

A Fast and Easy Regression Technique for k -NN Classification Without Using Negative Pairs

Yutaro Shigeto, Masashi Shimbo^(✉), and Yuji Matsumoto

Nara Institute of Science and Technology, Ikoma, Nara, Japan
{yutaro-s, shimbo, matsu}@is.naist.jp

Abstract. This paper proposes an inexpensive way to learn an effective dissimilarity function to be used for k -nearest neighbor (k -NN) classification. Unlike Mahalanobis metric learning methods that map both query (unlabeled) objects and labeled objects to new coordinates by a single transformation, our method learns a transformation of labeled objects to new points in the feature space whereas query objects are kept in their original coordinates. This method has several advantages over existing distance metric learning methods: (i) In experiments with large document and image datasets, it achieves k -NN classification accuracy better than or at least comparable to the state-of-the-art metric learning methods. (ii) The transformation can be learned efficiently by solving a standard ridge regression problem. For document and image datasets, training is often more than two orders of magnitude faster than the fastest metric learning methods tested. This speed-up is also due to the fact that the proposed method eliminates the optimization over “negative” object pairs, i.e., objects whose class labels are different. (iii) The formulation has a theoretical justification in terms of reducing hubness in data.

1 Introduction

Let \mathcal{X} be a feature space and \mathcal{Y} be a set of class labels. The k -nearest neighbor (k -NN) classifier predicts the class label of an unknown query object $\mathbf{x} \in \mathcal{X}$ by its nearest neighbors. Given \mathbf{x} and a set of labeled objects $\mathcal{D} = \{(\mathbf{x}_i, y_i)\}_{i=1}^n$ where $\mathbf{x}_i \in \mathcal{X}$ is the feature vector of the i th object and $y_i \in \mathcal{Y}$ its class label, the classifier first computes the distance between \mathbf{x} and each labeled object \mathbf{x}_i . It then predicts the class label \hat{y} of \mathbf{x} by the majority among its k nearest labeled objects. When $k = 1$, the decision rule of the k -NN (1-NN) classifier is simply:

$$\hat{y} = \arg \min_{y_i : (\mathbf{x}_i, y_i) \in \mathcal{D}} f(\mathbf{x}, \mathbf{x}_i), \quad (1)$$

where function $f : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$ is some distance/dissimilarity function.

Obviously, the choice of function f affects the accuracy of classification. Therefore, many researchers [2, 11, 13, 15] have tackled *metric learning*, which is the task of learning a suitable distance function from data.

For Euclidean object space $\mathcal{X} = \mathbb{R}^d$, metric learning is usually formulated as the task of finding a Mahalanobis distance. In this formulation, the distance between two objects $\mathbf{x}, \mathbf{z} \in \mathbb{R}^d$ is defined by

$$f(\mathbf{x}, \mathbf{z}) = \sqrt{(\mathbf{x} - \mathbf{z})^T \mathbf{M} (\mathbf{x} - \mathbf{z})}, \quad (2)$$

with some positive (semi)definite matrix \mathbf{M} . By defining matrix \mathbf{L} by $\mathbf{M} = \mathbf{L}^T \mathbf{L}$, we can write the distance in Eq. (2) as

$$f(\mathbf{x}, \mathbf{z}) = \|\mathbf{L}\mathbf{x} - \mathbf{L}\mathbf{z}\|. \quad (3)$$

This equation shows that learning Mahalanobis distance is equivalent to learning a suitable linear transformation \mathbf{L} .

In the context of k -NN classification, distance needs to be measured only between query (unlabeled) objects and labeled objects, as can be seen from Eq. (1); when distance $f(\mathbf{x}, \mathbf{z})$ is computed, the first object \mathbf{x} is always a query object, and the second object \mathbf{z} is always a labeled object \mathbf{x}_i . Moreover, function f need not be metric and can be any measure of dissimilarity; for instance, f being asymmetric is perfectly acceptable.

In this paper, we learn one such dissimilarity function. The idea is to compute a transformation of labeled objects to new points while unlabeled objects are kept at their original points. Thus, our objective is to find a suitable matrix \mathbf{W} that defines a dissimilarity function

$$f(\mathbf{x}, \mathbf{z}) = \|\mathbf{x} - \mathbf{W}\mathbf{z}\|, \quad (4)$$

where \mathbf{x} is a query object, and \mathbf{z} is a labeled object.

Because the coordinates of query objects are fixed, our formulation might appear less flexible than Mahalanobis distance learning (Eq. (3)). However, as shown in a subsequent section, it gives a better k -NN classification accuracy than Mahalanobis distance learning methods on many datasets that feature high-dimensional space. Moreover, optimizing \mathbf{W} in Eq. (4) is much easier and substantially (often more than two orders of magnitude) faster.

The effectiveness of the proposed approach has a theoretical foundation in terms of reducing *hubness* in data [7, 9]. Recent studies have shown that the presence of hubs, which are a few objects that appear in the k -NNs of many objects, is an obstacle that can harm the performance of many vector space methods [7, 8, 10]. We show that metric learning is no exception, and the transformation of labeled objects restrains hubs from emerging. This approach is justified by a recent result of Shigeto et al. [9], who used regression to reduce hubness in zero-shot problems. In their work, the problem was cast as a task of cross-domain matching, whereas in this paper, we are concerned with improving the accuracy of k -NN classification in a single space.

Another notable feature of the proposed method is that it eliminates optimization over “negative” object pairs, i.e., objects belonging to difference classes. In other words, our method only attempts to make objects of the same class (“positive” object pairs) to be closer. Its objective function does not have any

constraints or terms that promote negative object pairs to be apart from each other. Such constraints are indispensable in Mahalanobis metric learning to prevent trivial solutions $\mathbf{M} = \mathbf{O}$ in Eq. (2) or $\mathbf{L} = \mathbf{O}$ in Eq. (3), and metric learning typically optimizes over a large number of negative object pairs. Moreover, incorporating negative pairs results in a non-convex optimization problem with respect to matrix \mathbf{L} . Existing metric learning methods [2, 5, 11, 13, 15] hence resorts to optimizing $\mathbf{M} = \mathbf{L}^T \mathbf{L}$ using computationally intensive methods such as semi-definite programming. By contrast, since we only transforms labeled objects, we need not worry about $\mathbf{W} = \mathbf{O}$ being the solution (see Eq. (4)), thus eliminating the need of negative pairs. This makes the solution easily obtained with ridge regression, which contributes to reduced computation time.

2 Related Work

We briefly review some of the metric learning methods, mostly those used in the experiments in Sect. 5. For comprehensive survey of the field, see [1, 6].

A majority of the metric learning methods adopt Mahalanobis distance (Eq. (3)) as the distance function, and minimize the training loss under various constraints. As mentioned earlier, these methods do not make distinction between unlabeled (query) objects and labeled objects, in the sense that their coordinates are transformed by the same matrix, \mathbf{L} in Eq. (3). Our approach differs from these methods in that it projects only the labeled objects to new coordinates.

There are various strategies for learning Mahalanobis distance. Xing et al. [13] formulated metric learning as a convex optimization problem, and demonstrated its effectiveness in clustering tasks. The *large-margin nearest neighbor* (LMNN) method [11] is probably the most popular of all metric learning methods. Its objective is to minimize distances between objects with the same label, and to penalize objects with different labels when they are closer than a certain distance. Hence objects from different classes are separated by a large margin. To make the problem convex, in Xing et al.’s method and LMNN, optimization is done over not \mathbf{L} but $\mathbf{M} = \mathbf{L}^T \mathbf{L}$, with semidefinite programming. Ying and Li [15] presented an eigenvalue optimization framework for learning Mahalanobis distance. Davis et al. [2] proposed *information-theoretic metric learning* (ITML). ITML minimizes the LogDet divergence subject to linear constraints. It thus requires no eigenvalue computation or semi-definite programming.

Although it has been shown that these methods work well in many applications, learning Mahalanobis distance typically incurs high computational cost. Indeed, as we show in an experiment (Sect. 5), these methods spend substantial time in optimizing \mathbf{M} , when applied to large datasets.

3 Proposed Method

In this section, we present our approach for improving the k -NN classification accuracy.

In nearly all metric learning methods, the objective function to be optimized involves a term that encourages objects of the same class to be placed closer. In the same vein, our method also optimizes the transformation matrix \mathbf{W} in Eq. (4) by minimizing the distance between objects of the same class. However, in our formulation, the learned transformation \mathbf{W} is only applied to labeled objects.

Our training procedure consists of two steps. We first make training object pairs for which the distance should be minimized. To this end, we follow Weinberger and Saul [11]: for each labeled object $\mathbf{x}_i \in \mathbb{R}^d$ in the training set, we define its “target” objects \mathcal{T}_i to be the k objects in the training set that belong to the same class as \mathbf{x}_i and are closest to \mathbf{x}_i as measured by the original Euclidean distance (i.e., the one before training). We then find a matrix $\mathbf{W} \in \mathbb{R}^{d \times d}$ that moves objects in \mathcal{T}_i towards \mathbf{x}_i , by solving the following optimization problem:

$$\min_{\mathbf{W}} \sum_{i=1}^n \sum_{\mathbf{z} \in \mathcal{T}_i} \|\mathbf{x}_i - \mathbf{W}\mathbf{z}\|^2 + \lambda \|\mathbf{W}\|_F^2, \quad (5)$$

where $\lambda \geq 0$ is a hyperparameter for regularization and $\|\cdot\|_F$ represents the Frobenius norm. Equation (5) is a familiar objective function of ridge regression, and we have the closed-form solution:

$$\mathbf{W} = \mathbf{X}\mathbf{J}\mathbf{X}^T(\mathbf{X}\mathbf{X}^T + \lambda\mathbf{I})^{-1}, \quad (6)$$

where $\mathbf{X} = [\mathbf{x}_1, \dots, \mathbf{x}_n] \in \mathbb{R}^{d \times n}$, and $\mathbf{J} \in \{0, 1\}^{n \times n}$ is an indicator matrix such that $[\mathbf{J}]_{i,j} = 1$ if $\mathbf{x}_j \in \mathcal{T}_i$ and 0 otherwise.

In the test phase, we first compute the image $\mathbf{x}'_i = \mathbf{W}\mathbf{x}_i$ of every labeled object \mathbf{x}_i by the learned matrix \mathbf{W} . We then carry out k -NN classification by regarding $\mathcal{D}' = \{(\mathbf{x}'_i, y_i)\}$ as the labeled objects in place of the original ones, $\mathcal{D} = \{(\mathbf{x}_i, y_i)\}$. In the case of 1-NN classification, for example, this amounts to using the dissimilarity function f given by Eq. (4) in the decision rule of Eq. (1), i.e.,

$$\hat{y} = \arg \min_{y_i: (\mathbf{x}'_i, y_i) \in \mathcal{D}'} \|\mathbf{x} - \mathbf{x}'_i\|^2 = \arg \min_{y_i: (\mathbf{x}_i, y_i) \in \mathcal{D}} \|\mathbf{x} - \mathbf{W}\mathbf{x}_i\|^2. \quad (7)$$

4 Proposed Method Reduces Hubness

In this section, we argue that the proposed method is by design less susceptible to producing hubs [7] in the transformed labeled objects. This property is desirable, as hubs have been recognized as one of the major factors that harm the effectiveness of nearest neighbor methods.

4.1 Hubness Phenomenon

The hubness phenomenon [7] states that a small number of objects in the dataset, called *hubs*, may occur as the nearest neighbor of many objects. The presence

of hubs will diminish the utility of nearest-neighbor methods, because the lists of nearest neighbors frequently contain the same hub objects regardless of the query.

Hubs occur in data because of an inherent bias present in Euclidean space, called *spatial centrality* [7]; i.e., objects closest to the mean of the data tend to be the nearest neighbors of many objects. This bias is known to grow stronger with the dimensionality of the space.

The following proposition by Shigeto et al. [9] quantifies the degree of spatial centrality as a function of the dimension of the space and the variance of data distribution, when the feature values of query and data follow zero-mean Gaussian distributions with different variances. Let $E_{\mathcal{X}}[\cdot]$ and $\text{Var}_{\mathcal{X}}[\cdot]$ respectively denote the expectation and variance under a distribution \mathcal{X} .

Proposition 1 [9, Proposition 1]. *Let \mathbf{z} be a d -dimensional random vector sampled i.i.d. from a normal distribution with zero means and diagonal covariance matrix $s^2\mathbf{I}$; i.e., $\mathbf{z} \sim \mathcal{Z}$, where $\mathcal{Z} = \mathcal{N}(\mathbf{0}, s^2\mathbf{I})$. Further let $\sigma = \sqrt{\text{Var}_{\mathcal{Z}}[\|\mathbf{z}\|^2]}$ be the standard deviation of the squared norm $\|\mathbf{z}\|^2$.*

Consider two fixed samples \mathbf{z}_1 and \mathbf{z}_2 of random vector \mathbf{z} , such that the squared norms of \mathbf{z}_1 and \mathbf{z}_2 are $\gamma\sigma$ apart. In other words,

$$\|\mathbf{z}_2\|^2 - \|\mathbf{z}_1\|^2 = \gamma\sigma.$$

Let \mathbf{x} be a point sampled from a distribution \mathcal{X} with zero mean.

Then, the expected difference between the squared distances from \mathbf{x} to \mathbf{z}_1 and \mathbf{z}_2 is given by

$$\begin{aligned} \Delta &= E_{\mathcal{X}} [\|\mathbf{x} - \mathbf{z}_2\|^2] - E_{\mathcal{X}} [\|\mathbf{x} - \mathbf{z}_1\|^2] \\ &= \gamma s^2 \sqrt{2d}. \end{aligned} \tag{8}$$

The quantity in Eq. (8) can be interpreted as the degree of the spatial centrality bias present in the data, which causes hub formation. If \mathbf{z}_1 is closer to the origin (data mean) than \mathbf{z}_2 is, $\Delta > 0$ because in this case $\gamma > 0$.

This implies that a query object \mathbf{x} sampled from \mathcal{X} is more likely to be closer to object \mathbf{z}_1 than to \mathbf{z}_2 ; i.e., given query object \mathbf{x} , \mathbf{z}_1 is more likely to become its nearest neighbor. Because this reasoning applies to any pair of objects \mathbf{z}_1 and \mathbf{z}_2 in the dataset, it can be concluded that the objects closest to the data mean is most likely to be a hub.

Further, the factor s^2 in Eq. (8) suggests that, for a fixed query distribution \mathcal{X} , the data distribution \mathcal{Y} with smaller variance s is preferable to reduce spatial centrality, and hence hubness as well.

4.2 Hubness and the Proposed Method

Ridge regression reduces the variance of mapped feature values (observables) relative to that of target (response) variables; see, for example, Shigeto et al. [9, Proposition 2]. Thus, in our model of Eq. (5), the variance of the components of

the mapped objects $\mathbf{W}\mathbf{z}$ tends to be smaller than that of \mathbf{x} . From the discussion on hubness in Sect. 4.1, reducing the variance of data objects (which correspond to the image $\mathbf{W}\mathbf{z}$ of the labeled objects \mathbf{z} in the proposed method) relative to the query (unlabeled object \mathbf{x}) can reduce the spatial centrality. By combining these arguments, we expect that the proposed approach should alleviate the emergence of hubs, and, consequently, improve the accuracy of k -NN classification.

Note that we could think of a different regression problem in which query object \mathbf{x} , not labeled object \mathbf{z} , is mapped to new coordinates:

$$\min_{\mathbf{W}} \sum_{i=1}^n \sum_{\mathbf{z} \in \mathcal{T}_i} \|\mathbf{W}\mathbf{x}_i - \mathbf{z}\|^2 + \lambda \|\mathbf{W}\|_{\text{F}}^2. \quad (9)$$

This would result in function f as follows:

$$f(\mathbf{x}, \mathbf{z}) = \|\mathbf{W}\mathbf{x} - \mathbf{z}\|^2. \quad (10)$$

However, this dissimilarity function is useless as it actually *promotes* hubness. The variance of the transformed query objects shrinks as a result of regression. Thus, in this model, the variance of the labeled objects is made larger than the transformed query objects, but this is not a desirable situation according to Proposition 1. We also verify this in one of the experiments in Sect. 5.

5 Experiments

We evaluate the proposed approach on various classification tasks. The objective of these experiments is to investigate whether the proposed approach can reduce the emergence of hubs, and improve the performance of k -NN classification. The performance is measured against several popular metric learning methods.

5.1 Experimental Setups

Dataset Description. Three types of datasets were used for our evaluation: UCI, document, and image datasets.

From the UCI machine learning datasets,¹ we chose balance-scale, glass, ionosphere, iris, and wine, as they are frequently used for evaluation in metric learning literature [2, 5, 11, 15]. However, they are mostly toy problems, and their small feature dimensions, the numbers of labels and objects do not necessarily reflect real-world problems. We therefore used document and image datasets also for our evaluation.

For document and image classification, support vector machines are known to provide state-of-the-art accuracy. Notice, however, that our goal is not to design a state-of-the-art classifier. Rather, the main objective of this experiments is to evaluate the performances of the proposed method in comparison with metric learning methods, and to show its usefulness for k -NN classification.

¹ <http://archive.ics.uci.edu/ml/>.

Table 1. Dataset statistics. In document and image datasets, “original dim.” indicates the number of raw dimensions before applying PCA.

(a) UCI datasets.					
dataset	ionosphere	balance-scale	iris	wine	glass
#objects	351	625	150	178	214
#classes	2	3	3	3	6
dimension	34	4	4	13	9

(b) Document datasets.					(c) Image datasets.				
dataset	RCV	News	Reuters	TDT	dataset	AwA	CUB	SUN	aPY
#objects	9625	18846	8213	10021	#objects	30475	11788	14340	15339
#classes	4	20	41	56	#classes	50	200	717	32
dimension	300	300	300	300	dimension	300	300	300	300
original dim.	29992	26214	18933	36771	original dim.	4096	4096	4096	4096

For document classification tasks, we used four publicly available document datasets: RCV1-v2 (RCV), 20 newsgroups (News), Reuters21578 (Reuters), and TDT2 (TDT).² In Reuters21578 and TDT2, we removed minority classes that hold less than 10 objects in the dataset. After this removal, Reuters21578 and TDT2 had 56 and 41 classes, respectively.

For image classification, we used the following image datasets: aPascal & aYahoo (aPY), Animals with Attributes (AwA), Caltech-UCSD Birds-200-2011 (CUB), and SUN Attribute.³

The computational cost of metric learning methods is heavily dependent on the dimension of the feature space. In our preliminary experiment, training of the metric learning methods (LMNN, ITML, and DML-eig; see below) did not complete in a reasonable time on document and image datasets. We therefore had to use principal component analysis to reduce the dimensionality of features to 300 for these datasets.⁴

The dataset statistics are summarized in Table 1.

All data (set of feature vectors) were centered before training. For the wine dataset, we further converted the features to z-scores, following the remark on the UCI website that a k -NN classifier achieved a high accuracy with this standardization.

² Datasets were downloaded from <http://www.cad.zju.edu.cn/home/dengcai/Data/TextData.html>.

³ We used the publicly available features from <https://zimingzhang.files.wordpress.com/2014/10/cnn-features1.key>.

⁴ We also conducted experiments where the dimensionality of features was set to 100. The results are not presented here due to lack of space, but the same trend was observed.

Each dataset was randomly split into training (70%) and test (30%) sets. Experiments were repeated on four different random splits, for which we report the average performance.

Compared Methods. We trained distance/dissimilarity functions using the following methods, and carried out k -NN classification on the datasets above.

- original metric: Euclidean distance in the original feature space, without any training. This is the baseline.
- LMNN: Large margin nearest neighbor classification [11]. This method has often been used in distance metric learning experiments as a baseline.
- ITML: Information theoretic metric learning [2].
- DML-eig: Distance metric learning with eigenvalue optimization [15].
- Move-labeled (proposed method): Learning to move labeled objects toward query objects. This is our proposed approach that optimizes the ridge regression objective of Eq. (5), and predicts the labels using Eq. (7).
- Move-query: Learning to move query (unlabeled) objects toward fixed labeled objects. This is the method discussed in Sect. 4.2 and optimizes Eq. (9). Like Move-labeled it is also based on ridge regression, but the roles of the input and response variables are exchanged. The resulting dissimilarity function of Eq. (10) is then used for k -NN classification.

Notice that Move-query was tested only to verify the claim made in Sect. 4.2; i.e., although both Move-labeled and Move-query are based on ridge regression, the proposed method, Move-labeled, is expected to perform well by reducing hubness, whereas Move-query is expected to do the contrary, by *promoting* hubness. It was therefore not meant as a competitive method.

LMNN, ITML, and DML-eig learn a Mahalanobis distance. For these methods, we used the publicly available MATLAB implementations provided by the respective authors⁵. We implemented the proposed method also in MATLAB⁶, for fair evaluation of running time.

For LMNN, Move-labeled, and Move-query, the number of target objects for each training object was set to 1; i.e., for each object \mathbf{x}_i in the training set, we made a training pair $(\mathbf{x}_i, \mathbf{z})$ whose distance should be minimized, where \mathbf{z} is the object nearest to \mathbf{x}_i among those with the same class label as \mathbf{x}_i in the training set, with the distance measured by the original Euclidean metric. For the parameters of ITML on UCI datasets, the default values in the authors' implementation were used, and for document and image datasets, we followed Jain et al. [5]. For DML-eig, the default setting in the authors' code was used for obtaining pairwise constraints. We calibrated the parameter k of k -NN classification to be used at the test time and all other parameters (γ in ITML, μ in DML-eig, and λ in Move-labeled and Move-query) by cross validation on the training set.

⁵ LMNN: <https://bitbucket.org/mlcircus/lmnn/downloads>,
 ITML: <http://www.cs.utexas.edu/~pjain/itml/>,
 DML-eig: <http://www.albany.edu/~yy298919/software.html>.

⁶ This code will be made available at our homepage.

Evaluation Criteria. The methods were evaluated in three respects: (i) the accuracy of k -NN classification using the distance/dissimilarity measure learned by each method, (ii) training time, and (iii) the degree of hubness in the data with respect to the learned distance/dissimilarity.

Following the literature [4, 7–10], we used the skewness of N_{10} distribution as the measure of hubness in the data. The N_{10} distribution is the empirical distribution of the number $N_{10}(i)$ of times each labeled object i is found in the 10-nearest neighbors of query (unlabeled) objects, and its skewness is defined as follows:

$$(N_{10} \text{ skewness}) = \frac{\sum_{i=1}^n (N_{10}(i) - E[N_{10}])^3 / n}{\text{Var}[N_{10}]^{\frac{3}{2}}}$$

where n is the total number of labeled objects, and $E[N_{10}]$ and $\text{Var}[N_{10}]$ are respectively the empirical mean and variance of $N_{10}(i)$ over n labeled objects. A large N_{10} skewness value indicates the existence of labeled objects that frequently appear in the 10-nearest neighbor lists of query objects, i.e., hubs.

5.2 Experimental Results and Discussion

Skewness. Table 2 shows the skewness of N_{10} distribution. For all datasets, we observe that the proposed approach (Move-labeled) reduced N_{10} skewness considerably compared with the original Euclidean distance, meaning that it effectively suppressed the emergence of hub objects. N_{10} skewness was reduced by metric learning methods (LMNN, ITML, and DML-eig) on many datasets, most notably by DML-eig. Also, as expected from the discussion of Sect. 4.2, Move-query increased N_{10} skewness except for the iris dataset.

Accuracy. Table 3 shows the classification accuracy. In most datasets, both the metric learning methods and the proposed method outperformed the original distance metric. The proposed method is comparable with, or slightly better than, the metric learning methods. Although Move-query optimized the minimizing distance between objects in same class (our proposed method also optimized such distance), the method obtained poor results even compared with the original Euclidean metric except for the iris datasets.

Note that, in UCI datasets, we observed that the proposed method did not work well, and even Move-query were competitive with others. This is an expected result, because the UCI datasets did not have much hubness even with the original metric (see Table 2a). Hubs tend to be emerge in high dimensional space [4, 7, 8], but all the UCI datasets have a small dimensionality (see Table 1a). Consequently, hub reduction/promotion methods did not affect the result significantly.

Table 2. Skewness of N_{10} distribution: a high skewness indicates the emergence of hubs (smaller is better). The bold figure indicates the best performer.

(a) UCI datasets.					
method	ionosphere	balance-scale	iris	wine	glass
original metric	1.65	0.93	0.40	0.71	0.77
LMNN	1.05	0.63	0.39	0.61	0.74
ITML	0.96	0.79	0.10	0.43	0.70
DML-eig	0.78	0.66	0.41	0.38	0.59
Move-labeled (proposed)	1.04	0.56	0.32	0.55	0.59
Move-query	1.67	1.13	0.32	0.89	1.18

(b) Document datasets.				
method	RCV	News	Reuters	TDT
original metric	13.35	21.93	7.61	4.89
LMNN	3.86	14.74	7.63	4.01
ITML	4.27	19.65	7.30	2.39
DML-eig	1.71	1.45	3.05	1.34
Move-labeled (proposed)	1.14	2.88	4.53	1.44
Move-query	21.57	33.36	17.49	6.71

(c) Image datasets.				
method	AwA	CUB	SUN	aPY
original metric	2.49	2.38	2.52	2.80
LMNN	3.10	2.96	2.80	3.94
ITML	2.42	2.27	2.37	2.69
DML-eig	1.90	1.77	2.39	2.17
Move-labeled (proposed)	1.24	0.97	1.02	1.23
Move-query	7.81	7.83	7.48	11.65

Training Time. To investigate the computational cost, we measured the elapsed real time needed to train the proposed method and the metric learning methods.

Table 4 shows the average training time in document and image datasets. We observe that the proposed approach has a clear advantage in terms of training cost. It was faster than any metric learning methods compared. Indeed, on all datasets except RCV, it was more than two orders of magnitude faster than the fastest metric learning methods. This can be explained by the fact that the metric learning methods take burden of optimizing over Mahalanobis metric. To enforce the constraint that the matrix \mathbf{M} in Eq. (2) should remain positive semi-definite, these methods pay high computational cost, e.g., to check the non-negativity of eigenvalues, at every training iteration. In contrast, the proposed approach has a closed-form solution; although this solution depends on matrix inverse, it needs to be calculated only once.

Table 3. Classification accuracy [%]: bold figures indicate the best performers for each dataset.

(a) UCI datasets.					
method	ionosphere	balance-scale	iris	wine	glass
original metric	86.8	89.5	97.2	98.1	68.1
LMNN	90.3	90.0	96.7	98.1	67.7
ITML	87.7	89.5	97.8	99.1	65.0
DML-eig	87.7	91.2	96.7	98.6	66.5
Move-labeled (proposed)	89.6	89.5	97.2	98.6	70.8
Move-query	79.7	89.4	97.2	96.3	62.3

(b) Document datasets.				
method	RCV	News	Reuters	TDT
original metric	92.1	76.9	89.5	96.1
LMNN	94.7	79.9	91.5	96.6
ITML	93.2	77.0	90.8	96.5
DML-eig	94.5	73.3	85.9	95.7
Move-labeled (proposed)	94.4	81.6	91.6	96.7
Move-query	89.1	70.0	85.9	95.4

(c) Image datasets.				
method	AwA	CUB	SUN	aPY
original metric	83.2	51.6	26.2	82.2
LMNN	83.0	54.7	24.4	81.8
ITML	83.1	51.3	26.0	82.4
DML-eig	82.0	53.5	22.4	81.6
Move-labeled (proposed)	84.1	52.4	28.3	83.4
Move-query	79.2	43.3	14.6	78.7

Table 4. Training time [sec]: bold figures indicate the best performer for each dataset.

(a) Document datasets.					(b) Image datasets.				
method	RCV	News	Reuters	TDT	method	AwA	CUB	SUN	aPY
LMNN	1713.0	1164.7	676.2	886.1	LMNN	1525.5	1098.2	15704.3	317.3
ITML	35.5	1512.5	124.1	169.0	ITML	1536.3	577.6	1126.4	9211.2
DML-eig	762.2	6145.9	2710.4	2350.6	DML-eig	2048.0	2084.7	2006.1	1787.1
proposed	6.0	7.0	4.6	16.1	proposed	9.5	1.5	4.1	6.4

6 Conclusion

In this paper, we have proposed a simple regression-based technique to improve k -NN classification accuracy. The results of our work can be summarized as follows:

- To improve k -NN classification accuracy, we proposed learning a transformation of labeled objects, without altering the coordinates of query (unlabeled) objects. This approach is justified from the perspective of reducing hubness in the labeled objects. Because our method is inherently designed to suppress hubness, it need not consider pairs of objects from different classes during training. The number of such pairs can be enormous and their use also renders the optimization problem non-convex, which is therefore a major obstacle to the scalability of metric learning methods.
- Our method deviates from the metric learning framework as the learned transformation \mathbf{W} does not provide a proper metric. In k -NN classification, however, labeled objects can be interpreted not as mere points but rather a representation of the (non-linear) decision boundaries between classes. Our approach follows this interpretation and changes the decision boundaries, through \mathbf{W} , to improve classification accuracy.
- In the experiments of Sect. 5, our approach indeed improved the classification accuracy when the data was high-dimensional and hubs emerged. It outperformed metric learning methods on most document and image datasets, and comparable on the rest. However, its effect was not evident on the UCI datasets, in which hubness was not noticeable because of the low dimensionality of the data.
- The experiments showed that our approach was substantially faster than the compared metric learning methods. For large document and image datasets, the speed-up was more than two orders of magnitude over the fastest metric learning methods, although the classification accuracy was better or comparable.

We have focused on multi-class classification problems in this paper, but hubness is known to be harmful in other situations, such as clustering and semi-supervised classification in high-dimensional space [7]. We plan to extend our approach to deal with these situations. We will also extend our method to learn nonlinear metrics.

Another direction of future work is to investigate the effect of our approach on kernel machines. Metric learning has been shown to be an effective preprocessing for kernel machines [3, 12, 14], and we will pursue a similar line using our approach.

Acknowledgments. We thank anonymous reviewers for their comments and suggestions. This work was partially supported by JSPS Kakenhi Grant No. 15H02749.

References

1. Bellet, A., Habrard, A., Sebban, M.: A survey on metric learning for feature vectors and structured data. arXiv preprint [arXiv:1306.6709](https://arxiv.org/abs/1306.6709) (2014)
2. Davis, J.V., Kulis, B., Jain, P., Sra, S., Dhillon, I.S.: Information-theoretic metric learning. In: ICML 2007, pp. 209–216 (2007)
3. Dhillon, P.S., Talukdar, P.P., Crammer, K.: Learning better data representation using inference-driven metric learning. In: ACL 2010, pp. 377–381 (2010)
4. Hara, K., Suzuki, I., Shimbo, M., Kobayashi, K., Fukumizu, K., Radovanović, M.: Localized centering: reducing hubness in large-sample data. In: AAAI 2015 (2015)
5. Jain, P., Kulis, B., Davis, J.V., Dhillon, I.S.: Metric and kernel learning using a linear transformation. JMLR **13**, 519–547 (2012)
6. Kulis, B.: Metric learning: a survey. Found. Trends Mach. Learn. **5**(4), 287–364 (2013)
7. Radovanović, M., Nanopoulos, A., Ivanović, M.: Hubs in space: popular nearest neighbors in high-dimensional data. JMLR **11**, 2487–2531 (2010)
8. Schnitzer, D., Flexer, A., Schedl, M., Widmer, G.: Local and global scaling reduce hubs in space. JMLR **13**, 2871–2902 (2012)
9. Shigeto, Y., Suzuki, I., Hara, K., Shimbo, M., Matsumoto, Y.: Ridge regression, hubness, and zero-shot learning. In: Appice, A., Rodrigues, P.P., Santos Costa, V., Soares, C., Gama, J., Jorge, A. (eds.) ECML PKDD 2015. LNCS (LNAI), vol. 9284, pp. 135–151. Springer, Cham (2015). doi:[10.1007/978-3-319-23528-8_9](https://doi.org/10.1007/978-3-319-23528-8_9)
10. Suzuki, I., Hara, K., Shimbo, M., Saerens, M., Fukumizu, K.: Centering similarity measures to reduce hubs. In: EMNLP 2013, pp. 613–623 (2013)
11. Weinberger, K.Q., Saul, L.K.: Distance metric learning for large margin nearest neighbor classification. JMLR **10**, 207–244 (2009)
12. Weinberger, K.Q., Tesauro, G.: Metric learning for kernel regression. In: AISTATS 2007, pp. 608–615 (2007)
13. Xing, E.P., Ng, A.Y., Jordan, M.I., Russell, S.: Distance metric learning, with application to clustering with side-information. In: NIPS 2002, pp. 505–512 (2002)
14. Xu, Z., Weinberger, K.Q., Chapelle, O.: Distance metric learning for kernel machines. arXiv preprint [arXiv:1208.3422](https://arxiv.org/abs/1208.3422) (2013)
15. Ying, Y., Li, P.: Distance metric learning with eigenvalue optimization. JMLR **13**, 1–26 (2012)

Advances in Knowledge Discovery and Data Mining
21st Pacific-Asia Conference, PAKDD 2017, Jeju, South
Korea, May 23-26, 2017, Proceedings, Part I

Kim, J.; Shim, K.; Cao, L.; Lee, J.-G.; Lin, X.; Moon, Y.-S.
(Eds.)

2017, XXXII, 841 p. 242 illus., Softcover

ISBN: 978-3-319-57453-0