



Yang, X. , Guo, Y., Dong, M. and Xue, J.-H. (2022) Towards certified robustness of distance metric learning. *IEEE Transactions on Neural Networks and Learning Systems*, 35(3), pp. 3834-3844. (doi: [10.1109/TNNLS.2022.3199902](https://doi.org/10.1109/TNNLS.2022.3199902))

The material cannot be used for any other purpose without further permission of the publisher and is for private use only.

There may be differences between this version and the published version. You are advised to consult the publisher's version if you wish to cite from it.

<https://eprints.gla.ac.uk/277559/>

Deposited on 25 August 2022

Enlighten – Research publications by members of the University of
Glasgow

<http://eprints.gla.ac.uk>

Towards Certified Robustness of Distance Metric Learning

Xiaochen Yang, Yiwen Guo, Mingzhi Dong and Jing-Hao Xue, *Senior Member, IEEE*

Abstract—Metric learning aims to learn a distance metric such that semantically similar instances are pulled together while dissimilar instances are pushed away. Many existing methods consider maximizing or at least constraining a distance margin in the *feature* space that separates similar and dissimilar pairs of instances to guarantee their generalization ability. In this paper, we advocate imposing an adversarial margin in the *input* space so as to improve the generalization and robustness of metric learning algorithms. We first show that, the adversarial margin, defined as the distance between training instances and their closest adversarial examples in the input space, takes account of both the distance margin in the feature space and the correlation between the metric and triplet constraints. Next, to enhance robustness to instance perturbation, we propose to enlarge the adversarial margin through minimizing a derived novel loss function termed the perturbation loss. The proposed loss can be viewed as a data-dependent regularizer and easily plugged into any existing metric learning methods. Finally, we show that the enlarged margin is beneficial to the generalization ability by using the theoretical technique of algorithmic robustness. Experimental results on 16 datasets demonstrate the superiority of the proposed method over existing state-of-the-art methods in both discrimination accuracy and robustness against possible noise.

Index Terms—Metric learning, Nearest neighbor, Adversarial perturbation, Robustness, Generalization ability

I. INTRODUCTION

METRIC learning focuses on learning similarity or dissimilarity between data. Research on metric learning originates from at least 2002, where [1] first proposes to formulate it as an optimization problem. Since then, many metric learning methods have been proposed for classification [2]–[4], clustering [5], and information retrieval [6], [7]. In particular, the methods have shown to be particularly superior in open-set classification and few-shot classification with notable applications in, for example, face verification [8], [9] and person re-identification [10], [11].

One commonly studied distance metric is the generalized Mahalanobis distance, which defines the distance between any two instances $\mathbf{x}_i, \mathbf{x}_j \in \mathbb{R}^p$ as

$$d_M(\mathbf{x}_i, \mathbf{x}_j) = \sqrt{(\mathbf{x}_i - \mathbf{x}_j)^T \mathbf{M} (\mathbf{x}_i - \mathbf{x}_j)},$$

X. Yang is with the School of Mathematics and Statistics, University of Glasgow, Glasgow, UK. E-mail: xiaochen.yang@glasgow.ac.uk.

Y. Guo is an independent researcher.

M. Dong is with the School of Computer Science, Fudan University, and Shanghai Key Laboratory of Data Science, Fudan University, Shanghai, China.

J.-H. Xue is with the Department of Statistical Science, University College London, London, UK.

The first two authors contribute equally to this work.

Manuscript received; revised.

where M is a positive semidefinite (PSD) matrix. Owing to its PSD property, M can be decomposed into $L^T L$. Thus, computing the Mahalanobis distance is equivalent to linearly transforming the instances from the input space to the feature space via L and then computing the Euclidean distance $\|L\mathbf{x}_i - L\mathbf{x}_j\|_2$ in the transformed space.

To learn a specific distance metric for each task, prior knowledge on instance similarity and dissimilarity should be provided as side information. Metric learning methods differ by the form of side information they use and the supervision encoded in similar and dissimilar pairs. For example, pairwise constraints enforce the distance between instances of the same class to be small (or smaller than a threshold value) and the distance between instances of different classes to be large (or larger than a threshold value) [1], [5]. The thresholds could be either pre-defined or learned for similar and dissimilar pairs [12], [13]. In triplet constraints $(\mathbf{x}_i, \mathbf{x}_j, \mathbf{x}_l)$, distance between the different-class pair $(\mathbf{x}_i, \mathbf{x}_l)$ should be larger than distance between the same-class pair $(\mathbf{x}_i, \mathbf{x}_j)$, and typically, plus a margin [14]–[17]. More recently, quadruplet constraints are proposed, which require the difference in the distance of two pairs of instances to exceed a margin [18], and $(N + 1)$ -tuple constraint extends the triplet constraint for multi-class classification [19].

The gap between thresholds in pairwise constraints and the margin in triplet and quadruplet constraints are both designed to learn a distance metric that could ensure good generalization of the subsequent k -nearest neighbor (k NN) classifier. However, such a distance margin imposed in the feature space does not consider the correlation between the data and the learned metric. Consequently, it may be insufficient to withstand a small perturbation of the instance occurred in the input space, thereby failing to certify the robustness or even possess the anticipated generalization benefit. As illustrated in Fig. 1(upper), while \mathbf{x}_i selects the same-class instance \mathbf{x}_j as its nearest neighbor in the feature space, a tiny perturbation from \mathbf{x}_i to \mathbf{x}'_i in the input space can be magnified by the learned distance metric, leading to a change in its nearest neighbor from \mathbf{x}_j to the different-class instance \mathbf{x}_l . When the NN algorithm is used as the classifier, the perturbation results in an incorrect label prediction.

In this paper, we propose a simple yet effective method to enhance the generalization ability of metric learning algorithms and their robustness against instance perturbation. As shown in Fig. 1(bottom), the principal idea is to enlarge the adversarial margin, defined as the distance between a training instance and its closest adversarial example in the input space [20].

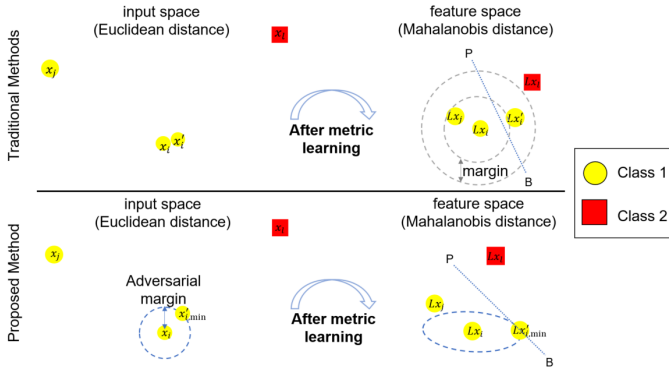


Fig. 1: *Upper*: Traditional methods aim to separate the same-class pair $(\mathbf{x}_i, \mathbf{x}_j)$ and the different-class pair $(\mathbf{x}_i, \mathbf{x}_l)$ by a margin in the feature space. While \mathbf{x}_i has \mathbf{x}_j as its nearest neighbor (NN) in the feature space and is correctly predicted by using the NN classifier, the metric is sensitive to perturbation in the input space; a tiny perturbation from \mathbf{x}_i to \mathbf{x}'_i changes the NN to \mathbf{x}_l and leads to an incorrect prediction. *Bottom*: The proposed method aims to enlarge the adversarial margin in the input space, which equals to the Euclidean distance between \mathbf{x}_i and the closest point $\mathbf{x}_{i,\min}$ in the input space that lies on the decision boundary in the feature space (indicated by PB) and quantifies the maximum degree to which robustness can be certified.

In particular, our contributions are fourfold.

- 1) We identify that the distance margin, widely used in existing methods, is insufficient to withstand adversarial examples, and we introduce a direct measure of robustness termed the adversarial margin, which quantifies the maximum degree to which a training instance could be perturbed without changing the label of its nearest neighbor (or k nearest neighbors if required) in the feature space. Building on a geometric insight, we derive an analytically simple solution to the adversarial margin, which reveals the importance of an adaptive margin considering the correlation between the data and the distance metric (Section II-A-II-B).
- 2) We define a novel hinge-like perturbation loss to penalize the adversarial margin for being small. The proposed loss function serves as a general approach to enhancing robustness, as it can be optimized jointly with any existing triplet-based metric learning methods; the optimization problem suggests that our method learns a discriminative metric in a weighted manner and simultaneously functions as a data-dependent regularization (Section II-C).
- 3) We show the benefit of enlarging the adversarial margin to the generalization ability of the learned distance metric by using the theoretical technique of algorithmic robustness [21] (Theorem 1, Section II-D).
- 4) We conduct experiments on 16 datasets in both noise-free and noisy settings. Results show that the proposed method outperforms state-of-the-art robust metric learning methods in terms of classification accuracy and validate its robustness to possible noise in the input

space (Section IV).

Notation: Let $\{\mathbf{x}_i, y_i\}_{i=1}^n$ denote the set of training instance and label pairs, where $\mathbf{x}_i \in \mathcal{X} \subseteq \mathbb{R}^p$ and $y_i \in \mathcal{Y} = \{1, \dots, C\}$; \mathcal{X} is called the input space. Our framework is based on triplet constraints $\{\mathbf{x}_i, \mathbf{x}_j, \mathbf{x}_l\}$ and we adopt the following strategy for generating triplets [14]:

$$\mathcal{S} = \{(\mathbf{x}_i, \mathbf{x}_j) : \mathbf{x}_j \in \{k\text{NNs with the same class label of } \mathbf{x}_i\}\},$$

$$\mathcal{R} = \{(\mathbf{x}_i, \mathbf{x}_j, \mathbf{x}_l) : (\mathbf{x}_i, \mathbf{x}_j) \in \mathcal{S}, y_i \neq y_l\}.$$

\mathbf{x}_j is termed the target neighbor of \mathbf{x}_i and \mathbf{x}_l is termed the impostor. $|\mathcal{S}|$ and $|\mathcal{R}|$ denote the numbers of elements in the sets \mathcal{S} and \mathcal{R} , respectively. d_E and d_M denote the Euclidean and Mahalanobis distances, respectively; $\mathbf{M} \in \mathbb{S}_+^p$, where \mathbb{S}_+^p is the cone of $p \times p$ real-valued PSD matrices. $\mathbf{M}^2 = \mathbf{M}\mathbf{M}$. $\mathbb{1}[\cdot]$ denotes the indicator function and $[a]_+ = \max(a, 0)$ for $a \in \mathbb{R}$.

II. METHODOLOGY

In this section, we introduce our method for enhancing robustness of triplet-based metric learning algorithms through maximizing the adversarial margin. First, we review the existing distance margin and provide the rationale for enlarging the adversarial margin. Second, an explicit formula for the adversarial margin is derived. Third, we propose the perturbation loss to encourage a larger adversarial margin and present its optimization jointly with the existing large (distance) margin nearest neighbor (LMNN) algorithm. Lastly, we show that enlarging the adversarial margin is beneficial to the generalization ability of the learned distance metric.

A. Motivation for enlarging the adversarial margin

Suppose \mathbf{x}_i is a training instance and $\mathbf{x}_j, \mathbf{x}_l$ are the nearest neighbor of \mathbf{x}_i from the same class and from the different class respectively. Many triplet-based methods, such as LMNN [14], impose the following constraint on the triplet:

$$f(\mathbf{x}_i) := d_M^2(\mathbf{x}_i, \mathbf{x}_l) - d_M^2(\mathbf{x}_i, \mathbf{x}_j) \geq 1.$$

When the constraint is satisfied, \mathbf{x}_i will be correctly classified using the NN classifier. Moreover, the value one represents the unit margin at the distance level and is designed to robustify the model against small noises in training instances.

Nevertheless, the distance margin may be insufficient to withstand deliberately manipulated perturbations. Let $\Delta\mathbf{x}_i$ denote a perturbation of \mathbf{x}_i . When the perturbation size is constrained as $\|\Delta\mathbf{x}_i\|_2 \leq r$, $f(\mathbf{x}_i + \Delta\mathbf{x}_i)$ decreases the most from $f(\mathbf{x}_i)$ if $\Delta\mathbf{x}_i$ is chosen in the direction of $\mathbf{M}(\mathbf{x}_l - \mathbf{x}_j)$: $f(\mathbf{x}_i + \Delta\mathbf{x}_i) - f(\mathbf{x}_i) = 2\Delta\mathbf{x}_i^T \mathbf{M}(\mathbf{x}_j - \mathbf{x}_l) = -2r\|\mathbf{M}(\mathbf{x}_l - \mathbf{x}_j)\|_2$. Therefore, in order to correctly classify the perturbed instance $\mathbf{x}_i + \Delta\mathbf{x}_i$, it is required that $f(\mathbf{x}_i + \Delta\mathbf{x}_i)$ is positive, that is, $\|\mathbf{M}(\mathbf{x}_l - \mathbf{x}_j)\|_2$ should be small. One way to reduce this value is by regularizing the spectral norm of \mathbf{M} . However, it is demanding for the metric to satisfy the large distance margin for all triplets and meanwhile keep a small spectral norm.

To achieve robustness against instance perturbation, we suggest an alternative way by maximizing the adversarial margin, defined as the distance between the training instance

and its closest *adversarial example* [20]. More concretely, an adversarial example is a perturbed point whose nearest neighbor, identified based on the learned Mahalanobis distance, changes from an instance of the same class to one of a different class; consequently, it will be misclassified by the NN classifier and increase the risk of misclassification by k NN. In terms of previous notations, an adversarial example is a perturbed point $\mathbf{x}_i + \Delta\mathbf{x}_i$ such that $f(\mathbf{x}_i + \Delta\mathbf{x}_i) < 0$. If all adversarial examples of an instance are far away from the instance itself, i.e. there is no $\Delta\mathbf{x}_i$ such that $\|\Delta\mathbf{x}_i\|_2 \leq r$ and $f(\mathbf{x}_i + \Delta\mathbf{x}_i) < 0$, a high degree of robustness is achieved. Building on this rationale, we will first find the closest adversarial example and then push this point away from the training instance. Moreover, since the test instance can be regarded as a perturbed copy of training instances [21], improving robustness on correctly classified training instances also helps enhance the generalization ability of the learned metric.

B. Derivation of adversarial margin

We start by deriving a closed-form solution to the closest adversarial example. Given a training instance \mathbf{x}_i and the associated triplet constraint $(\mathbf{x}_i, \mathbf{x}_j, \mathbf{x}_l)$, we aim to find the closest point $\mathbf{x}_{i,\min}$ to \mathbf{x}_i in the input space that lies on the decision boundary formed by \mathbf{x}_j and \mathbf{x}_l in the feature space. Note that closeness is defined in the input space and will be calculated using the Euclidean distance since we target at changes on the original feature of an instance; and that the decision boundary is found in the feature space since k NNs are identified by using the Mahalanobis distance. Mathematically, we can formulate the closest adversarial example $\mathbf{x}_{i,\min}$ as follows:

$$\begin{aligned} \mathbf{x}_{i,\min} &= \arg \min_{\mathbf{x}'_i \in \mathbb{R}^p} (\mathbf{x}'_i - \mathbf{x}_i)^T (\mathbf{x}'_i - \mathbf{x}_i) \\ \text{s.t. } & (\mathbf{L}\mathbf{x}'_i - \frac{\mathbf{L}\mathbf{x}_j + \mathbf{L}\mathbf{x}_l}{2})^T (\mathbf{L}\mathbf{x}_l - \mathbf{L}\mathbf{x}_j) = 0 \end{aligned} \quad (1)$$

The objective function of Eq. 1 corresponds to minimizing the Euclidean distance from the training instance \mathbf{x}_i . The constraint represents the decision boundary, which is the perpendicular bisector of points $\mathbf{L}\mathbf{x}_j$ and $\mathbf{L}\mathbf{x}_l$. In other words, it is a hyperplane that is perpendicular to the line joining points $\mathbf{L}\mathbf{x}_j$ and $\mathbf{L}\mathbf{x}_l$ and passes their midpoint $\frac{\mathbf{L}\mathbf{x}_j + \mathbf{L}\mathbf{x}_l}{2}$; all points on the hyperplane are equidistant from $\mathbf{L}\mathbf{x}_j$ and $\mathbf{L}\mathbf{x}_l$.

Since Eq. 1 minimizes a convex quadratic function with an equality constraint, we can find an explicit formula for $\mathbf{x}_{i,\min}$ by using the method of Lagrangian multipliers; detailed derivation is provided in Appendix A:

$$\mathbf{x}_{i,\min} = \mathbf{x}_i + \frac{\left(\frac{\mathbf{x}_j + \mathbf{x}_l}{2} - \mathbf{x}_i\right)^T \mathbf{M}(\mathbf{x}_l - \mathbf{x}_j)}{(\mathbf{x}_l - \mathbf{x}_j)^T \mathbf{M}^2(\mathbf{x}_l - \mathbf{x}_j)} \mathbf{M}(\mathbf{x}_l - \mathbf{x}_j). \quad (2)$$

With the solution of $\mathbf{x}_{i,\min}$, we can now calculate the squared Euclidean distance between \mathbf{x}_i and $\mathbf{x}_{i,\min}$:

$$d_E^2(\mathbf{x}_i, \mathbf{x}_{i,\min}) = \frac{(d_M^2(\mathbf{x}_i, \mathbf{x}_l) - d_M^2(\mathbf{x}_i, \mathbf{x}_j))^2}{4d_{M^2}^2(\mathbf{x}_j, \mathbf{x}_l)}. \quad (3)$$

For clarity, we will call $d_E(\mathbf{x}_i, \mathbf{x}_{i,\min})$ the *adversarial margin*, in contrast to the distance margin as in LMNN. It represents

the maximum amount of tolerance for perturbation while retaining prediction correctness. The numerator of Eq. 3 is the square of the standard distance margin, and the denominator is the squared L_2 -norm of $\mathbf{M}(\mathbf{x}_l - \mathbf{x}_j)$. Therefore, in order to achieve a large adversarial margin, the metric should push \mathbf{x}_l away from the neighborhood of \mathbf{x}_i by expanding the distance in the direction that has a small correlation with $\mathbf{x}_l - \mathbf{x}_j$ (the optimal direction is orthogonal to $\mathbf{x}_l - \mathbf{x}_j$).

Remark 1: The objective function in Eq. 1 defines a hypersphere in the input space, which characterizes perturbations of equal magnitude in all directions, e.g., isotropic Gaussian noise. To model heterogeneous and correlated perturbation, we can extend the objective function by defining an arbitrary oriented hyperellipsoid, as discussed in Appendix A.

C. Metric learning via minimizing the perturbation loss

To improve robustness of distance metric, we design a perturbation loss to promote an increase in the adversarial margin. Two situations need to be distinguished here. Firstly, when the nearest neighbor of \mathbf{x}_i is an instance from the same class, we will penalize a small adversarial margin by using the hinge loss $[\tau^2 - d_E^2(\mathbf{x}_i, \mathbf{x}_{i,\min})]_+$. The reasons are that (a) the adversarial margin is generally smaller for hard instances that are close to the class boundary in contrast to those locating far away and (b) it is these hard instances that are more vulnerable to perturbation and demand an improvement in their robustness. Therefore, we introduce τ for directing attention to hard instances and controlling the desired margin. Secondly, in the other situation where the nearest neighbor of \mathbf{x}_i belongs to a different class, metric learning should focus on satisfying the distance requirement specified in the triplet constraint. In this case, we simply assign a large penalty of τ^2 to promote a non-increasing loss function. Integrating these two situations, we propose the following perturbation loss:

$$\begin{aligned} J_p &= \frac{1}{|\mathcal{R}|} \sum_{\mathcal{R}} \{[\tau^2 - \tilde{d}_E^2(\mathbf{x}_i, \mathbf{x}_{i,\min})]_+ \mathbb{1}[d_M^2(\mathbf{x}_i, \mathbf{x}_l) > d_M^2(\mathbf{x}_i, \mathbf{x}_j)] \\ &\quad + \tau^2 \mathbb{1}[d_M^2(\mathbf{x}_i, \mathbf{x}_l) \leq d_M^2(\mathbf{x}_i, \mathbf{x}_j)]\}, \end{aligned} \quad (4)$$

where $\sum_{\mathcal{R}}$ is an abbreviation for $\sum_{(\mathbf{x}_i, \mathbf{x}_j, \mathbf{x}_l) \in \mathcal{R}}$. To prevent the denominator of Eq. 3 from being zero, which may happen when different-class instances \mathbf{x}_j and \mathbf{x}_l are close to each other, we add a small constant ϵ ($\epsilon = 1e-10$) to the denominator; that is, $\tilde{d}_E^2(\mathbf{x}_i, \mathbf{x}_{i,\min}) = \frac{(d_M^2(\mathbf{x}_i, \mathbf{x}_l) - d_M^2(\mathbf{x}_i, \mathbf{x}_j))^2}{4(d_{M^2}^2(\mathbf{x}_j, \mathbf{x}_l) + \epsilon)}$.

The proposed perturbation loss can be readily included in the objective function of any metric learning methods and is particularly useful to triplet-based methods. When the same triplet set is used for supervising metric learning and deriving adversarial examples, our method can encourage the triplets to meet the distance margin by learning a discriminative metric. For this reason, we adapt LMNN as an example for its wide use and effective classification performance. The objective

function of LMNN with the perturbation loss is as follows:

$$\begin{aligned} \min_{\mathbf{M} \in \mathbb{S}_+^p} J &= J_{\text{LMNN}} + \lambda J_{\text{P}}, \\ J_{\text{LMNN}} &= (1 - \mu) \frac{1}{|\mathcal{S}|} \sum_{\mathcal{S}} d_{\mathbf{M}}^2(\mathbf{x}_i, \mathbf{x}_j) \\ &\quad + \mu \frac{1}{|\mathcal{R}|} \sum_{\mathcal{R}} [1 + d_{\mathbf{M}}^2(\mathbf{x}_i, \mathbf{x}_j) - d_{\mathbf{M}}^2(\mathbf{x}_i, \mathbf{x}_l)]_+, \end{aligned} \quad (5)$$

where $\sum_{\mathcal{S}}$ stands for $\sum_{(\mathbf{x}_i, \mathbf{x}_j) \in \mathcal{S}}$. The weight parameter $\lambda > 0$ controls the importance of perturbation loss (J_{P}) relative to the loss function of LMNN (J_{LMNN}). $\mu \in (0, 1)$ balances the impacts between pulling together target neighbors and pushing away impostors.

We adopt the projected gradient descent algorithm to solve the above optimization problem. The gradient of J_{P} and J_{LMNN} are given as follows:

$$\begin{aligned} \frac{\partial J_{\text{P}}}{\partial \mathbf{M}} &= \frac{1}{|\mathcal{R}|} \sum_{\mathcal{R}} \alpha_{ijl} \left\{ \frac{d_{\mathbf{M}}^2(\mathbf{x}_i, \mathbf{x}_l) - d_{\mathbf{M}}^2(\mathbf{x}_i, \mathbf{x}_j)}{2(d_{\mathbf{M}^2}^2(\mathbf{x}_j, \mathbf{x}_l) + \epsilon)} (\mathbf{X}_{ij} - \mathbf{X}_{il}) \right. \\ &\quad \left. + \frac{(d_{\mathbf{M}}^2(\mathbf{x}_i, \mathbf{x}_l) - d_{\mathbf{M}}^2(\mathbf{x}_i, \mathbf{x}_j))^2}{4(d_{\mathbf{M}^2}^2(\mathbf{x}_j, \mathbf{x}_l) + \epsilon)^2} (\mathbf{M}\mathbf{X}_{jl} + \mathbf{X}_{jl}\mathbf{M}) \right\}, \\ \frac{\partial J_{\text{LMNN}}}{\partial \mathbf{M}} &= \frac{1 - \mu}{|\mathcal{S}|} \sum_{\mathcal{S}} \mathbf{X}_{ij} + \frac{\mu}{|\mathcal{R}|} \sum_{\mathcal{R}} \beta_{ijl} (\mathbf{X}_{ij} - \mathbf{X}_{il}), \end{aligned} \quad (6)$$

where $\alpha_{ijl} = \mathbb{1}[d_{\mathbf{M}}^2(\mathbf{x}_i, \mathbf{x}_l) > d_{\mathbf{M}}^2(\mathbf{x}_i, \mathbf{x}_j), \tilde{d}_{\mathbf{E}}(\mathbf{x}_i, \mathbf{x}_{i,\min}) \leq \tau]$, $\beta_{ijl} = \mathbb{1}[1 + d_{\mathbf{M}}^2(\mathbf{x}_i, \mathbf{x}_j) - d_{\mathbf{M}}^2(\mathbf{x}_i, \mathbf{x}_l) \geq 0]$; $\mathbf{X}_{ij} = (\mathbf{x}_i - \mathbf{x}_j)(\mathbf{x}_i - \mathbf{x}_j)^T$ and $\mathbf{X}_{il}, \mathbf{X}_{jl}$ are defined similarly. The gradient of J_{P} is a sum of two descent directions. The first direction $\mathbf{X}_{ij} - \mathbf{X}_{il}$ agrees with LMNN, indicating that our method updates the metric toward better discrimination in a weighted manner. The second direction $\mathbf{M}\mathbf{X}_{jl} + \mathbf{X}_{jl}\mathbf{M}$ controls the scale of \mathbf{M} ; the metric will descend at a faster pace in the direction of a larger correlation between \mathbf{M} and \mathbf{X}_{jl} . This suggests our method functions as a data-dependent regularization. Let \mathbf{M}^t denote the Mahalanobis matrix learned at the t th iteration. The distance matrix will be updated as

$$\mathbf{M}^{t+1} = \mathbf{M}^t - \gamma \left(\frac{\partial J_{\text{LMNN}}}{\partial \mathbf{M}^t} + \lambda \frac{\partial J_{\text{P}}}{\partial \mathbf{M}^t} \right),$$

where γ denotes the learning rate. Following [14]'s work, γ is increased by 1% if the loss function decreases and decreased by 50% otherwise. To guarantee the PSD property, we factorize \mathbf{M}^{t+1} as $\mathbf{V}\mathbf{\Lambda}\mathbf{V}^T$ via eigendecomposition and truncate all negative eigenvalues to zero, i.e. $\mathbf{M}^{t+1} = \mathbf{V} \max(\mathbf{\Lambda}, 0) \mathbf{V}^T$.

Remark 2: The proposed perturbation loss is a generic approach to improving robustness against possible perturbation. In Appendix B, we illustrate examples of incorporating the perturbation loss into two different types of triplet-based methods, sparse compositional metric learning (SCML) [15] and proxy neighborhood component analysis (ProxyNCA++) [22]. SCML revises the structure of the Mahalanobis distance by representing it as a sparse and non-negative combination of rank-one basis elements, which typically results in less number of parameters to be estimated. ProxyNCA++ revises the construction of triplet constraints by replacing nearest instances \mathbf{x}_j and \mathbf{x}_l with nearest proxy points. The proxies are learned to represent each class, and the resulting method

is shown to generalize well on small datasets [23], robust to outliers and noisy labels [24], and improves computational efficiency on large-scale datasets.

Remark 3: Learning a distance metric for extremely high-dimensional data will result in a large number of parameters to be estimated and potentially suffer from overfitting. In order to reduce the input dimensionality, PCA is often applied to pre-process the data prior to metric learning [14], [25]. In Appendix A-A, we extend the proposed method such that the distance metric learned in the low-dimensional PCA subspace could still achieve robustness against perturbation in the original high-dimensional input space. The decision boundary of NN classifier (i.e. the constraint of Eq. 1) is revised in order to take account of the linear transformation matrix induced by the Mahalanobis distance and that of PCA. The proposed extension will be evaluated in Section IV-C.

D. Generalization benefit

From the perspective of algorithmic robustness [21], enlarging the adversarial margin could potentially improve the generalization ability of triplet-based metric learning methods. The following generalization bound, i.e. the gap between the generalization error \mathcal{L} and the empirical error ℓ_{emp} , follows from the pseudo-robust theorem of [26]. Preliminaries and derivations are given in Appendix C.

Theorem 1. Let \mathbf{M}^* be the optimal solution to Eq. 5. Then for any $\delta > 0$, with probability at least $1 - \delta$ we have:

$$\begin{aligned} |\mathcal{L}(\mathbf{M}^*) - \ell_{emp}(\mathbf{M}^*)| \\ \leq \frac{\hat{n}(t_s)}{n^3} + B \left(\frac{n^3 - \hat{n}(t_s)}{n^3} + 3\sqrt{\frac{2K \ln 2 + 2 \ln 1/\delta}{n}} \right), \end{aligned} \quad (7)$$

where $\hat{n}(t_s)$ denotes the number of triplets whose adversarial margins are larger than τ , B is a constant denoting the upper bound of the loss function (i.e. Eq. 5), and K denotes the number of disjoint sets that partition the input-label space and equals to $|\mathcal{Y}|(1 + \frac{\tau}{p})^p$.

Enlarging the desired adversarial margin τ will affect two quantities in Eq. 7, namely K and $\hat{n}(t_s)$. First, since K equals to $|\mathcal{Y}|(1 + \frac{\tau}{p})^p$, increasing τ will cause K to decrease at a polynomial rate of the input dimensionality p . Moreover, as the right-hand side of Eq. 7 is a function of K ($\mathcal{O}(K^{1/2})$), this means that the upper bound of generalization gap reduces at a rate of $p^{1/2}$. Hence, for datasets with a relative large number of features, a small improvement in the adversarial margin can greatly benefit the generalization ability of the learned metric.

Secondly, when τ increases, less triplets will satisfy the condition that their adversarial margin is larger than τ ; that is, $\hat{n}(t_s)$ decreases with τ . Meanwhile, since $B > 1$, the upper bound is a decreasing function of $\hat{n}(t_s)$. Therefore, enlarging τ leads to an increase in the upper bound. However, the rate of such increase depends on the datasets. For example, if most instances in the dataset are well separated and have a margin in the original input space, enlarging the desired adversarial margin τ will not have a large impact on $\hat{n}(t_s)$, the upper bound, and thus the generalization gap.

In summary, for datasets with many features and most instances being separable, we expect an improvement in the

generalization ability of the learned distance metric from enlarging the adversarial margin.

III. RELATED WORK

A. Robust metric learning

To make machine learning models more secure and trustworthy, robustness to input perturbations is a crucial dimension [27]. More importantly, designing such robust metric learning algorithms is particularly vital to safety-critical applications, such as healthcare [28], network intrusion detection [29], and surveillance systems based on faces [30], gaits [31] and other biometric traits [32].

Existing approaches to improving the robustness of Mahalanobis distances can be categorized into four main types. The first type of method imposes structural assumption or regularization over M so as to avoid overfitting [25], [33]–[37]. Methods with structural assumption are proposed for classifying images and achieve robustness by exploiting the structural information of images; however, such information is generally unavailable in the symbolic datasets that will be studied in this paper. Regularization-based methods are proposed to reduce the risk of overfitting to feature noise in the training set. Our proposal, which is aimed to withstand test-time perturbation, does not conflict with these methods and can be combined with them to learn a more effective and robust distance metric; an example is shown in Section IV-C.

The second type of method adopts loss functions that are less sensitive to outlier samples or noisy labels. In most metric learning methods, loss functions are founded on the squared L_2 -norm distance for computational efficiency. However, such choice may be sensitive to outliers. To overcome this limitation, several remedies have been proposed, such as using L_1 -norm distances [38] and metric based on the signal-to-noise ratio [39], or replacing the square function with the maximum correntropy criterion [40].

The third type of method studies robustness to training noise [41], [42]. These methods explicitly model the noise distribution or identify clean latent examples, and consequently, use the expected Mahalanobis distance to adjust the value of the distance margin for each triplet. Our method can also be viewed as imposing a data-dependent and dynamic margin – to achieve the same adversarial margin, triplets that have a higher correlation between $x_l - x_j$ and the metric M should satisfy a larger distance margin. However, the focus of our work is orthogonal to the aforementioned two types of method.

The last type of method generates hard instances through adversarial learning and trains a metric to fare well in the new hard problem [43], [44]. While sharing the aim of improving metric robustness, our method is intrinsically different from them. Their methods approach the task at a data-level, where real examples are synthesized based on the criterion of incurring large losses. Our method tackles perturbation at a model-level, where a loss function is derived by considering the definition of robustness with respect to the decision maker k NN. By preventing change in the nearest neighbor in a strict manner, our method is capable of obtaining a certification on adversarial margin.

B. Adversarial robustness of deep metric learning

More recently, deep metric learning has been investigated intensively, which replaces the linear projection induced by the Mahalanobis distance with deep neural networks. While deep neural networks improve the discriminability between classes, they are found to be non-robust and vulnerable to adversarial examples. Robust optimization [20], [45] is one of the most effective approaches to improving adversarial robustness, which trains the network to be robust against adversarial perturbations that are mostly constructed via gradient-based optimization; [46] adapts it to deep metric learning by considering the interdependence between data points in pairwise or triplet constraints. Another way to enhance robustness and generalization ability is by attaining a large margin in the input space, which dates back to support vector machines [47] and inspires this work. Due to the hierarchical nonlinear nature of deep networks, the input-space margin cannot be computed exactly and a variety of approximations have been proposed [48]–[51]. In this work, we investigate such margin in the framework of metric learning, defines it specifically with respect to the k NN classifier, and provide an exact and analytical solution to the margin. The analytical solution to the margin provides fascinating insights into essential factors for the robustness of distance metrics.

C. Adversarial robustness of k NN classifiers

While the notion of adversarial examples applies to k NN classifiers, existing methods for deep neural networks cannot be implemented directly due to the non-differential nature of the classifier. [52] and [53] propose continuous substitutes of k NN, from which gradient-based adversarial examples can be constructed to attack the classifier. [54] formulates a series of quadratic programming (QP) problems and proposes an efficient algorithm to search exhaustively over all training samples and compute the minimal adversarial perturbation for the 1-NN classifier. In addition, the dual solution to these QP problems can be used for robustness verification. [55] proposes to improve adversarial robustness for k NN by pruning the training set in order to satisfy the condition defined through the robustness radius, i.e., the norm of the minimal adversarial perturbation. Our work also aims to robustify k NN, but achieves it through enlarging the adversarial margin.

IV. EXPERIMENTS

In this section, we first present two toy examples to illustrate the difference in the learning mechanisms of LMNN and the proposed method dubbed LMNN-PL. Next, we compare LMNN-PL with state-of-the-art methods on 16 benchmark datasets (13 low/medium-dimensional and three high-dimensional), and investigate the relationship between adversarial margin, generalization ability and robustness. Finally, the computational aspect of our method is discussed.

A. Comparisons between LMNN and LMNN-PL

We design two experiments to compare the metrics learned with the objective of enhancing class discriminability and of




Fig. 2: Comparison of learning mechanisms of LMNN and LMNN-PL when features exhibit different separability.




Fig. 3: Comparison of learning mechanisms of LMNN and LMNN-PL when confronting the problem of multicollinearity.

certified robustness. In the first example, a two-dimensional binary classification dataset is simulated, as shown in Fig. 2a. The positive class includes 100 instances drawn uniformly from $[-3, 0]$ in the horizontal (abbr. 1st) direction and $[0, 1]$ in the vertical (abbr. 2nd) direction. The negative class consists of two clusters, where the first cluster includes 100 instances drawn from $U(-3, 0)$ and $U(-0.6, -0.5)$ in the 1st and 2nd directions respectively, and the second cluster includes 20 instances drawn from $U(0, 0.1)$ and $U(0, 1)$ in the two directions respectively. By design, instances of positive and negative classes can be separated in both directions, while the separability in the 1st direction is much smaller than the 2nd direction. Figs. 2b and 2c show the instances in the projected feature space with metrics learned from LMNN and LMNN-PL, respectively; the projection direction is indicated by the unit vector of red and blue lines; and the metric and the average of adversarial margins ($\bar{d}_{\mathcal{E}}(\mathbf{x}_i, \mathbf{x}_{i,\min})$) are given in the caption. The objective of LMNN is to satisfy the distance margin. Thus, it expands the distance in both directions. Moreover, since the 1st direction has a small separability in the original instance space, this direction is assigned with a larger weight. In contrast, LMNN-PL controls the scale of M . Moreover, a notable difference is that the 2nd direction is assigned with a larger weight than the 1st direction, which is again caused by the small separability in the 1st direction. As

any perturbation in the 1st direction is highly likely to result in a misclassification, the proposed method diverts more attention to robust features, i.e. the 2nd direction. Due to the easiness of the task, all metrics lead to the same classification accuracy of 99.09% on a separate test set.

In the second example, we simulate a three-dimensional binary classification dataset, as shown in Fig. 3a. Each class includes 100 instances. The first two dimensions are drawn from multivariate Gaussian distributions with $\boldsymbol{\mu}_p = [0.45, 0.45]$, $\boldsymbol{\mu}_n = [-0.45, -0.45]$, $\boldsymbol{\Sigma}_p = \boldsymbol{\Sigma}_n = \begin{bmatrix} 1 & -0.9 \\ -0.9 & 1 \end{bmatrix}$; the third dimension equals the sum of the first two dimensions, plus white Gaussian noise with standard deviation of 0.01. By design, the dataset exhibits the problem of strong multicollinearity. This issue has little influence on LMNN as the data is nearly separable in all directions. However, it will affect the adversarial margin. Specifically, if the metric assigns equal weights to all dimensions, then the perturbation should be small in all directions so as to guarantee that the perturbed instance stays on the correct side of the decision boundary. In contrast, if the metric assigns weights only to the third dimension, then the perturbation in the first two dimensions will not cause any change in the learned feature space and hence a larger magnitude of perturbation can be tolerated. This expectation is supported by the empirical result in Fig. 3c, where the distance in the third dimension is more important

than the first two dimensions. LMNN achieves an accuracy of 95.50% and our method achieves an accuracy of 96.00%.

In summary, our method learns a discriminative metric, and meanwhile, imposes a data-dependent regularization on the metric. It also achieves larger adversarial margins than LMNN, demonstrating the effectiveness of the proposed perturbation loss.

B. Experiments on UCI data

1) *Data description and experimental setting*: In this experiment, we study 13 datasets from UCI machine learning repository [56]. Information on sample size, feature dimension and class information is listed in Table V of Appendix D-A. All datasets are pre-processed with mean-centering and standardization, followed by L_2 normalization to unit length. To evaluate the performance, we use 70-30% training-test partitions and report the average result over 20 rounds of random split. The only exception is the Credit dataset, where we only run the experiment once as the sample size is relatively large.

We evaluate the effectiveness of the proposed perturbation loss by incorporating it into LMNN, SCML, and ProxyNCA++ (abbreviated to PNCA); the resulting methods are denoted by LMNN-PL, SCML-PL and PNCA-PL, respectively. In addition, we conduct a thorough study by setting LMNN as the backbone and comparing LMNN-PL with two types of methods. First, we consider different regularizers on M . Specifically, we replace the regularizer in LMNN from $\sum_S d_M^2(\mathbf{x}_i, \mathbf{x}_j)$ to the log-determinant divergence (LDD) [12], which encourages learning a metric toward the identity matrix, the capped trace norm (CAP) [25], which encourages a low-rank matrix, and the spectral norm (SN), which has been used to improve adversarial robustness of deep neural networks [57]. Second, we compare with the robust metric learning method DRIFT [41], which models the perturbation distribution explicitly.

Hyperparameters of LMNN-PL are tuned via random search [58]. We randomly sample 50 sets of values from the following ranges: $\mu \in U(0.1, 0.9)$, $\tau \in U(0, P_{90\%}\{d_E(\mathbf{x}_i, \mathbf{x}_{i,\min})\})$, $\lambda \in U(0, 4/\tau^2)$. $U(a, b)$ denotes the uniform distribution. $P_{k\%}\{d_E(\mathbf{x}_i, \mathbf{x}_{i,\min})\}$ denotes the k th percentile of $d_E(\mathbf{x}_i, \mathbf{x}_{i,\min})$, where the distance is calculated for all i in the triplet constraints with respect to the Euclidean distance. Setting the upper bound of the desired margin τ via the percentile avoids unnecessary large values, matching our intention to enlarge the adversarial margin primarily for hard instances. The upper bound of the weight parameter λ depends on the realization of τ to ensure that magnitudes of perturbation loss and LMNN loss are at the same level. The optimal hyperparameters from five-fold cross-validation on the training data or a separate validation set are used to learn the metric. SCML-PL and PNCA-PL are tuned in a similar manner. More details on the training procedure of the proposed and other methods are given in Appendix D-B. The MATLAB code for our method is available at <http://github.com/xyang6/LMNNPL>.

For LMNN-based and SCML-based methods, we use 3NN as the classifier; for PNCA-based methods, we use the nearest

prototype classifier. Classification accuracy is used as the evaluation criterion, except for two highly imbalanced datasets (Ecoli and Yeast), G-means is used.

2) *Evaluation on classification performance*: Table I reports the mean value and standard deviation of classification accuracy or G-means for imbalanced datasets (indicated by an asterisk). LMNN-PL outperforms LMNN on 12 out of 13 datasets. Among the methods with LMNN as the backbone, our method achieves the highest accuracy on 8 datasets and second highest accuracy on the 4 datasets. These experimental results demonstrate the benefit of perturbation loss to generalization of the learned metric. Similarly, we see that SCML-CL outperforms or performs equally well with SCML on 9 datasets. The advantage of PNCA-PL becomes less distinct as it is superior to PNCA only on 7 datasets. However, this is fairly reasonable as the decision boundary formed by very few proxies is much smoother than the one from 3NN and hence the method is less likely to overfit to training data.

3) *Investigation into robustness*: To test robustness, we add zero-mean Gaussian noise with a diagonal covariance matrix and equal variances to test data; the noise intensity is controlled via the signal-to-noise ratio (SNR) and chosen as 5 dB. In addition, considering the small sample size of UCI datasets, we augment test data by adding multiple rounds of random noise until its size reaches 10,000. As shown in Table II, the proposed methods achieve higher classification accuracy or G-means than the corresponding baselines on almost all datasets. Moreover, LMNN-PL is superior to existing regularization techniques or robust metric learning methods on at least 9 datasets. These results clearly demonstrate the efficacy of adding perturbation loss for improving robustness against instance perturbation. Additional experiments with other noise types and intensities are reported in Appendix D-C, where we observe similar advantages of the proposed loss.

C. Experiments on high-dimensional data

As mentioned in Remark 3, we extend LMNN-PL for high-dimensional data with PCA being used as a pre-processing step. To verify its effectiveness, we test it on the following three datasets:

- 1) Isolet [56]: The dataset is a spoken letter database and is available from UCI. It includes 7,797 instances, grouped into four training sets and one test set. We apply PCA to reduce the feature dimension from 617 to 170, accounting for 95% of total variance. All methods are trained four times, one time on each training set, and evaluated on the pre-given test set.
- 2) MNIST-2k [59]: The dataset includes the first 2,000 training images and first 2,000 test images of the MNIST database. PCA is applied to reduce the dimension from 784 to 141, retaining 95% of total variance. All methods are trained and tested once on the pre-given training/test partition.
- 3) APS Failure [56]: This is a multivariate dataset with a highly imbalanced class distribution. The training set includes 60,000 instances, among which 1,000 belong to the positive class. The test set includes 16,000 instances

TABLE I: Classification accuracy (or G-means indicated by an asterisk next to the dataset name; mean±standard deviation) of 3NN with different metric learning methods on clean datasets.

Dataset	Euclidean	LMNN-based					SCML-based		PNCA-based		
		LMNN	LDD	CAP	SN	DRIFT	LMNN-PL	SCML	SCML-PL	PNCA	PNCA-PL
Australian	82.76±2.38	83.70±2.43	84.18±2.37	83.97±2.45	83.77±2.50	84.47±2.02	84.47±1.63	84.76±2.08	84.42±2.18	86.30±1.82	86.15±2.24
Breast cancer	97.17±1.33	97.12±1.25	96.95±1.51	97.00±1.08	97.05±1.32	96.98±1.16	97.02±1.30	97.00±1.09	97.07±1.24	97.02±1.30	97.05±1.24
Ecoli*	85.86±8.10	86.42±7.94	85.29±9.78	83.54±10.09	86.44±7.95	<u>86.45±6.54</u>	87.04±7.38	85.53±7.06	86.69±6.58	84.80±6.84	85.54±5.60
Fourclass	75.12±2.35	75.10±2.31	75.15±2.32	75.02±2.48	75.10±2.31	75.08±2.34	75.12±2.35	75.10±2.27	75.12±2.35	75.39±2.21	72.97±5.36
Haberman	72.25±4.41	72.19±3.89	<u>72.42±3.95</u>	71.52±3.54	72.30±4.57	72.02±3.94	72.64±4.29	72.75±3.79	72.36±4.38	75.28±3.55	75.67±3.85
Iris	87.11±4.92	87.11±5.08	87.67±4.70	86.67±5.49	87.22±5.24	85.89±4.46	87.33±4.73	86.89±6.40	87.44±5.31	84.44±6.29	83.22±5.97
Segment	94.79±0.65	95.31±0.89	95.58±0.81	95.51±0.70	95.38±0.83	95.75±0.65	<u>95.64±0.83</u>	92.61±6.65	93.95±1.47	94.73±0.93	94.52±0.99
Sonar	85.16±4.19	86.67±4.10	<u>87.22±3.90</u>	87.22±4.38	86.67±4.04	86.19±4.43	87.78±3.53	82.38±4.15	84.13±4.61	83.25±5.95	83.65±4.83
Voting	93.78±1.76	95.80±1.78	95.80±1.41	<u>95.92±1.45</u>	95.84±1.74	95.31±1.32	96.15±1.56	95.84±1.58	96.26±1.28	95.84±1.65	95.65±1.66
WDBC	96.29±1.61	96.99±1.30	96.96±1.43	<u>96.99±1.51</u>	96.93±1.34	96.70±1.16	97.13±1.33	97.25±1.30	97.25±1.52	97.37±1.49	97.37±0.94
Wine	95.28±2.36	97.31±1.94	96.67±1.76	96.85±2.26	97.41±1.84	97.69±1.79	97.69±1.89	97.69±1.79	97.22±2.04	97.04±2.71	97.22±1.95
Yeast*	70.33±10.50	69.84±10.26	70.26±10.51	70.29±10.52	69.86±10.29	70.32±10.51	70.32±10.51	68.81±11.35	69.90±10.35	66.01±13.21	69.41±10.36
Credit	76.40	76.41	76.68	76.96	76.50	76.87	<u>76.89</u>	76.45	76.29	81.15	81.07
# outperform	-	12	11	12	12	12	-	9	-	7	-

For methods with LMNN as the backbone, the best ones are shown in bold and the second best ones are underlined; for methods with SCML or PNCA as the backbone, the best ones are shown in bold. '# outperform' counts the number of datasets where LMNN-PL (SCML-PL, PNCA-PL resp.) outperforms or performs equally well with LMNN-based (SCML, PNCA resp.) methods.

TABLE II: Classification accuracy (or G-means indicated by an asterisk) of 3NN noise-contaminated datasets. Gaussian noise with an SNR of 5 dB is added to test data.

Dataset	Euclidean	LMNN-based					SCML-based		PNCA-based		
		LMNN	LDD	CAP	SN	DRIFT	LMNN-PL	SCML	SCML-PL	PNCA	PNCA-PL
Australian	82.28±1.67	82.46±1.58	83.02±1.58	82.36±1.56	82.56±1.54	82.58±1.45	83.50±1.56	82.93±1.65	83.42±1.68	82.79±2.37	83.66±2.28
Breast cancer	96.79±1.05	96.25±1.09	<u>96.69±1.09</u>	96.35±1.02	96.29±1.11	96.66±1.00	96.71±1.08	96.40±1.05	96.65±1.06	96.20±1.37	96.77±1.14
Ecoli*	79.95±7.67	74.96±7.16	<u>79.13±7.67</u>	74.46±9.36	75.19±7.23	77.86±7.76	80.04±7.51	76.49±6.78	78.38±7.92	75.52±6.08	77.04±6.20
Fourclass	69.11±1.12	67.62±1.23	68.77±1.14	67.63±1.12	68.55±1.30	69.03±1.13	69.01±1.17	68.07±1.16	68.86±1.06	70.42±2.17	69.39±4.60
Haberman	69.93±1.88	69.84±1.79	69.92±1.87	69.23±2.00	69.90±1.88	69.09±2.49	69.89±1.90	69.65±1.63	69.88±1.83	74.32±3.22	74.32±3.10
Iris	79.75±3.26	78.61±2.97	<u>78.87±3.16</u>	77.79±3.27	78.70±3.08	78.43±3.09	79.04±3.09	78.16±3.58	79.01±3.12	77.95±4.53	78.20±3.90
Segment	88.18±0.64	81.02±3.55	86.15±1.26	85.34±2.47	82.10±3.41	86.63±1.09	84.72±2.62	60.18±9.73	61.33±9.05	78.27±2.83	80.28±3.35
Sonar	83.47±3.21	83.56±4.27	86.18±2.95	85.41±2.82	83.52±4.28	84.65±3.30	85.00±3.15	77.01±4.23	79.49±3.80	80.74±4.36	81.74±3.43
Voting	93.19±1.15	94.00±1.00	94.25±1.14	<u>94.37±1.17</u>	94.06±1.00	93.95±1.12	94.64±1.21	93.99±1.15	94.64±1.09	92.61±1.64	93.46±1.77
WDBC	95.92±1.30	91.71±1.90	96.30±0.94	96.16±1.08	92.46±1.80	96.04±0.86	96.11±0.88	95.74±1.30	96.21±1.16	96.03±1.54	96.22±1.15
Wine	94.20±1.46	93.33±1.63	94.03±1.39	<u>93.97±1.47</u>	93.45±1.70	94.66±1.15	<u>94.51±1.20</u>	94.01±1.56	94.61±1.32	94.19±1.94	93.48±1.48
Yeast*	69.36±10.47	54.13±8.24	68.62±10.43	66.48±10.18	55.49±9.45	<u>69.64±10.49</u>	69.82±10.44	55.96±7.64	60.47±10.39	61.41±17.97	63.59±18.33
Credit	76.28	76.16	76.22	76.05	<u>76.30</u>	76.37	76.15	75.93	75.55	78.24	79.13
# outperform	-	12	9	10	11	9	-	12	-	11	-

with 375 positive ones. The training set is further split into 40,000 instances for training and 20,000 instances for selecting hyperparameters. All methods are tested once on the test set. Applying PCA reduces the feature dimension from 161 to 79. Due to the large sample size, we only evaluate LMNN and LMNN-PL on this dataset.

In addition to aforementioned methods, we introduce CAP-PL, which comprises the triplet loss of LMNN, the regularizer of CAP, and the proposed perturbation loss. CAP enforces M to be low-rank, which is a suitable constraint for high-dimensional data. With the inclusion of perturbation loss, we expect the learned compact metric to be more robust to perturbation. For a fair comparison, in CAP-PL, we use the same rank and regularization weight as CAP, and tune τ, λ from 10 randomly sampled sets of values.

Table III compares the generalization and robustness performance of LMNN, CAP, SCML, and our method; the generalization performance of other methods are inferior to LMNN-PL and reported in Table VIII of Appendix D-C. First, on all three original datasets, our method achieves better performance than the baseline methods, validating its efficacy in improving the generalization ability of the learned metric. Secondly, when the SNR is 20 dB, the average perturbation

size is smaller than the average adversarial margin. In this case, our method maintains its superiority. When the SNR is 5 dB, the average perturbation size is larger than the average adversarial margin. Nonetheless, our method produces even larger gain in classification performance for LMNN on all datasets except APS Failure with the Gaussian noise, for SCML on MNIST and on Isolet with the Gaussian noise, for CAP on Isolet. These results suggest that adversarial margin is indeed a contributing factor in enhancing robustness. Thirdly, CAP-PL obtains higher accuracy on both clean and noise-contaminated data than LMNN-PL. This supports our discussion in Section III that regularization and perturbation loss impose different requirements on M and combining them has the potential for learning a more effective distance metric.

D. Computational cost

We now analyze the computational complexity of LMNN-PL. According to Eq. 6, our method requires additional calculations on $d_{M^2}^2(x_j, x_i)$ and MX_{jl} . Given n training instances, k target neighbors and p features, the computational complexities of $d_{M^2}^2(x_j, x_i)$ and MX_{jl} are $O(np^2 + n^2p)$ and $O(n^2p^2)$, respectively. The total complexity is $O(p^3 + n^2p^2 + kn^2p)$, same as that of LMNN.




Fig. 4: Sensitivity of LMNN-PL to hyperparameters (indicated by the straight line). The optimal accuracy and parameter value found via CV are indicated by the dashed line and asterisk, respectively.

TABLE III: Generalization and robustness of selected metric learning methods on high-dimensional datasets.

Isolet						
Method	Clean	IG,SNR=20 (0.0809)	IG,SNR=5 (0.4233)	AG,SNR=20 (0.0588)	AG,SNR=5 (0.3181)	Adv. margin
LMNN	90.14±4.45	90.09±4.15	86.02±3.48	90.17±4.03	87.81±3.87	0.1095
LMNN-PL	91.08±3.71	91.02±3.77	87.91±3.30	91.05±3.73	89.40±3.76	0.1249
SCML	90.73±4.10	90.33±4.21	86.50±4.18	90.51±4.14	88.50±3.71	0.0683
SCML-PL	90.83±4.16	90.67±4.12	86.55±3.75	90.83±4.16	88.41±4.07	0.0822
CAP	91.05±3.66	91.13±3.85	88.97±4.00	91.10±3.73	89.90±3.87	0.1514
CAP-PL	91.58±3.96	91.52±3.86	89.91±3.74	91.47±3.91	90.65±3.73	0.1559
MNIST						
Method	Clean	IG,SNR=20 (0.0540)	IG,SNR=5 (0.2939)	AG,SNR=20 (0.0649)	AG,SNR=5 (0.3482)	Adv. margin
LMNN	90.55	90.00	88.40	90.10	88.40	0.1528
LMNN-PL	91.15	91.35	90.80	91.45	90.35	0.2235
SCML	88.95	88.75	87.35	88.85	86.45	0.1217
SCML-PL	89.15	89.20	88.50	89.35	88.05	0.1432
CAP	91.65	91.80	91.40	91.80	90.70	0.2219
CAP-PL	92.00	91.90	90.85	91.95	90.65	0.2264
APS Failure						
Method	Clean	IG,SNR=20 (0.0906)	IG,SNR=5 (0.5097)	AG,SNR=20 (0.0996)	AG,SNR=5 (0.5604)	Adv. margin
LMNN	80.69	80.20	75.66	81.18	75.13	0.1773
LMNN-PL	80.89	82.15	74.13	82.33	77.92	0.2583

Columns 3-6 report methods' robustness against isotropic Gaussian noise (IG) and anisotropic Gaussian noise (AG). Values in brackets give the average perturbation size, calculated as the mean value of the L_2 -norm of noises ($\|\Delta \mathbf{x}_i\|_2$).

TABLE IV: Average training time (in seconds) of LMNN-based methods.

	LMNN	LDD	CAP	SN	DRIFT	LMNN-PL
Australian	13.44	0.83	3.07	7.60	1.00	2.15
Segment	27.48	10.45	11.47	24.66	5.12	19.54
Sonar	4.93	4.08	4.65	30.39	0.92	6.75
WDBC	9.38	2.94	5.22	16.54	5.12	8.17
Credit	724.42	34.65	115.22	966.63	130.36	138.63
Isolet	339.57	207.69	176.50	540.26	NA	190.55
MNIST	369.55	68.98	289.18	197.50	37.51	391.04

Table IV compares the running time of LMNN-based methods on five UCI datasets that are large in sample size or in dimensionality and two high-dimensional datasets. The computational cost of our method is comparable to LMNN.

E. Parameter sensitivity

The proposed LMNN-PL includes three hyperparameters – μ for the weight of similarity constraints, λ for the weight of the perturbation loss, and τ for the desired adversarial margin. We investigate their influences on the classification performance by varying one hyperparameter and fixing the other two at their optimal values. Fig. 4 shows the accuracy on MNIST evaluated over the range of the hyperparameter. The performance changes smoothly with respect to μ . It is stable over a wide range of λ . When λ equals 0, LMNN-PL fails to learn a metric and returns a zero matrix. The performance is most affected by τ . Indeed, τ plays the central role in LMNN-PL as it determines the distribution of adversarial margins. A small value of τ has little influence on the objective function as the adversarial margin of most instances may already exceed it before optimization, and a large value may greatly reduce the number of triplets that satisfy the loss condition in the definition of pseudo-robustness (i.e. $\hat{n}(t_s)$ in Theorem 1). Therefore, we shall strive to search for its optimal value.

V. CONCLUSION

In this paper, we propose to enhance the robustness and generalization of distance metrics. This is easily achievable by taking advantage of the linear transformation induced by the Mahalanobis distance. Specifically, we find an explicit formula for the adversarial margin, which is defined as the Euclidean distance between benign instances and their closest adversarial examples, and advocate to enlarge it through penalizing the perturbation loss designed on the basis of the derivation. Experiments verify that our method effectively enlarges the adversarial margin, sustains classification excellence, and enhances robustness to instance perturbation. The proposed loss term is generic in nature and could be readily embedded in other Mahalanobis-based metric learning methods. In the future, we will consider extending the idea to metric learning methods with nonlinear feature extraction and/or nonlinear metric learning methods.

REFERENCES

- [1] E. Xing, M. Jordan, S. J. Russell, and A. Ng, "Distance metric learning with application to clustering with side-information," *Advances in Neural Information Processing Systems*, vol. 15, 2002.
- [2] B. Kulis, "Metric learning: A survey," *Foundations and Trends in Machine Learning*, vol. 5, no. 4, pp. 287–364, 2013.
- [3] A. Bellet, A. Habrard, and M. Sebban, "Metric learning," *Synthesis Lectures on Artificial Intelligence and Machine Learning*, vol. 9, no. 1, pp. 1–151, 2015.

- [4] D. Li and Y. Tian, "Survey and experimental study on metric learning methods," *Neural Networks*, vol. 105, pp. 447–462, 2018.
- [5] S. Xiang, F. Nie, and C. Zhang, "Learning a Mahalanobis distance metric for data clustering and classification," *Pattern Recognition*, vol. 41, no. 12, pp. 3600–3612, 2008.
- [6] B. McFee and G. R. Lanckriet, "Metric learning to rank," in *International Conference on Machine Learning*, 2010.
- [7] J. Yang, D. She, Y.-K. Lai, and M.-H. Yang, "Retrieving and classifying affective images via deep metric learning," in *AAAI Conference on Artificial Intelligence*, 2018.
- [8] M. Koestinger, M. Hirzer, P. Wohlhart, P. M. Roth, and H. Bischof, "Large scale metric learning from equivalence constraints," in *IEEE Conference on Computer Vision and Pattern Recognition*, 2012.
- [9] J. Hu, J. Lu, Y.-P. Tan, J. Yuan, and J. Zhou, "Local large-margin multi-metric learning for face and kinship verification," *IEEE Transactions on Circuits and Systems for Video Technology*, vol. 28, no. 8, pp. 1875–1891, 2017.
- [10] L. Ma, X. Yang, and D. Tao, "Person re-identification over camera networks using multi-task distance metric learning," *IEEE Transactions on Image Processing*, vol. 23, no. 8, pp. 3656–3670, 2014.
- [11] G. Zou, G. Fu, X. Peng, Y. Liu, M. Gao, and Z. Liu, "Person re-identification based on metric learning: a survey," *Multimedia Tools and Applications*, vol. 80, no. 17, pp. 26 855–26 888, 2021.
- [12] J. V. Davis, B. Kulis, P. Jain, S. Sra, and I. S. Dhillon, "Information-theoretic metric learning," in *International Conference on Machine Learning*. ACM, 2007.
- [13] J. T. Kwok and I. W. Tsang, "Learning with idealized kernels," in *International Conference on Machine Learning*, 2003.
- [14] K. Q. Weinberger and L. K. Saul, "Distance metric learning for large margin nearest neighbor classification," *Journal of Machine Learning Research*, vol. 10, no. Feb, pp. 207–244, 2009.
- [15] Y. Shi, A. Bellet, and F. Sha, "Sparse compositional metric learning," in *AAAI Conference on Artificial Intelligence*, 2014.
- [16] K. Song, F. Nie, J. Han, and X. Li, "Parameter free large margin nearest neighbor for distance metric learning," in *AAAI Conference on Artificial Intelligence*, 2017.
- [17] M. Dong, Y. Wang, X. Yang, and J.-H. Xue, "Learning local metrics and influential regions for classification," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 42, no. 6, pp. 1522–1529, 2020.
- [18] M. T. Law, N. Thome, and M. Cord, "Quadruplet-wise image similarity learning," in *IEEE International Conference on Computer Vision*, 2013.
- [19] K. Sohn, "Improved deep metric learning with multi-class n-pair loss objective," in *Advances in Neural Information Processing Systems*, 2016.
- [20] C. Szegedy, W. Zaremba, I. Sutskever, J. Bruna, D. Erhan, I. Goodfellow, and R. Fergus, "Intriguing properties of neural networks," *arXiv preprint arXiv:1312.6199*, 2013.
- [21] H. Xu and S. Mannor, "Robustness and generalization," *Machine Learning*, vol. 86, no. 3, pp. 391–423, 2012.
- [22] E. W. Teh, T. DeVries, and G. W. Taylor, "ProxyNCA++: Revisiting and revitalizing proxy neighborhood component analysis," in *European Conference on Computer Vision*. Springer, 2020, pp. 448–464.
- [23] J. Snell, K. Swersky, and R. Zemel, "Prototypical networks for few-shot learning," in *Advances in Neural Information Processing Systems*, 2017, pp. 4080–4090.
- [24] S. Kim, D. Kim, M. Cho, and S. Kwak, "Proxy anchor loss for deep metric learning," in *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020, pp. 3238–3247.
- [25] Z. Huo, F. Nie, and H. Huang, "Robust and effective metric learning using capped trace norm: Metric learning via capped trace norm," in *ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 2016.
- [26] A. Bellet and A. Habrard, "Robustness and generalization for metric learning," *Neurocomputing*, vol. 151, pp. 259–267, 2015.
- [27] L. Floridi, "Establishing the rules for building trustworthy AI," *Nature Machine Intelligence*, vol. 1, no. 6, pp. 261–262, 2019.
- [28] Q. Suo, W. Zhong, F. Ma, Y. Ye, M. Huai, and A. Zhang, "Multi-task sparse metric learning for monitoring patient similarity progression," in *IEEE International Conference on Data Mining*, 2018, pp. 477–486.
- [29] R. Aliakbarisani, A. Ghasemi, and S. F. Wu, "A data-driven metric learning-based scheme for unsupervised network anomaly detection," *Computers & Electrical Engineering*, vol. 73, pp. 71–83, 2019.
- [30] M. Guillaumin, J. Verbeek, and C. Schmid, "Is that you? metric learning approaches for face identification," in *IEEE International Conference on Computer Vision*, 2009, pp. 498–505.
- [31] X. Ben, P. Zhang, W. Meng, R. Yan, M. Yang, W. Liu, and H. Zhang, "On the distance metric learning between cross-domain gaits," *Neurocomputing*, vol. 208, pp. 153–164, 2016.
- [32] I. Omara, H. Zhang, F. Wang, A. Hagag, X. Li, and W. Zuo, "Metric learning with dynamically generated pairwise constraints for ear recognition," *Information*, vol. 9, no. 9, p. 215, 2018.
- [33] R. Jin, S. Wang, and Y. Zhou, "Regularized distance metric learning: Theory and algorithm," in *Advances in Neural Information Processing Systems*, 2009.
- [34] D. Lim, G. Lanckriet, and B. McFee, "Robust structural metric learning," in *International Conference on Machine Learning*, 2013.
- [35] M. T. Law, N. Thome, and M. Cord, "Fantope regularization in metric learning," in *IEEE Conference on Computer Vision and Pattern Recognition*, 2014.
- [36] L. Luo and H. Huang, "Matrix variate Gaussian mixture distribution steered robust metric learning," in *AAAI Conference on Artificial Intelligence*, 2018.
- [37] K. Liu, L. Brand, H. Wang, and F. Nie, "Learning robust distance metric with side information via ratio minimization of orthogonally constrained $\ell_{2,1}$ -norm distances," in *International Joint Conference on Artificial Intelligence*, 2019.
- [38] H. Wang, F. Nie, and H. Huang, "Robust distance metric learning via simultaneous ℓ_1 -norm minimization and maximization," in *International Conference on Machine Learning*, 2014.
- [39] T. Yuan, W. Deng, J. Tang, Y. Tang, and B. Chen, "Signal-to-noise ratio: A robust distance metric for deep metric learning," in *IEEE Conference on Computer Vision and Pattern Recognition*, 2019.
- [40] J. Xu, L. Luo, C. Deng, and H. Huang, "New robust metric learning model using maximum correntropy criterion," in *ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, 2018, pp. 2555–2564.
- [41] H.-J. Ye, D.-C. Zhan, X.-M. Si, and Y. Jiang, "Learning Mahalanobis distance metric: considering instance disturbance helps," in *International Joint Conference on Artificial Intelligence*, 2017.
- [42] Q. Qian, J. Tang, H. Li, S. Zhu, and R. Jin, "Large-scale distance metric learning with uncertainty," in *IEEE Conference on Computer Vision and Pattern Recognition*, 2018.
- [43] S. Chen, C. Gong, J. Yang, X. Li, Y. Wei, and J. Li, "Adversarial metric learning," in *International Joint Conference on Artificial Intelligence*, 2018.
- [44] Y. Duan, W. Zheng, X. Lin, J. Lu, and J. Zhou, "Deep adversarial metric learning," in *IEEE Conference on Computer Vision and Pattern Recognition*, 2018.
- [45] A. Ben-Tal, L. El Ghaoui, and A. Nemirovski, *Robust Optimization*. Princeton University Press, 2009.
- [46] T. K. Panum, Z. Wang, P. Kan, E. Fernandes, and S. Jha, "Adversarial deep metric learning," 2021. [Online]. Available: <https://openreview.net/forum?id=Kzg0XmE6mxu>
- [47] C. Cortes and V. Vapnik, "Support-vector networks," *Machine Learning*, vol. 20, no. 3, pp. 273–297, 1995.
- [48] S. An, M. Hayat, S. H. Khan, M. Bennamoun, F. Boussaid, and F. Sohel, "Contractive rectifier networks for nonlinear maximum margin classification," in *IEEE International Conference on Computer Vision*, 2015.
- [49] G. Elsayed, D. Krishnan, H. Mobahi, K. Regan, and S. Bengio, "Large margin deep networks for classification," in *Advances in Neural Information Processing Systems*, 2018.
- [50] Z. Yan, Y. Guo, and C. Zhang, "Adversarial margin maximization networks," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2019.
- [51] G. W. Ding, Y. Sharma, K. Y. C. Lui, and R. Huang, "MMA training: Direct input space margin maximization through adversarial training," in *International Conference on Learning Representations*, 2020.
- [52] N. Papernot, P. McDaniel, and I. Goodfellow, "Transferability in machine learning: from phenomena to black-box attacks using adversarial samples," *arXiv preprint arXiv:1605.07277*, 2016.
- [53] C. Sitawarin and D. Wagner, "On the robustness of deep k-nearest neighbors," in *IEEE Security and Privacy Workshops*. IEEE, 2019, pp. 1–7.
- [54] L. Wang, X. Liu, J. Yi, Z.-H. Zhou, and C.-J. Hsieh, "Evaluating the robustness of nearest neighbor classifiers: A primal-dual perspective," *arXiv preprint arXiv:1906.03972*, 2019.
- [55] Y.-Y. Yang, C. Rashtchian, Y. Wang, and K. Chaudhuri, "Robustness for non-parametric classification: A generic attack and defense," in *International Conference on Artificial Intelligence and Statistics*. PMLR, 2020, pp. 941–951.
- [56] D. Dua and C. Graff, "UCI machine learning repository," 2017. [Online]. Available: <http://archive.ics.uci.edu/ml>

- [57] P. L. Bartlett, D. J. Foster, and M. J. Telgarsky, "Spectrally-normalized margin bounds for neural networks," *Advances in Neural Information Processing Systems*, vol. 30, pp. 6240–6249, 2017.
- [58] J. Bergstra and Y. Bengio, "Random search for hyper-parameter optimization," *Journal of Machine Learning Research*, vol. 13, no. Feb, pp. 281–305, 2012.
- [59] D. Cai, X. He, J. Han, and T. S. Huang, "Graph regularized nonnegative matrix factorization for data representation," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 33, no. 8, pp. 1548–1560, 2010, data: <http://www.cad.zju.edu.cn/home/dengcai/Data/MLData.html>.
- [60] K. B. Petersen and M. S. Pedersen, "The matrix cookbook, nov 2012," vol. 3274, 2012, uRL: <http://www2.imm.dtu.dk/pubdb/p.php>.
- [61] R. Tibshirani, "CMU Machine Learning 10-725 Lecture Slides: Proximal Gradient Descent and Acceleration," 2015, uRL: <https://www.stat.cmu.edu/~ryantibs/convexopt-S15/lectures/08-prox-grad.pdf>. Last visited on 18/May/2020.
- [62] J. Goldberger, G. E. Hinton, S. Roweis, and R. R. Salakhutdinov, "Neighbourhood components analysis," *Advances in neural information processing systems*, vol. 17, 2004.
- [63] M. J. Wainwright, *Metric entropy and its uses*, ser. Cambridge Series in Statistical and Probabilistic Mathematics. Cambridge University Press, 2019, pp. 121–158.

APPENDIX A

DERIVATION OF THE CLOSEST ADVERSARIAL EXAMPLE, ADVERSARIAL MARGIN AND GRADIENT OF PERTURBATION LOSS

In the following derivations, we consider heterogeneous and correlated perturbations; the homogeneous and uncorrelated perturbations, discussed in the main text, can be seen as a special case.

First, we characterize heterogeneous and correlated perturbations through hyperellipsoids. In terms of the quadratic form, an arbitrarily oriented hyperellipsoid, centered at $\boldsymbol{\mu} \in \mathbb{R}^p$, is defined by the solutions to the equation

$$\{\boldsymbol{x} \in \mathbb{R}^p : (\boldsymbol{x} - \boldsymbol{\mu})^T \mathbf{A}_0 (\boldsymbol{x} - \boldsymbol{\mu}) = r^2\},$$

where \mathbf{A}_0 is a positive definite matrix. By the Cholesky decomposition, $\mathbf{A}_0 = \mathbf{A}\mathbf{A}^T$. Therefore, finding the closest adversarial example of \boldsymbol{x}_i on the hyperellipsoid is equivalent to finding the point \boldsymbol{x}'_i that defines the smallest hypersphere given by $(\mathbf{A}^T(\boldsymbol{x}'_i - \boldsymbol{x}_i))^T (\mathbf{A}^T(\boldsymbol{x}'_i - \boldsymbol{x}_i)) = r^2$.

The closest adversarial example is identified by optimizing Eq. 8; its equivalent form is given in Eq. 9:

$$\begin{aligned} \boldsymbol{x}_{i,\min} &= \arg \min_{\boldsymbol{x}'_i \in \mathbb{R}^p} (\boldsymbol{x}'_i - \boldsymbol{x}_i)^T \mathbf{A}_0 (\boldsymbol{x}'_i - \boldsymbol{x}_i) \\ \text{s.t. } (\mathbf{L}\boldsymbol{x}'_i - \frac{\mathbf{L}\boldsymbol{x}_j + \mathbf{L}\boldsymbol{x}_l}{2})^T (\mathbf{L}\boldsymbol{x}_l - \mathbf{L}\boldsymbol{x}_j) &= 0. \end{aligned} \quad (8)$$

$$\begin{aligned} \boldsymbol{x}_{i,\min} &= \arg \min_{\boldsymbol{x}'_i \in \mathbb{R}^p} (\mathbf{A}^T(\boldsymbol{x}'_i - \boldsymbol{x}_i))^T (\mathbf{A}^T(\boldsymbol{x}'_i - \boldsymbol{x}_i)) \\ \text{s.t. } (\mathbf{L}\boldsymbol{x}'_i - \frac{\mathbf{L}\boldsymbol{x}_j + \mathbf{L}\boldsymbol{x}_l}{2})^T (\mathbf{L}\boldsymbol{x}_l - \mathbf{L}\boldsymbol{x}_j) &= 0. \end{aligned} \quad (9)$$

Applying the method of Lagrangian multiplier, we transform the above problem to the following Lagrangian function by introducing the Lagrangian multiplier λ and then solve it by setting the first partial derivatives to zero:

$$\min_{\boldsymbol{x}'_i} (\mathbf{A}^T(\boldsymbol{x}'_i - \boldsymbol{x}_i))^T (\mathbf{A}^T(\boldsymbol{x}'_i - \boldsymbol{x}_i)) - \lambda (\mathbf{L}\boldsymbol{x}'_i - \frac{\mathbf{L}\boldsymbol{x}_j + \mathbf{L}\boldsymbol{x}_l}{2})^T (\mathbf{L}\boldsymbol{x}_l - \mathbf{L}\boldsymbol{x}_j)$$

$$\begin{aligned} \frac{\delta}{\delta \boldsymbol{x}'_i} : 2\mathbf{A}\mathbf{A}^T(\boldsymbol{x}'_i - \boldsymbol{x}_i) - \lambda \mathbf{L}^T \mathbf{L}(\boldsymbol{x}_l - \boldsymbol{x}_j) &= 0 \\ \boldsymbol{x}'_i &= \boldsymbol{x}_i + \frac{\lambda}{2} \mathbf{A}_0^{-1} \mathbf{L}^T \mathbf{L}(\boldsymbol{x}_l - \boldsymbol{x}_j) \\ (\mathbf{L}\boldsymbol{x}_i + \frac{\lambda}{2} \mathbf{L}\mathbf{A}_0^{-1} \mathbf{L}^T \mathbf{L}(\boldsymbol{x}_l - \boldsymbol{x}_j) - \frac{\mathbf{L}\boldsymbol{x}_j + \mathbf{L}\boldsymbol{x}_l}{2})^T (\mathbf{L}\boldsymbol{x}_l - \mathbf{L}\boldsymbol{x}_j) &= 0 \\ \frac{\lambda}{2} &= \frac{(\frac{\boldsymbol{x}_j + \boldsymbol{x}_l}{2} - \boldsymbol{x}_i)^T \mathbf{L}^T \mathbf{L}(\boldsymbol{x}_l - \boldsymbol{x}_j)}{(\boldsymbol{x}_l - \boldsymbol{x}_j)^T \mathbf{L}^T \mathbf{L}\mathbf{A}_0^{-1} \mathbf{L}^T \mathbf{L}(\boldsymbol{x}_l - \boldsymbol{x}_j)} \\ \boldsymbol{x}_{i,\min} &= \boldsymbol{x}_i + \frac{(\frac{\boldsymbol{x}_j + \boldsymbol{x}_l}{2} - \boldsymbol{x}_i)^T \mathbf{L}^T \mathbf{L}(\boldsymbol{x}_l - \boldsymbol{x}_j)}{(\boldsymbol{x}_l - \boldsymbol{x}_j)^T \mathbf{L}^T \mathbf{L}\mathbf{A}_0^{-1} \mathbf{L}^T \mathbf{L}(\boldsymbol{x}_l - \boldsymbol{x}_j)} \mathbf{A}_0^{-1} \mathbf{M}(\boldsymbol{x}_l - \boldsymbol{x}_j) \end{aligned}$$

The Hessian matrix equals $2\mathbf{A}_0$, which is positive definite, and hence $\boldsymbol{x}_{i,\min}$ is the minimum point. Replacing $\mathbf{A}_0 = \mathbf{I}$ (identity matrix) gives Eq. 2.

Next, we calculate the squared adversarial margin by first simplifying $\boldsymbol{x}_{i,\min}$ and then computing r^2 as follows:

$$\begin{aligned} & \left(\frac{\boldsymbol{x}_j + \boldsymbol{x}_l}{2} - \boldsymbol{x}_i \right)^T \mathbf{M}(\boldsymbol{x}_l - \boldsymbol{x}_j) \\ &= \frac{1}{2} [(\boldsymbol{x}_j - \boldsymbol{x}_i) + (\boldsymbol{x}_l - \boldsymbol{x}_i)]^T \mathbf{M}[(\boldsymbol{x}_l - \boldsymbol{x}_i) - (\boldsymbol{x}_j - \boldsymbol{x}_i)] \\ &= \frac{1}{2} [d_M^2(\boldsymbol{x}_i, \boldsymbol{x}_l) - d_M^2(\boldsymbol{x}_i, \boldsymbol{x}_j)] \\ r^2 &= (\boldsymbol{x}_i - \boldsymbol{x}_{i,\min})^T \mathbf{A}_0 (\boldsymbol{x}_i - \boldsymbol{x}_{i,\min}) \\ &= \left(\frac{d_M^2(\boldsymbol{x}_i, \boldsymbol{x}_l) - d_M^2(\boldsymbol{x}_i, \boldsymbol{x}_j)}{2(\boldsymbol{x}_l - \boldsymbol{x}_j)^T \mathbf{M}\mathbf{A}_0^{-1} \mathbf{M}(\boldsymbol{x}_l - \boldsymbol{x}_j)} \mathbf{A}_0^{-1} \mathbf{L}^T \mathbf{L}(\boldsymbol{x}_l - \boldsymbol{x}_j) \right)^T \\ & \quad \mathbf{A}_0 \left(\frac{d_M^2(\boldsymbol{x}_i, \boldsymbol{x}_l) - d_M^2(\boldsymbol{x}_i, \boldsymbol{x}_j)}{2(\boldsymbol{x}_l - \boldsymbol{x}_j)^T \mathbf{M}\mathbf{A}_0^{-1} \mathbf{M}(\boldsymbol{x}_l - \boldsymbol{x}_j)} \mathbf{A}_0^{-1} \mathbf{L}^T \mathbf{L}(\boldsymbol{x}_l - \boldsymbol{x}_j) \right) \\ &= \frac{(d_M^2(\boldsymbol{x}_i, \boldsymbol{x}_l) - d_M^2(\boldsymbol{x}_i, \boldsymbol{x}_j))^2}{4 \left((\boldsymbol{x}_l - \boldsymbol{x}_j)^T \mathbf{M}\mathbf{A}_0^{-1} \mathbf{M}(\boldsymbol{x}_l - \boldsymbol{x}_j) \right)^2} \\ & \quad (\boldsymbol{x}_l - \boldsymbol{x}_j)^T \mathbf{L}^T \mathbf{L}\mathbf{A}_0^{-1} \mathbf{A}_0 \mathbf{A}_0^{-1} \mathbf{L}^T \mathbf{L}(\boldsymbol{x}_l - \boldsymbol{x}_j) \\ &= \frac{(d_M^2(\boldsymbol{x}_i, \boldsymbol{x}_l) - d_M^2(\boldsymbol{x}_i, \boldsymbol{x}_j))^2}{4 \left((\boldsymbol{x}_l - \boldsymbol{x}_j)^T \mathbf{M}\mathbf{A}_0^{-1} \mathbf{M}(\boldsymbol{x}_l - \boldsymbol{x}_j) \right)} \end{aligned}$$

Substituting $\mathbf{A}_0 = \mathbf{I}$ gives Eq. 3.

Finally, we derive the gradient of J_P with respect to \mathbf{M} . When $d_M^2(\boldsymbol{x}_i, \boldsymbol{x}_l) > d_M^2(\boldsymbol{x}_i, \boldsymbol{x}_j)$ and $r \geq \tau$ (i.e. $\tilde{d}_E(\boldsymbol{x}_i, \boldsymbol{x}_{i,\min}) \geq \tau$ in the hypersphere case), or $d_M^2(\boldsymbol{x}_i, \boldsymbol{x}_l) \leq d_M^2(\boldsymbol{x}_i, \boldsymbol{x}_j)$, the gradient equals zero. When $d_M^2(\boldsymbol{x}_i, \boldsymbol{x}_l) > d_M^2(\boldsymbol{x}_i, \boldsymbol{x}_j)$ and $r < \tau$, the gradient of J_P equals the gradient of $-r^2$ (i.e. $-\tilde{d}_E^2(\boldsymbol{x}_i, \boldsymbol{x}_{i,\min})$ in the hypersphere case), which can be calculated by using the quotient rule and the derivative of trace [60]:

$$\begin{aligned} & \frac{\partial}{\partial \mathbf{M}} (d_M^2(\boldsymbol{x}_i, \boldsymbol{x}_l) - d_M^2(\boldsymbol{x}_i, \boldsymbol{x}_j))^2 \\ &= 2(d_M^2(\boldsymbol{x}_i, \boldsymbol{x}_l) - d_M^2(\boldsymbol{x}_i, \boldsymbol{x}_j)) (\mathbf{X}_{il} - \mathbf{X}_{ij}) \\ \frac{\partial}{\partial \mathbf{M}} (\boldsymbol{x}_l - \boldsymbol{x}_j)^T \mathbf{M}\mathbf{A}_0^{-1} \mathbf{M}(\boldsymbol{x}_l - \boldsymbol{x}_j) &= \frac{\partial}{\partial \mathbf{M}} \text{tr}(\mathbf{X}_{jl} \mathbf{M}\mathbf{A}_0^{-1} \mathbf{M}) \\ &= \mathbf{X}_{jl} \mathbf{M}\mathbf{A}_0^{-1} + \mathbf{A}_0^{-1} \mathbf{M}\mathbf{X}_{jl} \\ \frac{\partial}{\partial \mathbf{M}} \frac{(d_M^2(\boldsymbol{x}_i, \boldsymbol{x}_l) - d_M^2(\boldsymbol{x}_i, \boldsymbol{x}_j))^2}{4 \left((\boldsymbol{x}_l - \boldsymbol{x}_j)^T \mathbf{M}\mathbf{A}_0^{-1} \mathbf{M}(\boldsymbol{x}_l - \boldsymbol{x}_j) + \epsilon \right)} & \\ &= \frac{2(d_M^2(\boldsymbol{x}_i, \boldsymbol{x}_l) - d_M^2(\boldsymbol{x}_i, \boldsymbol{x}_j)) (\mathbf{X}_{il} - \mathbf{X}_{ij})}{4 \left((\boldsymbol{x}_l - \boldsymbol{x}_j)^T \mathbf{M}\mathbf{A}_0^{-1} \mathbf{M}(\boldsymbol{x}_l - \boldsymbol{x}_j) + \epsilon \right)} \\ & \quad - \frac{(d_M^2(\boldsymbol{x}_i, \boldsymbol{x}_l) - d_M^2(\boldsymbol{x}_i, \boldsymbol{x}_j))^2 (\mathbf{X}_{jl} \mathbf{M}\mathbf{A}_0^{-1} + \mathbf{A}_0^{-1} \mathbf{M}\mathbf{X}_{jl})}{4 \left((\boldsymbol{x}_l - \boldsymbol{x}_j)^T \mathbf{M}\mathbf{A}_0^{-1} \mathbf{M}(\boldsymbol{x}_l - \boldsymbol{x}_j) + \epsilon \right)^2}, \end{aligned}$$

where $\text{tr}(\cdot)$ denotes the trace operator. $\mathbf{X}_{ij} = (\boldsymbol{x}_i - \boldsymbol{x}_j)(\boldsymbol{x}_i - \boldsymbol{x}_j)^T$ and $\mathbf{X}_{il}, \mathbf{X}_{jl}$ are defined similarly. Substituting $\mathbf{A}_0 = \mathbf{I}$ gives Eq. 6.

A. Derivation of closest adversarial example, adversarial margin, and gradient of perturbation loss in the high-dimensional case

The closet adversarial example, adversarial margin and gradient of the perturbation loss with dimensionality reduction are derived by following the same principle as in Appendix A.

The method of Lagrangian multiplier is applied to derive a closed-form solution to the closest adversarial example:

$$\begin{aligned} & \min_{\mathbf{x}'_i} (\mathbf{A}^T(\mathbf{x}'_i - \mathbf{x}_i))^T (\mathbf{A}^T(\mathbf{x}'_i - \mathbf{x}_i)) \\ & \quad - \lambda (\mathbf{x}'_i - \frac{\mathbf{x}_j + \mathbf{x}_l}{2})^T \mathbf{D}^T \mathbf{L}^T \mathbf{L} \mathbf{D} (\mathbf{x}_l - \mathbf{x}_j) \\ \frac{\delta}{\delta \mathbf{x}'_i} : \mathbf{x}'_i = & \mathbf{x}_i + \frac{\lambda}{2} \mathbf{A}_0^{-1} \mathbf{D}^T \mathbf{L}^T \mathbf{L} \mathbf{D} (\mathbf{x}_l - \mathbf{x}_j) \\ \frac{\lambda}{2} = & \frac{(\frac{\mathbf{x}_j + \mathbf{x}_l}{2} - \mathbf{x}_i)^T \mathbf{D}^T \mathbf{L}^T \mathbf{L} \mathbf{D} (\mathbf{x}_l - \mathbf{x}_j)}{(\mathbf{x}_l - \mathbf{x}_j)^T \mathbf{D}^T \mathbf{L}^T \mathbf{L} \mathbf{D} \mathbf{A}_0^{-1} \mathbf{D}^T \mathbf{L}^T \mathbf{L} \mathbf{D} (\mathbf{x}_l - \mathbf{x}_j)} \\ \mathbf{x}_{i,\min}^{\text{PCA}} = & \mathbf{x}_i + \frac{(\frac{\tilde{\mathbf{x}}_l + \tilde{\mathbf{x}}_l}{2} - \tilde{\mathbf{x}}_i)^T \mathbf{M} (\tilde{\mathbf{x}}_l - \tilde{\mathbf{x}}_j)}{(\tilde{\mathbf{x}}_l - \tilde{\mathbf{x}}_j)^T \mathbf{M} \mathbf{D} \mathbf{A}_0^{-1} \mathbf{D}^T \mathbf{M} (\tilde{\mathbf{x}}_l - \tilde{\mathbf{x}}_j)}} \\ & \mathbf{A}_0^{-1} \mathbf{D}^T \mathbf{M} (\tilde{\mathbf{x}}_l - \tilde{\mathbf{x}}_j), \end{aligned}$$

where $\tilde{\mathbf{x}}$ denotes $\mathbf{D}\mathbf{x}$.

The squared adversarial margin is calculated from the definition of the hyperellipsoid:

$$\begin{aligned} r^2 = & (\mathbf{x}_i - \mathbf{x}_{\min}^{\text{PCA}})^T \mathbf{A}_0 (\mathbf{x}_i - \mathbf{x}_{\min}^{\text{PCA}}) \\ = & \frac{(d_M^2(\tilde{\mathbf{x}}_i, \tilde{\mathbf{x}}_l) - d_M^2(\tilde{\mathbf{x}}_i, \tilde{\mathbf{x}}_j))^2}{4((\tilde{\mathbf{x}}_l - \tilde{\mathbf{x}}_j)^T \mathbf{L}^T \mathbf{L} \mathbf{D} \mathbf{A}_0^{-1} \mathbf{D}^T \mathbf{L}^T \mathbf{L} (\tilde{\mathbf{x}}_l - \tilde{\mathbf{x}}_j))^2} \\ & \cdot (\tilde{\mathbf{x}}_l - \tilde{\mathbf{x}}_j)^T \mathbf{L}^T \mathbf{L} \mathbf{D} \mathbf{A}_0^{-1} \mathbf{A}_0 \mathbf{A}_0^{-1} \mathbf{D}^T \mathbf{L}^T \mathbf{L} (\tilde{\mathbf{x}}_l - \tilde{\mathbf{x}}_j) \\ = & \frac{(d_M^2(\tilde{\mathbf{x}}_i, \tilde{\mathbf{x}}_l) - d_M^2(\tilde{\mathbf{x}}_i, \tilde{\mathbf{x}}_j))^2}{4(\tilde{\mathbf{x}}_l - \tilde{\mathbf{x}}_j)^T \mathbf{M} \mathbf{D} \mathbf{A}_0^{-1} \mathbf{D}^T \mathbf{M} (\tilde{\mathbf{x}}_l - \tilde{\mathbf{x}}_j)} \end{aligned}$$

The perturbation loss is defined similarly to Eq. 4 as follows:

$$\begin{aligned} J_P^{\text{PCA}} = & \frac{1}{|\mathcal{R}|} \sum_{\mathcal{R}} \{ [\tau^2 - \tilde{d}_E^2(\mathbf{x}_i, \mathbf{x}_{i,\min}^{\text{PCA}})] + \mathbb{1}_{\{d_M^2(\tilde{\mathbf{x}}_i, \tilde{\mathbf{x}}_l) > d_M^2(\tilde{\mathbf{x}}_i, \tilde{\mathbf{x}}_j)\}} \\ & + \tau^2 \mathbb{1}_{\{d_M^2(\tilde{\mathbf{x}}_i, \tilde{\mathbf{x}}_l) \leq d_M^2(\tilde{\mathbf{x}}_i, \tilde{\mathbf{x}}_j)\}} \}, \end{aligned}$$

where $\tilde{d}_E^2(\mathbf{x}_i, \mathbf{x}_{i,\min}^{\text{PCA}}) = \frac{(d_M^2(\tilde{\mathbf{x}}_i, \tilde{\mathbf{x}}_l) - d_M^2(\tilde{\mathbf{x}}_i, \tilde{\mathbf{x}}_j))^2}{4((\tilde{\mathbf{x}}_l - \tilde{\mathbf{x}}_j)^T \mathbf{M} \mathbf{D} \mathbf{D}^T \mathbf{M} (\tilde{\mathbf{x}}_l - \tilde{\mathbf{x}}_j) + \epsilon)}$.

The gradient of J_P^{PCA} is given as:

$$\begin{aligned} \frac{\partial J_P^{\text{PCA}}}{\partial \mathbf{M}} = & \frac{1}{|\mathcal{R}|} \sum_{\mathcal{R}} \alpha_{ijl} \left\{ \frac{(d_M^2(\tilde{\mathbf{x}}_i, \tilde{\mathbf{x}}_l) - d_M^2(\tilde{\mathbf{x}}_i, \tilde{\mathbf{x}}_j)) (\tilde{\mathbf{X}}_{ij} - \tilde{\mathbf{X}}_{il})}{2((\tilde{\mathbf{x}}_l - \tilde{\mathbf{x}}_j)^T \mathbf{M} \mathbf{D} \mathbf{A}_0^{-1} \mathbf{D}^T \mathbf{M} (\tilde{\mathbf{x}}_l - \tilde{\mathbf{x}}_j) + \epsilon)} \right. \\ & + \frac{(d_M^2(\tilde{\mathbf{x}}_i, \tilde{\mathbf{x}}_l) - d_M^2(\tilde{\mathbf{x}}_i, \tilde{\mathbf{x}}_j))^2}{4((\tilde{\mathbf{x}}_l - \tilde{\mathbf{x}}_j)^T \mathbf{M} \mathbf{D} \mathbf{A}_0^{-1} \mathbf{D}^T \mathbf{M} (\tilde{\mathbf{x}}_l - \tilde{\mathbf{x}}_j) + \epsilon)^2} \\ & \left. \cdot (\tilde{\mathbf{X}}_{jl} \mathbf{M} \mathbf{D} \mathbf{A}_0^{-1} \mathbf{D}^T + \mathbf{D} \mathbf{A}_0^{-1} \mathbf{D}^T \mathbf{M} \tilde{\mathbf{X}}_{jl}) \right\}. \end{aligned}$$

APPENDIX B

ADDITIONAL ILLUSTRATIONS OF ROBUSTIFYING METRIC LEARNING METHODS

In this section, we incorporate the proposed perturbation loss into two triplet-based metric learning methods, sparse compositional metric learning (SCML) and proxy neighborhood component analysis (PNCA).

A. Sparse compositional metric learning with perturbation loss

We start by briefly revisiting SCML [15]. The core idea is to represent the Mahalanobis distance as a non-negative combination of K basis elements; that is,

$$\mathbf{M} = \sum_{k=1}^K w_k \mathbf{b}_k \mathbf{b}_k^T, \quad \mathbf{w} \geq 0,$$

where the basis set $\{\mathbf{b}_k\}_{k=1}^K$ is generated by using the Fisher discriminative analysis at several local regions. To learn a discriminative metric with good generalization ability, the learning objective comprises a margin-based hinge loss function and an L_1 -norm regularization term as follows:

$$\min_{\mathbf{w}} J_{\text{SCML}} = \frac{1}{|\mathcal{R}|} \sum_{\mathcal{R}} [1 + d_{\mathbf{w}}^2(\mathbf{x}_i, \mathbf{x}_j) - d_{\mathbf{w}}^2(\mathbf{x}_i, \mathbf{x}_l)]_+ + \eta \|\mathbf{w}\|_1,$$

where $\eta \geq 0$ controls the degree of sparsity.

Now, we illustrate how to incorporate the perturbation loss into SCML and solve the associated optimization problem. The solution to the closest adversarial example and the form of perturbation loss remain the same as LMNN-PL; the learning of the Mahalanobis distance is replaced by learning the sparse coefficients, and the optimization problem is solved via the accelerated proximal gradient descent algorithm.

First, we incorporate the proposed perturbation loss J_P (Eq. 4) into the original objective function J_{SCML} :

$$\min_{\mathbf{M} \in \mathbb{S}_+^p} J = J_{\text{SCML}} + \lambda J_P. \quad (10)$$

The squared adversarial margin of Eq. 3 is now a function of \mathbf{w} :

$$d_E^2(\mathbf{x}_i, \mathbf{x}_{i,\min}) = \frac{[d_{\mathbf{w}}^2(\mathbf{x}_i, \mathbf{x}_l) - d_{\mathbf{w}}^2(\mathbf{x}_i, \mathbf{x}_j)]^2}{4d_{\mathbf{w}2}^2(\mathbf{x}_j, \mathbf{x}_l)}$$

$$d_{\mathbf{w}}^2(\mathbf{x}_i, \mathbf{x}_j) = (\mathbf{x}_i - \mathbf{x}_j)^T \left(\sum_{k=1}^K w_k \mathbf{b}_k \mathbf{b}_k^T \right) (\mathbf{x}_i - \mathbf{x}_j)$$

$$d_{\mathbf{w}2}^2(\mathbf{x}_j, \mathbf{x}_l) = (\mathbf{x}_j - \mathbf{x}_l)^T \left(\sum_{k_1=1}^K \sum_{k_2=1}^K w_{k_1} w_{k_2} \mathbf{b}_{k_1} \mathbf{b}_{k_1}^T \mathbf{b}_{k_2} \mathbf{b}_{k_2}^T \right) (\mathbf{x}_j - \mathbf{x}_l).$$

Next, to optimize problem (10), we apply the accelerated proximal gradient descent algorithm with a backtracking step-size rule [61]. The gradient of J_P with respect to \mathbf{w} is as follows:

$$\begin{aligned} \frac{\partial J_P}{\partial \mathbf{w}_k} = & \frac{1}{|\mathcal{R}|} \sum_{\mathcal{R}} \alpha_{ijl} \left\{ \frac{d_M^2(\mathbf{x}_i, \mathbf{x}_l) - d_M^2(\mathbf{x}_i, \mathbf{x}_j)}{2(d_{\mathbf{M}2}^2(\mathbf{x}_j, \mathbf{x}_l) + \epsilon)} \text{tr}(\mathbf{b}_k \mathbf{b}_k^T (\mathbf{X}_{ij} - \mathbf{X}_{il})) \right. \\ & \left. + \frac{[d_M^2(\mathbf{x}_i, \mathbf{x}_l) - d_M^2(\mathbf{x}_i, \mathbf{x}_j)]^2}{4(d_{\mathbf{M}2}^2(\mathbf{x}_j, \mathbf{x}_l) + \epsilon)^2} \text{tr}((\mathbf{b}_k \mathbf{b}_k^T \mathbf{M} + \mathbf{M} \mathbf{b}_k \mathbf{b}_k^T) \mathbf{X}_{jl}) \right\}, \end{aligned}$$

where ϵ is a small constant added to the denominator of $d_E^2(\mathbf{x}_i, \mathbf{x}_{i,\min})$.

B. Proxy neighborhood component analysis with perturbation loss

Neighborhood component analysis is first proposed in [62], and in this paper, we build on one of its latest variants, ProxyNCA++ [22]. ProxyNCA++ aims to maximize the proxy assignment probability as follows:

$$J_{\text{ProxyNCA++}} = -\log \left(\frac{\exp(-d(\frac{\mathbf{x}_i}{\|\mathbf{x}_i\|_2}, \frac{f(\mathbf{x}_i)}{\|f(\mathbf{x}_i)\|_2}) \cdot \frac{1}{T})}{\sum_{f(\mathbf{a}) \in A} \exp(-d(\frac{\mathbf{x}_i}{\|\mathbf{x}_i\|_2}, \frac{f(\mathbf{a})}{\|f(\mathbf{a})\|_2}) \cdot \frac{1}{T})} \right),$$

where $f(\mathbf{a})$ is a proxy function which maps each instance \mathbf{a} to its class proxy, A denotes the set of all proxies, and T is the temperature scaling parameter. As the method is proposed in the context of deep metric learning, \mathbf{x}_i is assumed to be the feature vector after the embedding network and d is chosen as the Euclidean distance. In our work, we will use the original features and the Mahalanobis distance.

There are two sets of learning parameters in ProxyNCA++, the Mahalanobis distance and the proxies. We optimize them alternatively, and in each step, we apply the gradient descent algorithm. The gradient of the ProxyNCA++ term is analogous to the classical NCA and can be found in [62].

The perturbation loss has been revised slightly to include the adversarial margin from instances whose nearest proxy is from a different class:

$$J_P = \frac{1}{n} \sum_{\mathbf{x}_i} \left\{ [\tau^2 - \tilde{d}_E^2(\mathbf{x}_i, \mathbf{x}_{i,\min})]_+ \mathbb{1}[d_M^2(\mathbf{x}_i, \mathbf{a}^-) > d_M^2(\mathbf{x}_i, \mathbf{a}^+)] \right. \\ \left. + (\tau^2 + \tilde{d}_E^2(\mathbf{x}_i, \mathbf{x}_{i,\min})) \mathbb{1}[d_M^2(\mathbf{x}_i, \mathbf{a}^-) \leq d_M^2(\mathbf{x}_i, \mathbf{a}^+)] \right\},$$

where \mathbf{a}^+ and \mathbf{a}^- denote the nearest proxy to \mathbf{x}_i from the same class and different class, respectively. The gradient of the perturbation loss with respect to M is given in Eq. 6; the gradient of the perturbation loss with respect to the proxies is given as follows:

$$\frac{\partial J_P}{\partial \mathbf{a}_i} = \sum_{\mathbf{x} \in A_i^+} \left\{ \frac{d_M^2(\mathbf{x}, \mathbf{a}^-) - d_M^2(\mathbf{x}, \mathbf{a}_i)}{(d_{M^2}^2(\mathbf{a}_i, \mathbf{a}^-) + \epsilon)} M(\mathbf{a}_i - \mathbf{x}) \right. \\ \left. + \frac{[d_M^2(\mathbf{x}, \mathbf{a}^-) - d_M^2(\mathbf{x}, \mathbf{a}_i)]^2}{2(d_{M^2}^2(\mathbf{a}_i, \mathbf{a}^-) + \epsilon)^2} M^2(\mathbf{a}_i - \mathbf{a}^-) \right\} \\ + \sum_{\mathbf{x} \in A_i^-} \left\{ - \frac{d_M^2(\mathbf{x}, \mathbf{a}_i) - d_M^2(\mathbf{x}, \mathbf{a}^+)}{(d_{M^2}^2(\mathbf{a}^+, \mathbf{a}_i) + \epsilon)^2} M(\mathbf{a}_i - \mathbf{x}) \right. \\ \left. + \frac{[d_M^2(\mathbf{x}, \mathbf{a}_i) - d_M^2(\mathbf{x}, \mathbf{a}^+)]^2}{2(d_{M^2}^2(\mathbf{a}^+, \mathbf{a}_i) + \epsilon)^2} M^2(\mathbf{a}_i - \mathbf{a}^+) \right\},$$

where $\mathbf{x} \in A_i^+$ is a set of instances whose nearest proxy is \mathbf{a}_i , and \mathbf{x} and \mathbf{a}_i have the same class label; $\mathbf{x} \in A_i^-$ denotes a set of instances whose different-class nearest proxy is \mathbf{a}_i .

APPENDIX C

PRELIMINARIES AND PROOF OF GENERALIZATION BENEFIT

A. Preliminaries

Definition 2. [26] An algorithm \mathcal{A} is $(K, \epsilon(\cdot), \hat{n}(\cdot))$ pseudo-robust for $K \in \mathbb{N}$, $\epsilon(\cdot) : (\mathcal{Z} \times \mathcal{Z} \times \mathcal{Z})^n \rightarrow \mathbb{R}$ and $\hat{n}(\cdot) : (\mathcal{Z} \times \mathcal{Z} \times \mathcal{Z})^n \rightarrow \{1, \dots, n^3\}$ if $\mathcal{Z} = (\mathcal{X} \times \mathcal{Y})$ can be partitioned into K disjoint sets, denoted by $\{C_k\}_{k=1}^K$, such that for all training samples $\mathbf{s} \in \mathcal{Z}^n$ drawn independently and identically distributed (IID) from the probability distribution \mathcal{P} , there exists a subset of training triplets $\hat{t}_s \subseteq t_s$, with $|\hat{t}_s| = \hat{n}(t_s)$, such that the following holds: $\forall (\mathbf{s}_1, \mathbf{s}_2, \mathbf{s}_3) \in \hat{t}_s$, $\forall \mathbf{z}_1, \mathbf{z}_2, \mathbf{z}_3 \in \mathcal{Z}$, $\forall i, j, l = 1, \dots, K$, if $\mathbf{s}_1, \mathbf{z}_1 \in C_i$, $\mathbf{s}_2, \mathbf{z}_2 \in C_j$ and $\mathbf{s}_3, \mathbf{z}_3 \in C_l$, then

$$|\ell(\mathcal{A}_{t_s}, \mathbf{s}_1, \mathbf{s}_2, \mathbf{s}_3) - \ell(\mathcal{A}_{t_s}, \mathbf{z}_1, \mathbf{z}_2, \mathbf{z}_3)| \leq \epsilon(t_s).$$

Theorem 3. [26] If \mathcal{A} is $(K, \epsilon(\cdot), \hat{n}(\cdot))$ pseudo-robust and the training triplets t_s come from a sample generated by n

IID draws from \mathcal{P} , then for any $\delta > 0$, with probability at least $1 - \delta$ we have:

$$|\mathcal{L}(\mathcal{A}_{t_s}) - \ell_{emp}(\mathcal{A}_{t_s})| \\ \leq \frac{\hat{n}(t_s)}{n^3} \epsilon(t_s) + B \left(\frac{n^3 - \hat{n}(t_s)}{n^3} + 3\sqrt{\frac{2K \ln 2 + 2 \ln 1/\delta}{n}} \right), \quad (11)$$

where \mathcal{L} denotes the expected error, ℓ_{emp} is the empirical error, and B is a constant denoting the upper bound of the loss function ℓ . $|\mathcal{L} - \ell_{emp}|$ is termed the generalization gap.

Definition 4. [63] A δ -cover of a set Θ with respect to a metric ρ is a set $\{\theta^1, \dots, \theta^N\} \subset \Theta$ such that for each $\theta \in \Theta$, there exists some $i \in \{1, \dots, N\}$ such that $\rho(\theta, \theta^i) \leq \delta$. The δ -covering number $N(\delta, \Theta, \rho)$ is the cardinality of the smallest δ -cover.

B. Theorem and proof

Theorem 5. Let M^* be the optimal solution to Eq. 5. Then for any $\delta > 0$, with probability at least $1 - \delta$ we have:

$$|\mathcal{L}(M^*) - \ell_{emp}(M^*)| \\ \leq \frac{\hat{n}(t_s)}{n^3} + B \left(\frac{n^3 - \hat{n}(t_s)}{n^3} + 3\sqrt{\frac{2K \ln 2 + 2 \ln 1/\delta}{n}} \right),$$

where $\hat{n}(t_s)$ denotes the number of triplets whose adversarial margins are larger than τ , B is a constant denoting the upper bound of the loss function (i.e. Eq. 5), and K denotes the number of disjoint sets that partition the input-label space and equals to $|\mathcal{Y}|(1 + \frac{\tau}{\epsilon})^p$.

Proof. After embedding the perturbation loss, learning algorithms that minimize the classical triplet loss, i.e. $[1 + d_M^2(\mathbf{x}_i, \mathbf{x}_j) - d_M^2(\mathbf{x}_i, \mathbf{x}_l)]_+ \cdot \mathbb{1}_{\{y_i=y_j \neq y_l\}}$, are $(|\mathcal{Y}|(1 + \frac{\tau}{\epsilon})^p, 1, \hat{n}(\cdot; \tau))$ pseudo-robust. $\epsilon = 1$ since, by definition of adversarial margin, any \mathbf{x} that falls into the Euclidean ball with center \mathbf{x}_i and a radius of the desired margin τ will satisfy $d_M^2(\mathbf{x}, \mathbf{x}_i) > d_M^2(\mathbf{x}, \mathbf{x}_j)$. Therefore, any change in the triplet loss is bounded by 1. The value of K can be determined via the covering number [26]. The instance space \mathcal{X} can be partitioned by using the covering number $N(\tau, \mathcal{X}, \|\cdot\|_2)$. By normalizing all instances to have unit L_2 -norm, we obtain a finite covering number as $N \leq (1 + \frac{\tau}{\epsilon})^p$ [63]. The label space \mathcal{Y} can be partitioned into $|\mathcal{Y}|$ sets. Therefore, the number of disjoint sets, i.e. K , is always smaller than $|\mathcal{Y}|(1 + \frac{\tau}{\epsilon})^p$. \square

APPENDIX D

EXPERIMENTAL SETUP AND ADDITIONAL RESULTS

A. Datasets

Table V lists information on sample size, feature dimension and class information. The class distribution of datasets Ecoli, Yeast and APS Failure is highly imbalanced and their imbalanced ratios (IR) are also listed.

B. Experimental setting

Hyperparameter tuning of compared methods: LMNN, SCML, and DRIFT are implemented by using the official codes provided by the authors; all parameters are set as default apart from the trade-off parameter. For LMNN, the trade-off parameter μ is chosen from $\{0.1, 0.2, \dots, 0.9\}$. For




Fig. 5: Robustness of LMNN-based methods on the Australian dataset with different noise types and intensities.

TABLE V: Characteristics of the datasets.

Dataset	# instances (training,test,valid.)	# features (reduced dimensions)	# classes	# rounds
UCI data				
Australian	690	14	2	20
Breast cancer	683	9	2	20
Ecoli (0-1-4-7 vs 2-3-5-6)	336	7	2 (IR:10.59)	20
Fourclass	862	2	2	20
Haberman	306	3	2	20
Iris	150	4	3	20
Segment	2310	19	7	20
Sonar	208	60	2	20
Voting	435	16	2	20
WDBC	569	30	2	20
Wine	178	13	3	20
Yeast (CYT vs POX)	483	8	2 (IR:23.15)	20
Credit (Taiwan)	30000 (10k,10k,10k)	23	2	1
High-dimensional data				
Isolet	7797 (1560,1558)	617 (170)	26	4
MNIST	4000 (2000,2000)	784 (141)	10	1
APS Failure	76000 (40k,20k,16k)	171 (79)	2 (IR:58.90)	1

SCML, the weight of the regularization term η is chosen from $\{10^{-5}, 10^{-4}, \dots, 10^3\}$, and the number of bases is set as 200, 400 and 1000 for UCI datasets whose sample size is smaller than 500, larger than 500, and large-scale/high-dimensional datasets (i.e. Credit, Isolet, and MNIST), respectively. For LDD, the regularizer weight is chosen from $\{10^{-6}, 10^{-5}, \dots, 10^2\}$. For CAP, the regularizer weight is chosen from $\{10^{-3}, 10^{-2}, \dots, 10\}$, and the rank of \mathbf{M} is chosen from 10 values equally spaced between 1 and p . For PNCA, the temperature scaling parameter is chosen from $\{10^{-2}, 10^{-1}, \dots, 10^2\}$, and the number of proxies per class is chosen from $[1, 2, 3, 5, 10]$ for all UCI datasets except $[1, 3, 5, 10, 20]$ for Credit. For DRIFT, we search the grid suggested by the authors. Trade-off parameters are tuned via five-fold cross-validation on the training data or on a separate validation set as listed in Table V. For all these methods, triplet constraints are generated from 3 target neighbors and 10 nearest impostors, calculated under the Euclidean distance.

Experimental setting of LMNN-PL: \mathbf{M} is initialized as the identity matrix. The learning rate γ is initialized to 1. Following [14]’s work, γ is increased by 1% if the loss function decreases and decreased by 50% otherwise. The

training stops if the relative change in the objective function is smaller than the threshold of 10^{-7} or reaches the maximum number of iterations of 1000.

Experimental setting of SCML-PL: SCML-PL is tuned in the same manner as LMNN-PL via random search; the range of η and the number of bases are same as SCML, and the ranges of τ and λ are same as LMNN-PL. The method is optimised via the accelerated proximal gradient descent algorithm with a backtracking stepsize rule [61]. The initial learning rate is set as 1 and the shrinkage factor is set as 0.8. \mathbf{w} is initialized as the unit vector.

Experimental setting of PNCA-PL: Ranges of hyperparameters are same as PNCA and LMNN-PL, except that λ is chosen from $\{10^{-2}, 10^{-1}, \dots, 10^2\}$. In addition, for training stability, we update proxies only in the first 200 iterations and fix them afterwards.

C. Additional experimental results

In addition to the isotropic Gaussian noise presented in the main text, we test the robustness performance against anisotropic Gaussian noise and adversarial perturbations. Anisotropic Gaussian noise is generated from a zero-mean Gaussian with a diagonal covariance matrix and unequal variances estimated from the training data; its noise intensity is determined via SNR. Adversarial perturbations are computed according to Eq. 2 with the nearest target neighbor and impostor found from using the learned distance metric; the magnitude of perturbations is controlled via L_2 norm.

To start with, we conduct an in-depth experiment on the Australian dataset by altering the noise intensity. Fig. 5a plots the classification accuracy of LMNN-based methods under different levels of isotropic Gaussian noise (equal variances). When the noise intensity is low, the performance of LMNN and LMNN-PL remain stable. When the noise intensity increases to the SNR of 10 dB or 5 dB, the performances of both method degrade. Owing to the enlarged adversarial margin, the influence on LMNN-PL is slightly smaller than that on LMNN. When the SNR equals 1 dB, the performance gain from using LMNN-PL becomes smaller. This result is reasonable as the desired margin τ is selected according to the criterion of classification accuracy and hence may be too small to withstand a high level of noise. LMNN-PL surpasses

TABLE VI: Classification accuracy (or G-means indicated by an asterisk) of 3NN on datasets contaminated with anisotropic Gaussian noise (SNR=5 dB).

Dataset	Euclidean	LMNN-based					SCML-based		PNCA-based		
		LMNN	LDD	CAP	SN	DRIFT	LMNN-PL	SCML	SCML-PL	PNCA	PNCA-PL
Australian	81.88±1.72	82.13±1.52	82.57±1.55	81.82±1.52	82.20±1.54	81.97±1.53	82.90±1.53	82.59±1.70	82.84±1.64	83.21±1.69	83.68±1.81
Breast cancer	96.76±1.06	96.24±1.06	96.66±1.07	96.27±1.03	96.27±1.07	96.61±0.97	96.69±1.07	96.34±1.03	96.63±1.03	96.05±1.39	96.74±1.13
Ecoli*	78.28±7.53	75.13±7.58	<u>76.97±7.64</u>	72.64±9.06	75.27±7.60	76.18±7.65	78.58±7.53	75.02±6.89	77.14±7.97	90.91±3.17	91.40±2.55
Fourclass	69.20±1.11	67.74±1.25	68.84±1.14	67.84±1.19	68.61±1.26	69.13±1.02	69.04±1.11	68.22±1.10	68.96±1.11	70.54±2.14	69.49±4.62
Haberman	70.30±1.89	70.21±1.84	70.25±1.82	69.39±2.06	70.23±1.85	69.31±2.50	70.25±1.90	69.98±1.64	70.24±1.85	74.72±3.15	74.76±3.09
Iris	79.84±3.28	78.75±2.96	<u>79.04±3.17</u>	77.90±3.31	78.84±3.08	78.57±3.09	79.20±3.08	78.32±3.60	79.18±3.13	78.13±4.55	78.34±3.91
Segment	86.27±0.70	79.03±3.37	<u>83.49±1.17</u>	82.77±2.49	80.04±3.25	83.88±1.33	82.13±2.70	61.28±9.78	62.86±8.76	76.30±2.74	78.28±3.32
Sonar	83.50±3.19	83.54±4.30	86.18±2.93	85.44±2.79	83.50±4.32	84.65±3.30	84.99±3.13	76.91±4.32	79.49±3.80	80.71±4.35	81.76±3.43
Voting	93.19±1.14	94.01±1.00	<u>94.24±1.13</u>	94.37±1.17	94.05±1.01	93.94±1.12	94.64±1.21	93.99±1.15	94.65±1.09	92.61±1.64	93.45±1.77
WDBC	95.89±1.31	92.01±1.65	96.30±0.94	96.14±1.11	92.67±1.69	96.02±0.88	96.07±0.89	95.75±1.29	96.22±1.14	96.02±1.51	96.24±1.14
Wine	94.13±1.47	93.27±1.62	93.97±1.38	93.87±1.49	93.38±1.68	94.55±1.15	94.44±1.21	93.92±1.55	94.52±1.33	94.08±1.93	93.37±1.49
Yeast*	69.91±10.46	59.00±8.78	68.63±10.66	65.74±10.26	59.58±8.98	<u>70.15±10.56</u>	70.23±10.48	56.43±8.74	60.32±9.07	61.84±18.45	63.70±18.69
Credit	76.01	<u>76.07</u>	<u>76.22</u>	74.52	76.29	<u>76.12</u>	76.06	75.73	75.73	78.89	79.73
# outperform	-	12	9	10	12	9	-	13	-	11	-

TABLE VII: Classification accuracy (or G-means indicated by an asterisk) of 3NN on datasets contaminated with adversarial perturbations (L_2 -norm=0.2).

Dataset	Euclidean	LMNN-based					SCML-based		PNCA-based		
		LMNN	LDD	CAP	SN	DRIFT	LMNN-PL	SCML	SCML-PL	PNCA	PNCA-PL
Australian	79.09±2.57	79.98±2.95	80.79±2.72	76.95±3.86	80.14±2.78	78.87±3.13	81.78±2.53	79.64±3.11	81.08±2.58	83.89±2.91	84.35±3.04
Breast cancer	96.29±1.57	95.44±1.69	96.00±1.84	95.68±1.44	95.46±1.77	95.85±1.67	96.00±1.53	95.59±1.60	95.71±1.66	96.37±1.91	96.66±1.67
Ecoli*	80.60±11.55	72.82±16.67	80.70±11.42	68.78±21.39	73.84±16.43	78.59±10.55	<u>79.97±13.37</u>	74.80±14.62	75.20±13.99	71.58±20.12	79.07±9.35
Fourclass	66.27±2.64	60.73±4.97	66.53±2.34	64.34±3.18	64.05±4.44	66.22±2.66	<u>66.35±3.00</u>	64.90±3.18	66.49±2.89	72.80±3.32	71.29±5.21
Haberman	57.58±4.26	57.25±5.28	57.19±4.67	<u>57.53±5.07</u>	57.19±5.05	59.16±6.17	57.30±4.55	56.40±5.37	57.02±3.75	74.94±4.39	75.45±4.38
Iris	78.67±5.37	76.33±6.70	<u>76.67±6.22</u>	76.78±5.42	76.11±7.06	76.44±6.93	76.78±6.39	74.00±6.99	76.67±6.26	78.22±8.14	77.67±7.33
Segment	76.10±1.70	56.80±6.21	<u>67.86±3.12</u>	67.37±6.19	59.21±6.54	71.31±2.83	65.53±5.17	29.49±9.00	30.56±12.95	68.71±4.31	71.61±4.66
Sonar	71.75±6.11	42.38±16.87	<u>64.13±6.36</u>	58.33±6.89	41.67±16.92	64.44±5.46	56.03±8.43	28.17±11.22	40.24±11.81	60.40±5.93	61.90±11.01
Voting	90.88±2.06	90.88±2.97	<u>91.11±2.62</u>	91.03±2.70	89.47±3.05	91.22±2.26	92.71±1.99	90.08±2.46	92.25±2.92	87.79±3.70	89.89±6.93
WDBC	93.57±1.89	59.27±12.73	91.05±2.08	90.56±3.02	65.03±12.00	<u>90.91±2.40</u>	90.41±3.58	87.51±4.15	90.73±2.97	90.85±3.24	91.64±2.64
Wine	90.28±3.23	84.07±6.02	90.09±3.08	86.94±3.53	84.63±6.16	87.59±3.40	88.43±4.07	86.02±5.39	88.98±4.55	88.70±3.75	86.20±4.55
Yeast*	70.21±10.47	53.42±26.38	70.19±10.50	68.15±12.59	52.15±25.54	70.14±10.49	70.19±10.49	59.83±13.44	62.61±18.96	61.65±19.36	66.65±19.05
Credit	67.22	67.14	67.21	65.02	67.29	<u>67.23</u>	66.96	66.03	66.44	80.60	80.27
# outperform	-	12	6	9	12	8	-	13	-	9	-

all other LMNN-based methods until the noise intensity is very large. Fig. 5b plots the accuracy under anisotropic Gaussian noise (unequal variances). Compared with the case of isotropic Gaussian noise, the degradation of all methods is more pronounced in this case, but the pattern remains similar. Fig. 5c presents the results under adversarial perturbations. The proposed method achieves the highest accuracy over the range of perturbation size. The method LDD is also quite robust to adversarial perturbations. This should not be surprising as it encourages learning a metric close to the Euclidean distance, and the Euclidean distance is less sensitive to perturbation than the discriminative Mahalanobis distance.

Tables VI and VII list the classification accuracy or G-means for all datasets contaminated by the anisotropic Gaussian noise with the SNR of 5 dB and the adversarial perturbations with the perturbation size of 0.2, respectively. Comparing the case of adversarial perturbations with that of Gaussian-type noises, we see three clear differences. First, LMNN performs much worse than the Euclidean distance in the presence of adversarial perturbations, especially on datasets with a large number of features relative to the sample size such as Sonar and WDBC. A potential reason is that the learned metric will stretch the distances in a few directions for discriminability, and hence adding perturbations in these directions is very likely to change the decision of k NN. Secondly, robustness

TABLE VIII: Generalization and robustness of metric learning methods on high-dimensional datasets (additional results).

Isolet						
Method	Clean	IG,SNR=20	IG,SNR=5	AG,SNR=20	AG,SNR=5	Margin
Euclidean	84.16±4.09	83.93±4.30	82.73±3.70	83.93±4.30	83.18±3.59	0.1009
AML	86.75±3.16	86.59±3.49	85.97±3.69	86.69±3.59	86.24±3.82	0.1261
LDD	90.91±3.90	90.81±4.12	87.97±3.83	90.75±4.12	89.13±4.05	0.1333
MNIST						
Method	Clean	IG,SNR=20	IG,SNR=5	AG,SNR=20	AG,SNR=5	Margin
Euclidean	88.70	88.90	88.25	89.05	88.05	0.2091
AML	89.25	88.70	88.85	89.30	89.20	0.2142
LDD	90.85	90.85	87.90	90.95	90.30	0.2232
DRIFT	90.85	90.75	87.45	90.65	89.45	0.2054

DRIFT is unable to learn a metric on Isolet and hence is not reported.

to adversarial perturbations differs markedly across the types of methods. SCML-based deteriorates drastically on some datasets while PNCA-based methods are much more robust. Thirdly, the proposed method is effective in safeguarding the baseline methods against adversarial perturbations on most of the datasets.

Table VIII is a supplement to Table III of main text. It reports the performance of Euclidean, AML, LDD and DRIFT on high-dimensional datasets.

D. Convergence analysis

The objective function of the proposed perturbation loss (Eq. 4) includes several indicator functions. As a consequence, we cannot guarantee that the optimization algorithm will converge to a global or local optimum. Empirically, we decay the learning rate as training progresses and stop training when the change in the objective function is small or if it reaches the maximum iteration. Figure 6 shows the classification accuracy of the proposed LMNN-PL on the MNIST dataset along training. The test accuracy increases in the initial stage, oscillates afterwards, and stabilizes after around 500 iterations.




Fig. 6: Convergence curve of LMNN-PL on MNIST.