \begin{array}{c} \text{distance} \\ \underbrace{ \begin{array}{c} z_1 \\ ||z_2|| \end{array}} \\ \mathcal{L}_y^A \\ \end{bmatrix} } l$$
(103)

then the loss $\mathcal{L}_y = \mathcal{L}_y^A$ is the *Arccos* distance defined as

$$\mathcal{L}_{y}^{A}(z_{i}, z_{j}) := \frac{1}{2}y \left\| \frac{z_{i}}{\|z_{i}\|} - \frac{z_{j}}{\|z_{j}\|} \right\|^{2} = y \left(1 - \left\langle \frac{z_{i}}{\|z_{i}\|}, \frac{z_{j}}{\|z_{j}\|} \right\rangle \right)$$
(104)

and it holds

$$\nabla_{z_1}^2 \mathcal{L}_y^A(z_1, z_2) = \frac{y}{\|z_1\|^2} \left(\left\langle \frac{z_1}{\|z_1\|}, \frac{z_2}{\|z_2\|} \right\rangle \mathbb{I} + \frac{z_1^\top z_2 + z_2^\top z_1}{\|z_1\| \|z_2\|} - 3 \left\langle \frac{z_1}{\|z_1\|}, \frac{z_2}{\|z_2\|} \right\rangle \frac{z_1^\top z_1}{\|z_1\|^2} \right)$$
(105)

$$\nabla_{z_2} \nabla_{z_1} \mathcal{L}_y^A(z_1, z_2) = \frac{y}{\|z_1\| \|z_2\|} \left(-\mathbb{I} + \frac{z_1^\top z_1}{\|z_1\|^2} + \frac{z_2^\top z_2}{\|z_2\|^2} - \left\langle \frac{z_1}{\|z_1\|}, \frac{z_2}{\|z_2\|} \right\rangle \frac{z_1^\top z_2}{\|z_1\| \|z_2\|} \right)$$
(106)

$$\nabla_{z_2}^2 \mathcal{L}_y^A(z_1, z_2) = \frac{y}{\|z_2\|^2} \left(\left\langle \frac{z_1}{\|z_1\|}, \frac{z_2}{\|z_2\|} \right\rangle \mathbb{I} + \frac{z_1^\top z_2 + z_2^\top z_1}{\|z_1\| \|z_2\|} - 3 \left\langle \frac{z_1}{\|z_1\|}, \frac{z_2}{\|z_2\|} \right\rangle \frac{z_2^\top z_2}{\|z_2\|^2} \right)$$
(107)

F.3. Hessian approximations

In classification and regression tasks, the Generalized Gauss Newton approximation is sufficient to guarantee positive definiteness, this is because the hessian of the loss with respect to the Neural Network output $\nabla_z^2 \mathcal{L}_y(z)$ is positive definite

both for MSE and cross-entropy. The overall loss in those cases is a sum over data points, without subtraction, and thus a positive matrix. For the contrastive loss, on the other hand, positive pairs contribute positively while negative pairs contribute negatively, and thus there is no guarantee in general. Moreover, with the Arccos loss there is not even the guarantee that $\nabla_z^2 \mathcal{L}_y^A(z)$ is positive definite. Therefore, we consider three approximations of the hessian of the contrastive loss: full, positives, fixed, that ensures it to positive definite.

Full. The first possible approach is to forcefully ensure that the matrix only has positive values. This is formalized by applying an elementwise ReLU, that is

$$\left[\nabla_{\theta}^{2} \mathcal{L}(\theta; \mathcal{D})\right]_{nm} \approx \max\left(0, \left[\nabla_{\theta}^{2} \mathcal{L}(\theta; \mathcal{D})\right]_{nm}\right) \qquad \forall n, m$$
(108)

Positives. An alternative approach is to consider only the contribution of the positive pairs

$$\nabla_{\theta}^{2} \mathcal{L}(\theta; \mathcal{D}) \approx \sum_{\mathfrak{p}_{ij} \in \mathcal{D}_{\text{pos}}^{2}} \nabla_{\theta}^{2} \mathcal{L}_{y_{ij}}(f_{\theta}(x_{i}), f_{\theta}(x_{j}))$$
(109)

neglecting the contribution of the negative pairs $\sum_{\mathcal{D}_{neg}^2}$. This approximation will be far from the truth and there is no theoretical justification. This approximation strategy is inspired by Shi & Jain (2019), which only uses positive pairs to train the network responsible for predicting the variance.

Fixed. The third approach is to consider the contrastive term as a function of *one data point at a time*, assuming the other fixed. This idea can be formalized by making use of the stop gradient notation: sg. The per-observation contrastive Eq. 22 can be written as

$$\mathcal{L}_{y}(z_{1}, z_{2}) = \frac{1}{2} \mathcal{L}_{y}(\mathrm{sg}[z_{1}], z_{2}) + \frac{1}{2} \mathcal{L}_{y}(z_{1}, \mathrm{sg}[z_{2}]).$$
(110)

We highlight that, with this definition, the zero- and first-order derivative does not change, and thus the loss and gradient are exactly the same as for the standard contrastive loss. However, the second-order derivative loses the cross term

$$\nabla_{(z_1, z_2)}^2 \mathcal{L}_y(z_1, z_2) = \begin{pmatrix} \nabla_{z_1}^2 \mathcal{L}_y(z_1, z_2) & 0\\ 0 & \nabla_{z_2}^2 \mathcal{L}_y(z_1, z_2) \end{pmatrix}$$
(111)

and the Hessian of the loss can then be more compactly written as

$$\nabla^{2}_{\theta}\mathcal{L}(\theta;\mathcal{D}) = \sum_{\mathfrak{p}_{ij}\in\mathcal{D}^{2}} \nabla^{2}_{\theta}\mathcal{L}_{y_{ij}}(f_{\theta}(x_{i}), f_{\theta}(x_{j}))$$

$$= \sum_{\mathfrak{p}_{ij}\in\mathcal{D}^{2}} y_{ij} \begin{pmatrix} J_{\theta}f_{\theta}(x_{i}) \\ J_{\theta}f_{\theta}(x_{j}) \end{pmatrix}^{\top} \begin{pmatrix} \nabla^{2}_{z_{i}}\mathcal{L}_{y_{ij}}(z_{i}, z_{j}) & 0 \\ 0 & \nabla^{2}_{z_{j}}\mathcal{L}_{y_{ij}}(z_{i}, z_{j}) \end{pmatrix} \begin{pmatrix} J_{\theta}f_{\theta}(x_{i}) \\ J_{\theta}f_{\theta}(x_{j}) \end{pmatrix}$$

$$= 2 \sum_{\mathfrak{p}_{ij}\in\mathcal{D}^{2}} y_{ij} J_{\theta}f_{\theta}(x_{i})^{\top} \nabla^{2}_{z_{i}}\mathcal{L}_{y_{ij}}(z_{i}, z_{j}) J_{\theta}f_{\theta}(x_{i}).$$
(112)

Proposition F.1. Consider the *full* approximation. Using only the diagonal of the hessian \implies positive definiteness.

Proof. The eigenvalues of a diagonal matrix are exactly the values on the diagonal. Enforcing these elements to be positive is equivalent to enforcing that all eigenvalues are positive. This implies the positive definiteness of the matrix. \Box

Proposition F.2. Consider the **positives** approximation. Euclidean loss \implies positive definiteness.

Proof. Consider the hessian of Eq. 109 and substitute the expression Eq. 95

$$\nabla^2_{\theta} \mathcal{L}(\theta; \mathcal{D}) \approx \sum_{\mathfrak{p}_{ij} \in \mathcal{D}^2_{\text{nos}}} \nabla^2_{\theta} \mathcal{L}_{y_{ij}}(f_{\theta}(x_i), f_{\theta}(x_j))$$
(113)

$$= \sum_{\mathfrak{p}_{ij} \in \mathcal{D}^2_{\text{pos}}} y_{ij} \begin{pmatrix} J_{\theta} f_{\theta}(x_i) \\ J_{\theta} f_{\theta}(x_j) \end{pmatrix}^{\top} \begin{pmatrix} \mathbb{I} & -\mathbb{I} \\ -\mathbb{I} & \mathbb{I} \end{pmatrix} \begin{pmatrix} J_{\theta} f_{\theta}(x_i) \\ J_{\theta} f_{\theta}(x_j) \end{pmatrix}$$
(114)

$$= \frac{1}{|\mathcal{D}_{\text{pos}}^2|} \sum_{\mathfrak{p}_{ij} \in \mathcal{D}_{\text{pos}}^2} \begin{pmatrix} J_{\theta} f_{\theta}(x_i) \\ J_{\theta} f_{\theta}(x_j) \end{pmatrix}^{\top} \begin{pmatrix} \mathbb{I} & -\mathbb{I} \\ -\mathbb{I} & \mathbb{I} \end{pmatrix} \begin{pmatrix} J_{\theta} f_{\theta}(x_i) \\ J_{\theta} f_{\theta}(x_j) \end{pmatrix}$$
(115)

which is a sum of positive definite matrixes.

Proposition F.3. Consider the fixed approximation. Euclidean loss \implies positive definiteness.

Proof. Consider the hessian of Eq. 112 and substitute the expression Eq. 95

$$\nabla_{\theta}^{2} \mathcal{L}(\theta; \mathcal{D}) = 2 \sum_{\mathfrak{p}_{ij} \in \mathcal{D}^{2}} y_{ij} J_{\theta} f_{\theta}(x_{i})^{\top} J_{\theta} f_{\theta}(x_{i})$$
(116)

$$=2\sum_{x_i,c_i\in\mathcal{D}}\sum_{x_j,c_j\in\mathcal{D}}y_{ij}J_{\theta}f_{\theta}(x_i)^{\top}J_{\theta}f_{\theta}(x_i)$$
(117)

$$= 2 \sum_{x_i, c_i \in \mathcal{D}} \left(\sum_{x_j, c_j \in \mathcal{D}} y_{ij} \right) J_{\theta} f_{\theta}(x_i)^{\top} J_{\theta} f_{\theta}(x_i)$$
(118)

$$= 2 \sum_{x_i, c_i \in \mathcal{D}} \left(\frac{|\mathcal{D}_{\text{pos}}^2(x_i)|}{|\mathcal{D}_{\text{pos}}^2(x_i)|} - \frac{|\mathcal{D}_{\text{neg inside}}^2(x_i)|}{|\mathcal{D}_{\text{neg}}^2(x_i)|} \right) J_{\theta} f_{\theta}(x_i)^{\top} J_{\theta} f_{\theta}(x_i)$$
(119)

$$= 2 \sum_{x_i, c_i \in \mathcal{D}} \underbrace{\left(1 - \frac{|\mathcal{D}_{\text{neg inside}}^2(x_i)|}{|\mathcal{D}_{\text{neg}}^2(x_i)|}\right)}_{\ge 0} J_{\theta} f_{\theta}(x_i)^{\top} J_{\theta} f_{\theta}(x_i)$$
(120)

where $|\mathcal{D}_{pos}^2(x_i)|$ is the number of positive pairs containing the point x_i , and similarly for $|\mathcal{D}_{reg}^2(x_i)|$ and $|\mathcal{D}_{reg inside}^2(x_i)|$. The hessian is then equivalent to a sum of positive definite matrixes multiplied by a non-negative factor, and thus it is positive definite.

We highlight, as we can see from the last Proposition's proof, that if all the negatives are inside the margin, the hessian is exactly 0. The more negatives are outside the margin, the more positive the hessian is.

G. Experimental details

In this section, we provide details on the experiments. We highlight that code for all experiments and baselines are available at https://github.com/FrederikWarburg/bayesian-metric-learning.

G.1. Datasets

FashionMnist. We use the standard test-train split for FashionMnist (Xiao et al., 2017) and for MNIST (LeCun et al., 1998). We normalize the images in the range [0, 1] and do not perform any data augmentation during training. We train with RMSProp with learning rate 10^{-3} and default PyTorch settings, and with exponential learning rate decay with $\gamma = \exp(-0.1)$. We use a memory factor of 0.0001 and maximum 5000 pairs per train step. We trained for 20 epochs.

CIFAR10. We use the standard test-train split for CIFAR () and for SVHN (). We normalize the images in the range [0, 1] and do not perform any data augmentation during training. We train with RMSProp with learning rate 10^{-5} and default PyTorch settings, and with exponential learning rate decay with $\gamma = \exp(-0.1)$. We use a memory factor of 0.0001 and maximum 5000 pairs per train step. We trained for 20 epochs.

CUB200. The CUB-200-2011 dataset (Wah et al., 2011) consists of 11788 images of 200 bird species. The birds are captured from different perspectives and in different environments, making this a challenging dataset for image retrieval. We follow the procedure of Musgrave et al. (2020) and divide the first 100 classes into the training set and the last 100 classes into the test set. In this zero-shot setting, the trained models have not seen any of the bird species in the test set, and the learned features must generalize well across species. Similarly to Warburg et al. (2021), we use the Stanford Car-196 (Krause et al., 2013) dataset as OoD data. The Car-196 dataset is composed of 16185 images of 196 classes of cars. We conduct a similar split as Musgrave et al. (2020), and only evaluate on the last 98 classes. This constitutes a very challenging zero-shot OoD dataset, where the model at test time needs to distinguish cars from birds, however, the model has not seen any of the bird species at training time. We use imagenet normalization and during training augment with random resized crops and random horizontal flipping. We image 224x224 image resolution. We train with RMSProp with

learning rate 10^{-7} and default PyTorch settings, and with exponential learning rate decay with $\gamma = \exp(-0.1)$. We use a memory factor of 0.0001 and maximum 30 pairs per train step. We trained for 20 epochs.

LFW. We use the face recognition dataset LFW (Huang et al., 2007) with the standard zero-shot train/test split, CUB200 as OoD data. It is challenging because of the high number of classes and few observations per class. LFW (Huang et al., 2007) consists of 13 000 images of 5 749 people using the standard zero-shot train/test split. One reason reliable uncertainties are important for face recognition systems is to avoid granting access based on an erroneous prediction. An example of such failure happened with the initial release of the Apple Face ID software, which failed to recognize underrepresented groups that were missing or underrepresented in the training distribution (The Week, 2018). Reliable uncertainties and OoD detection might have mitigated such issues. We use imagenet normalization and during training augment with random resized crops and random horizontal flipping. We image 224x224 image resolution. We train with RMSProp with learning rate 10^{-7} and default PyTorch settings, and with exponential learning rate decay with $\gamma = \exp(-0.1)$. We use a memory factor of 0.0001 maximum 30 pairs per train step. We trained for 200 epochs.

MSLS. We use standard zero-shot train/val/test split (Warburg et al., 2020). We image 224x224 image resolution. We train with RMSProp with learning rate 10^{-7} and default PyTorch settings, and with exponential learning rate decay with $\gamma = \exp(-0.1)$. We use a memory factor of 0.0001 maximum 10 pairs per train step. We trained for 200 epochs.

G.2. Model Architectures

For all experiments, we use a last-layer diagonal LA. Across all experiments, our networks follow standard practices in image retrieval. For FashionMnist, the network is [conv2d(1, 32), relu, conv2d(32, 64), relu, maxpool2d(2), Flatten, Linear(9216)] and for CiFAR10, the network is [conv2d(1, 32), relu, conv2d(32, 64), relu, maxpool2d(2), conv2d(64, 64), relu, Flatten, Linear(9216)]. For CUB200 and LFW, we use a pretrained ResNet50 backbone followed by a Generalized-Mean pooling layer and a dimension-preserving learned whitening layer (linear layer). The weight posterior is learned only for the last layer

G.3. Metrics

Evaluating the models' uncertainty estimates is tricky. We propose four metrics that capture both *interpolation* and *extrapolation* behaviors uncertainty estimates. To measure *extrapolation* behavior, we measure the performance of out-ofdistribution (OOD) detection and report Area Under Receiver Operator Curve (AUROC) and Area Under Precision-Recall Curve (AUPRC). These metrics describe the model's ability to assign high uncertainty to observations it has not seen during training (e.g., we train a model on birds and use images of cars as out-of-distribution examples), and are often used in unsupervised representation learning (Miani et al., 2022). To measure *interpolation* behavior, we measure the models' ability to assign reliable uncertainties to in-distribution data. Similarly to Wang & Deng (2021), we measure the Area Under the Sparsification Curve. The sparsification curve is computed by iteratively removing the observation with the highest variance and recomputing the predictive performance (mAP@1). Lastly, for closed set datasets, we propose a method to compute Expected Calibration Error (ECE) for image retrieval. This is done by sampling 100 latent variables z_i from the predicted latent distributions p(z|I). We then perform nearest neighbor classification on these samples to obtain c_i predictions. We use the mode as the model prediction and the consensus with mode as the confidence, e.g., if they constitute 60 samples, then the prediction has 60% confidence. We compare the estimated confidence with accuracy with

$$ECE = \sum_{i}^{N} \frac{1}{|B_i|} |\operatorname{acc}(B_i) - \operatorname{conf}(B_i)|, \qquad (121)$$

where we $acc(B_i)$ and $conf(B_i)$ is the accuracy and confidence of the *i*th bin.

G.4. Details on baselines

If nothing else is stated we use the same architecture and training/evaluation code for all models. We also release all baseline models.

- Deterministic: we train with the contrastive loss.
- MC dropout: we train with dropout between all trainable layers (except for the ResNet50 blocks). We use a dropout rate 0.2. Dropout was enabled during test time.

- Deep Ensemble: Each ensemble consists of 5 models, each initialized with different seeds.
- PFE: We add an uncertainty module that learns the variance. The rest of the model is frozen during training and initialized with a deterministically trained model. The uncertainty model consists of [linear, batchnorm1d, relu, linear, batchnorm1d]. We experimented with the last batch normalization (BN) layer sharing parameters or just a standard BN layer. We found the former to work slightly better.
- HIB: We use the same uncertainty module as in PFE, and a kl weight of 10⁻⁴. We use 8 samples during training to compute the loss.