

Maximin Separation Probability Clustering *

Gao Huang^{†‡}, Jianwen Zhang[§], Shiji Song^{†‡}, Zheng Chen[§]

[†] Tsinghua National Laboratory for Information Science and Technology (TNList)

[‡] Department of Automation, Tsinghua University, Beijing 100084, China

[§] Microsoft Research, Beijing

huang-g09@mails.tsinghua.edu.cn, jiazhan@microsoft.com

shijis@tsinghua.edu.cn, zhengc@microsoft.com

Abstract

This paper proposes a new approach for discriminative clustering. The intuition is, for a good clustering, one should be able to learn a classifier from the clustering labels with high generalization accuracy. Thus we define a novel metric to evaluate the quality of a clustering labeling, named *Minimum Separation Probability* (MSP), which is a lower bound of the generalization accuracy of a classifier learnt from the clustering labeling. We take MSP as the objective to maximize and propose our approach *Maximin Separation Probability Clustering* (MSPC), which has several attractive properties, such as invariance to anisotropic feature scaling and intuitive probabilistic explanation for clustering quality. We present three efficient optimization strategies for MSPC, and analyze their interesting connections to existing clustering approaches, such as maximum margin clustering (MMC) and discriminative k -means. Empirical results on real world data sets verify that MSP is a robust and effective clustering quality measure. It is also shown that the proposed algorithms compare favorably to state-of-the-art clustering algorithms in both accuracy and efficiency.

Introduction

Clustering is an unsupervised learning paradigm that aims to discover interesting groups in given data sets. It arises in a wide range of contexts, and various clustering algorithms have been proposed in the literature. In recent years, there is a new branch of research efforts on clustering, which connects the objective of clustering with that of classification. Some of the representative approaches include Maximum Margin Clustering (MMC) (Xu et al. 2004; Valizadegan and Jin 2007; Zhang, Tsang, and Kwok 2007; Li et al. 2009; Zhou et al.), Maximum Volume Clustering (MVC) (Niu et al. 2013) and linear discriminative analysis (LDA) based clustering (De la Torre and Kanade 2006; Ding and Li 2007). These methods focus on maximizing discrimination between clusters, thus are usually referred to as *discriminative clustering* (DC). Empirically these methods

have shown encouraging results in terms of clustering accuracy. However, there are still two important issues which are not well addressed.

The first issue is the problem of anisotropic feature scaling (AFS), which arises frequently in practice when features correspond to variables with different units. In contrast to supervised classification, clustering algorithms are more susceptible to AFS due to the absence of labels. Supposing we amplify the values of a certain feature by a large factor (e.g., change the unit of time from minute to second), then the maximal margin in MMC, or the maximal volume in MVC, will be dominated by this feature, and the clustering results are very likely to be affected. A common solution to this problem is preprocessing the data so that the values of each feature lie in a proper range. Although effective in practice, data preprocessing is generally performed independently of the clustering process, making it difficult to be optimized.

The second issue is the lack of a consistent measure for evaluating clustering quality. In supervised classification, we can leverage labeled training data to evaluate how confident we are in the learned model. However, it is a challenge to define a quality measure for clustering without true labels. Though the clustering measures of existing clustering algorithms, such as energy in k -means, maximum margin in MMC, entropy or mutual information in information-based clustering, can be used as measures for clustering quality, they are basically problem dependent. For example, some of these measures are not invariant to problem size or are sensitive to feature scaling.

This paper tries to address the aforementioned issues by resorting to the key intuition of discriminative clustering: if a clustering is good, then we should be able to learn a classifier from the clustering labels with high generalization accuracy. However, it is usually difficult to evaluate the generalization performance without true labels. In this paper, we define a novel clustering metric named *Minimum Separation Probability* (MSP) based on the Minimax Probability Machine (MPM) (Lanckriet et al. 2003), and propose a clustering approach by maximizing MSP, which is called *Maximin¹ Separation Probability Clustering* (MSPC). It can be

*This work was done during Gao Huang's internship at Microsoft Research.

Copyright © 2015, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

¹The term *minimax* is used in (Lanckriet et al. 2003) since it considers minimizing the maximum probability of misclassification.

shown that MSPC has close relations to MMC and several other DC approaches. However, MSPC is invariant to AFS, and its objective yields an intuitive explanation for clustering quality. Empirical results demonstrate that MSP has strong positive correlation with clustering accuracy. Additionally, we develop efficient learning algorithms for MSPC, leading to an increase in speed by several orders-of-magnitude over MMC and MVC algorithms. Meanwhile, MSPC achieves impressive clustering accuracy on real data sets.

Our contributions. 1) We introduce a novel clustering approach MSPC, the objective of which provides an effective and robust measure for clustering quality. 2) We propose three algorithms to optimize the proposed model, which are advantageous both in efficiency and clustering accuracy. 3) We study connections between MSPC and several existing algorithms. This also builds a bridge between two important types of DC approaches, namely, MMC and LDA based clustering.

Related Work

Discriminative clustering. Maximum Margin Clustering (MMC) (Xu et al. 2004) extended the large margin principle in SVM to clustering. There are several improvements to MMC, which can handle multi-class clustering or with higher efficiency (Zhang, Tsang, and Kwok 2007; Zhao, Wang, and Zhang 2008; Li et al. 2009). Similarly, the Maximum Volume Clustering (MVC) (Niu et al. 2013) extended the large volume principle to clustering. There are several other related discriminative clustering approaches, which simultaneously perform supervised dimension reduction and clustering (De la Torre and Kanade 2006; Ding and Li 2007; Ye, Zhao, and Wu 2007). We will further discuss the connections between these algorithms and MSPC in the rest sections.

Minimax Probability Machine (MPM) was proposed by Lanckriet et al. (2003) for supervised classification. It learns a classifier by maximizing a lower bound of the generalization accuracy. Recently, Huang et al. (2014) revisited MPM and extended it for transductive learning and semi-supervised learning based on an efficient label-switching strategy. Our paper can be viewed as extending MPM to the unsupervised case where we need to simultaneously learn a clustering labeling and a classifier.

Clustering Quality Measure (CQM) has been studied in the applied statistics literature (Milligan 1981), where it is termed as *cluster validity*. Recently, CQM has been axiomatized for pairwise distance-based clustering by Ackerman and Ben-David (2008). However, they do not connect CQM with the commonly used clustering accuracy. The unsupervised classification accuracy in (Balasubramanian, Donmez, and Lebanon 2011) can also be viewed as a CQM, but it assumes label proportion is known.

Minimax Probability Machine

Considering binary classification, there is a labeled training data set $\{\mathcal{X}, \mathcal{Y}\} = \{\mathbf{x}_i, y_i\}_{i=1}^N$ ($\mathbf{x}_i \in \mathcal{R}^d, y_i \in \{1, 2\}$), where N is the number of samples, and d is the dimensionality. The Minimax Probability Machine (MPM) (Lanckriet

et al. 2003) assumes that the samples of each class \mathcal{X}_j ($j = 1, 2$) are independently generated by a random distribution, thus the class-conditional distributions can be modeled by two random variables \mathbf{X}_1 and \mathbf{X}_2 , respectively. MPM aims to maximize the probabilities that \mathbf{X}_1 and \mathbf{X}_2 respectively lie on two sides of a hyperplane $\mathcal{H}(\mathbf{w}, b)$ ($\mathbf{w} \in \mathcal{R}^d, b \in \mathcal{R}$), i.e., $\Pr(\mathbf{w}^\top \mathbf{X}_1 + b \leq 0)$ and $\Pr(\mathbf{w}^\top \mathbf{X}_2 + b \geq 0)$. As these probabilities are usually difficult to compute, MPM considers maximizing the worst case separation probability over all possible distributions of \mathbf{X}_j whose mean and covariance $\{\mu_j, \Sigma_j\}$ match the empirical moments $\{\hat{\mu}_j, \hat{\Sigma}_j\}$ of \mathcal{X}_j , leading to the following MPM formulation:

$$\begin{aligned} \max_{p, \mathbf{w} \in \mathcal{R}^d, b \in \mathcal{R}} \quad & p \\ \text{s.t.} \quad & \inf_{\mathbf{X}_1 \sim \{\hat{\mu}_1, \hat{\Sigma}_1\}} \Pr(\mathbf{w}^\top \mathbf{X}_1 + b \geq 0) \geq p, \\ & \inf_{\mathbf{X}_2 \sim \{\hat{\mu}_2, \hat{\Sigma}_2\}} \Pr(\mathbf{w}^\top \mathbf{X}_2 + b \leq 0) \geq p, \end{aligned} \quad (1)$$

where $\hat{\mu}_k = \frac{1}{N_k} \sum_{\mathbf{x}_i \in \mathcal{X}_k} \mathbf{x}_i$, $\hat{\Sigma}_k = \frac{1}{N_k} \sum_{\mathbf{x}_i \in \mathcal{X}_k} (\mathbf{x}_i - \hat{\mu}_k)(\mathbf{x}_i - \hat{\mu}_k)^\top$ and $N_k = |\mathcal{X}_k|$, with $k = 1, 2$.

A remarkable advantage of considering the worst case distribution is that the constraints in (1) can be converted into second order cone constraints due to the multivariate Chebyshev inequalities (Marshall, Olkin, and others 1960), without making specific assumptions on the distribution form of \mathbf{X}_k . As derived in (Lanckriet et al. 2003), the optimization in (1) can be reformulated as

$$\max_{\mathbf{w} \in \mathcal{R}^d} \quad \kappa := \frac{\mathbf{w}^\top (\hat{\mu}_1 - \hat{\mu}_2)}{\sqrt{\mathbf{w}^\top \hat{\Sigma}_1 \mathbf{w}} + \sqrt{\mathbf{w}^\top \hat{\Sigma}_2 \mathbf{w}}}, \quad (2)$$

where $\kappa = \sqrt{p/(1-p)}$. This optimization problem can be solved efficiently by standard second order cone programming (SOCP), or by the iterative least squares approach introduced in (Lanckriet et al. 2003), which has a worst case complexity of $O(d^3)^2$.

Maximin Separation Probability Clustering

Inspired by MPM, we define a novel metric to measure the quality of a candidate clustering labeling.

Definition 1. The minimum separation probability (MSP) of a clustering labeling \mathbf{y} on \mathbf{X} is the optimal p solved from problem (1) on the pseudo-labeled training set $\{\mathbf{X}, \mathbf{y}\}$.

The intuition under this metric is, from a good clustering labeling, we should be able to learn a classifier with high generalization accuracy, which is bounded from below by MSP. Note here we are actually assuming that the clusters induced by an arbitrary labeling \mathbf{y} are consisted of i.i.d samples which are generated from independent distributions. This assumption is stronger than that in MPM, but is common in discriminative clustering.

The proposed clustering approach is to find a best clustering labeling by directly maximizing MSP:

$$\max_{\mathbf{w} \in \mathcal{R}^d, \mathbf{y} \in \{1, 2\}^N} \quad \kappa := \frac{|\mathbf{w}^\top (\hat{\mu}_1(\mathbf{y}) - \hat{\mu}_2(\mathbf{y}))|}{\sum_{k=1, 2} \sqrt{\mathbf{w}^\top \hat{\Sigma}_k(\mathbf{y}) + \lambda \Lambda} \mathbf{w}} \quad (3)$$

²This complexity does not take into account the cost of computing the empirical moments, which has a complexity of $O(Nd^2)$

where $\lambda\mathbf{\Lambda}$ is a regularization term on the covariance matrices. Though the two covariance matrices can be regularized separately with different regularizers, we use a unified one for convenience. In our algorithms, we let $\mathbf{\Lambda} = \text{diag}(\hat{\mathbf{\Sigma}})$, where $\hat{\mathbf{\Sigma}}$ is the covariance of the whole training set \mathcal{X} . Note that the empirical mean and covariance of both clusters are functions of the label vector \mathbf{y} .

We call the formulation (3) *Maximin Separation Probability Clustering* (MSPC). In the subsequent part, we analyze its properties, and introduce optimization methods to solve it.

Invariant to feature scaling The following theorem shows that MSP is invariant to invertible linear transformation on training data, and is thus immune to anisotropic feature scaling.

Theorem 1. *Let $\mathbf{P} \in \mathcal{R}^{d \times d}$ be an invertible matrix. Then the two data sets $\mathcal{X} = \{\mathbf{x}_i\}_{i=1}^N$ and $\mathcal{X}^{\mathbf{P}} = \{\mathbf{P}^T \mathbf{x}_i\}_{i=1}^N$ have the same optimal MSP.*

Due to limited spaces, all proofs appear in the supplementary material.

Since MSP is not pairwise distance-based CQM, we cannot analyze it under the axiomatic framework proposed by (Ackerman and Ben-David 2008). However, it can be shown that a generalized version of MSP satisfy all the axioms for a valid CQM. We give the detailed discussions in the supplementary material.

Optimization of MSPC

The mixed integer problem (3) is intractable in practice. In this section, we introduce three efficient algorithms to solve it approximately.

Approach 1: Iterative MPM

Inspired by the iterative support vector regression method (IterSVR) (Zhang, Tsang, and Kwok 2007) for MMC, we propose to solve MSPC in a similar fashion. Firstly, k -means clustering is adopted to find an initial labeling. Then we iterate between solving \mathbf{w} and updating \mathbf{y} as the IterSVR algorithm. Note that with fixed labeling, solving \mathbf{w} in MSPC is actually training an ordinary MPM model. We denote this algorithm by MSPC^{MPM} , which is summarized in the first column of Table 1.

This simple algorithm works surprisingly well in our experiments. It also converges fast, with typically 10 iterations before terminating. Moreover, using the updating formula (2) in Appendix B, at each iteration we only need to compute the covariance of a small set instances, i.e., those who are switched from one cluster to the other at this iteration.

Approach 2: Maximizing a lower bound

In this subsection, we derive a lower bound of MSP, enabling the design of an efficient MSPC algorithm. Moreover, this bound allows us to build connections between MSPC and some existing algorithms, as given in the next section.

Let $r_1 \triangleq N_1/N$ and $r_2 \triangleq N_2/N$ denote the ratios of samples assigned to subsets \mathcal{X}_1 and \mathcal{X}_2 , respectively. We then have the following theorem which gives a lower bound for the objective in (3).

Theorem 2. *Let $r_1 \vee r_2 \triangleq \max\{r_1, r_2\}$ and $r_1 \wedge r_2 \triangleq \min\{r_1, r_2\}$. Then we have*

$$\kappa^2 \geq \frac{(\mathbf{w}^T(\hat{\mu}_1 - \hat{\mu}_2))^2}{2(r_1 \wedge r_2)^{-1} \mathbf{w}^T(\hat{\mathbf{\Sigma}} + \lambda\mathbf{\Lambda})\mathbf{w} - 2(r_1 \vee r_2)(\mathbf{w}^T(\hat{\mu}_1 - \hat{\mu}_2))^2}, \quad (4)$$

and the equality holds when $\mathbf{w}^T \hat{\mathbf{\Sigma}}_1 \mathbf{w} = \mathbf{w}^T \hat{\mathbf{\Sigma}}_2 \mathbf{w}$ and $r_1 = r_2$.

In this expression, r_k and $\hat{\mu}_k$ ($k = 1, 2$) are functions of the label vector \mathbf{y} . When \mathbf{y} is given, maximizing the lower bound in (4) reduces to

$$\max_{\mathbf{w} \in \mathcal{R}^d} \frac{\mathbf{w}^T(\hat{\mu}_1 - \hat{\mu}_2)(\hat{\mu}_1 - \hat{\mu}_2)^T \mathbf{w}}{\mathbf{w}^T(\hat{\mathbf{\Sigma}} + \lambda\mathbf{\Lambda})\mathbf{w}}. \quad (5)$$

The optimal solution to (5) can be obtained efficiently by finding the largest eigenvector of the generalized eigenvalue problem (GEP): $(\hat{\mu}_1 - \hat{\mu}_2)(\hat{\mu}_1 - \hat{\mu}_2)^T \mathbf{v} = \gamma(\hat{\mathbf{\Sigma}} + \lambda\mathbf{\Lambda})\mathbf{v}$.

Thus we can again adopt alternating optimization scheme to optimize the lower bound of κ . Specifically, with fixed \mathbf{y} , the optimal \mathbf{w} is obtained by solving a GEP. With fixed \mathbf{w} , the optimal label vector can be obtained as follows³. We first sort the samples according to the values of their projections on \mathbf{w} , i.e., $t_i = \mathbf{w}^T \mathbf{x}_i$. Then we assign the first N_1 samples to \mathcal{X}_1 and the rest to \mathcal{X}_2 , where N_1 is an integer between 1 and $N - 1$ that maximizes the lower bound in (4). The algorithm is summarized in the second column of Table 1.

Approach 3: A relaxation method

We give a relaxation of the formulation in Approach 2. The relaxed model can be solved without alternating optimization, and is efficient for high dimensional data.

Define a cluster indicator vector \mathbf{q} as $q_i = \sqrt{N_2/(NN_1)}$, if $i \in \mathcal{X}_1$, and $q_i = -\sqrt{N_1/(NN_2)}$, if $i \in \mathcal{X}_2$.

Thus we have $\|\mathbf{q}\| = 1$, and $\mathbf{w}^T(\hat{\mu}_1 - \hat{\mu}_2) = \sqrt{N/(N_1N_2)} \mathbf{w}^T \mathbf{X}^T \mathbf{q}$.

Now we relax \mathbf{q} to take real values under the constraint $\|\mathbf{q}\| = 1$, and assume that N_1 and N_2 are constants (consequently, r_1 and r_2 are constants as well). The optimization problem in Approach 2 is approximated by

$$\max_{\mathbf{w} \in \mathcal{R}^d, \|\mathbf{q}\|=1} \frac{\mathbf{w}^T \mathbf{X}^T \mathbf{q} \mathbf{q}^T \mathbf{X} \mathbf{w}}{\mathbf{w}^T(\hat{\mathbf{\Sigma}} + \lambda\mathbf{\Lambda})\mathbf{w}}. \quad (6)$$

Without loss of generality, we assume that the data matrix are preprocessed to be centered, i.e., $\mathbf{1}^T \mathbf{X} = \mathbf{0}$. Then we show that the optimal \mathbf{q} can be obtained by solving an eigenproblem, as stated by the following theorem.

Theorem 3. *An optimal solution of \mathbf{q} to (6) is the largest normalized eigenvector of the matrix $\mathbf{I}_N - (\mathbf{I}_N + \tilde{\mathbf{X}})^{-1}$ (or $\mathbf{X}(\mathbf{X}^T \mathbf{X} + \lambda\mathbf{\Lambda})^{-1} \mathbf{X}^T$), where $\tilde{\mathbf{X}} = \mathbf{X}(\lambda\mathbf{\Lambda})^{-1} \mathbf{X}^T$ and \mathbf{I}_N is the $N \times N$ identity matrix.*

According to Theorem 3, the relaxed cluster indicator vector \mathbf{q} can be optimized conveniently by computing the largest eigenvector of a $N \times N$ matrix. When the optimal \mathbf{q}^* is obtained, we can use the sign of its elements to construct our label vector. Alternatively, we can use the predict vector $\mathbf{X}\mathbf{w}^* = \mathbf{X}\mathbf{X}^T \mathbf{q}^*$ to obtain our final prediction. In our experiments, the later one is used. We denote the algorithm as MSPC^{EIG} , which is presented in the last column of Table 1.

³It can be shown that this process finds the optimal \mathbf{y} with given \mathbf{w} , in $O(N \log N + Nd)$ time.

Analysis of the optimization algorithms

Among the three algorithms, MSPC^{MPM} is the only one that tries to directly optimize the MSP defined in (3), and it usually yields the highest accuracy in our experiments. It has a complexity of $O(l(Nd^2 + d^3))$, where l denotes the number of iterations, Nd^2 comes from the covariance estimation, and d^3 accounts for MPM retaining. Typically, we have $l = 10$, thus the algorithm is very efficient when d is not large.

MSPC^{GEP} works in a similar fashion as MSPC^{MPM} , but it aims to maximize a lower bound of MSP. Since the lower bound is tight when two clusters are balanced, MSPC^{GEP} tends to give similar results as MSPC^{MPM} on balanced data sets. It has a computational complexity of $O(l(Nd^2 + d^\alpha + N \log N))$, where d^α corresponds to solving the GEP in (5) with $2 < \alpha < 3$, and $N \log N$ accounts for the cost of sorting the projections.

MSPC^{EIG} has a complexity of $O(N^2d + N^3)$, if $N < d$, or $O(N^\alpha + d^3)$, if $N > d$ ($2 < \alpha < 3$). Thus it is more efficient than MSPC^{MPM} and MSPC^{GEP} on high dimensional data where $N \ll d$.

To summarize, MSPC^{MPM} is the first choice if accuracy is the main concern; MSPC^{GEP} is especially suited for balanced and medium-dimensional data; MSPC^{EIG} is the most efficient one on high dimensional data.

Connections with Existing Algorithms

Relationships to MMC

We show that interesting connections exist between MSPC and MMC. Actually, MMC solutions ensures high MSP under certain conditions, and MSPC essentially favors a large *relative margin*.

Theorem 4. *Supposing that the optimal margin learned by MMC is δ , then the MSP (with $\lambda = 0$) is at least $p \geq \frac{1}{2(r_1 \wedge r_2)^{-1}(\sigma_{\max}/\delta)^2 - 2(r_1 \vee r_2) + 1}$, where σ_{\max}^2 is the largest eigenvalue of the covariance matrix $\hat{\Sigma}$, and r_1 and r_2 are the ratios of samples assigned to the two clusters, respectively.*

Theorem 4 states that to ensure high MSP, the *normalized margin* δ/σ_{\max} , instead of the absolute margin δ , should be large. For any data set, we could create an arbitrarily large margin by feature scaling, however, the normalized margin cannot be made arbitrary large since σ_{\max} is also sensitive to feature scaling. This theorem also explains why data preprocessing is important to MMC: it prevents σ_{\max} from growing too large, so that the subsequent margin maximization process could reliably maximize intrinsic separability.

Theorem 5. *Supposing that the two clusters yielded by MSPC are respectively generated from two independent distributions with mean μ_1 and μ_2 , and with a MSP ($\lambda = 0$) of p , then there exists a hyperplane $\{\mathbf{w}, b\}$ such that at least $1 - r$ fraction of the points from each distribution satisfy $|\mathbf{w}^\top \mathbf{x}_i + b| \geq \rho\Delta/2$ ($0 \leq \rho \leq 1$), where $\Delta = |\mathbf{w}^\top (\mu_1 - \mu_2)|$, and $r = \frac{p^{-1}-1}{(p^{-1}-1)+(1-\rho)^2/4}$.*

In MMC, we are usually maximizing a *soft* margin, by pushing *most* of the points away from the margin. From Theorem 5, we see that MSPC also ensures a large fraction (e.g., at least $1 - r$) of the samples are separated by a margin $\rho\Delta$, where Δ is actually the *average margin* between two clusters. If we define the $|\mathbf{w}^\top \mathbf{x}_i + b|/\Delta$ as the *relative margin*, then a high MSP solution ensures a large relative margin with high probability. Unlike the absolute margin in MMC, the relative margin is invariant to feature scaling.

Relationships to LDA-based clustering approaches

Several recent work (De la Torre and Kanade 2006; Ding and Li 2007; Ye, Zhao, and Wu 2007) addressed the advantages of integrating linear discriminant analysis (LDA) and k -means clustering. It can be shown that MSPC is related to these discriminative clustering approaches.

The objective function (binary clustering case) introduced in (De la Torre and Kanade 2006) is given by

$$\max_{\mathbf{w} \in \mathcal{R}^d, \mathbf{y} \in \{1,2\}^N} (\mathbf{w}^\top \mathbf{S} \mathbf{w})^{-1} \mathbf{w}^\top \mathbf{S}_b(\mathbf{y}) \mathbf{w}, \quad (7)$$

where \mathbf{S} is the *total scatter matrix*, and \mathbf{S}_b is the *between cluster scatter matrix*, which is a function of the label vector \mathbf{y} .

It can be observed that for binary clustering, \mathbf{S} and \mathbf{S}_b are respectively proportional to $\hat{\Sigma}$ and $(\hat{\mu}_1 - \hat{\mu}_2)(\hat{\mu}_1 - \hat{\mu}_2)^\top$ in our algorithm. Therefore, the two expressions in (4) and (7) are actually equivalent, if we omit the ratios r_j and the regularization term in the former. Thus the algorithm in (De la Torre and Kanade 2006) essentially maximizes an approximated lower bound of MSP (the objective of MSPC^{EIG}).

The objective function in (Ding and Li 2007) is identical to that in (De la Torre and Kanade 2006), except the *total scatter matrix* \mathbf{S} is replaced by the *within cluster scatter matrix* \mathbf{S}_w . Since we have the relation $\mathbf{S}_t = \mathbf{S}_w + \mathbf{S}_b$, the two objective are essentially equivalent. Thus the algorithm in (Ding and Li 2007) also maximize an approximated lower bound of MSP, but with a different optimization scheme.

The discriminative k -means algorithm proposed in (Ye, Zhao, and Wu 2007) has a similar objective as that in (Ding and Li 2007). It performs kernel k -means clustering with the kernel matrix $\mathbf{I}_N - (\mathbf{I}_N + (\lambda)^{-1} \mathbf{X} \mathbf{X}^\top)^{-1}$. Obviously, this kernel matrix can be obtained from Theorem 3 by setting the regularization matrix $\mathbf{\Lambda}$ to \mathbf{I}_N . Though the difference seems trivial, it is critical since the regularizer $\mathbf{\Lambda}$ preserves the property of invariant to AFS in our algorithm.

Empirical Study

Experiment setup All algorithms are implemented in MATLABTM, and are executed on an Intel i7 Quad Core CPU 3.39GHz machine with 16GB RAM. All of data sets are obtained from the UCI repository, except the handwritten digits data sets *MNIST*⁴ and *USPS*. For those data sets originally contain more than two classes, we select their first two classes to create binary clustering tasks, if not explicitly

⁴Available at <http://yann.lecun.com/exdb/mnist/>

MSPC ^{MPM}	MSPC ^{GEP}	MSPC ^{EIG}
1. Initialize \mathbf{y} by k -means; 2. Fix \mathbf{y} and optimize $\{\mathbf{w}, b\}$ with MPM; 3. Update the label vector \mathbf{y} according to $\text{sign}(\mathbf{w}^\top \mathbf{x}_i + b)$; 4. Repeat Step 2 and 3 until converge.	1. Initialize \mathbf{y} by k -means; 2. Fix \mathbf{y} and optimize \mathbf{w} by solving (5); 3. Fix \mathbf{w} , and find the optimal \mathbf{y} that maximizes the lower bound of κ ; 4. Repeat Step 2 and 3 until converge.	1. Calculate $\mathbf{I}_N - (\mathbf{I}_N + \tilde{\mathbf{X}})^{-1}$ or $\mathbf{X}(\mathbf{X}^\top \mathbf{X} + \lambda \mathbf{I})^{-1} \mathbf{X}^\top$; 2. Compute the largest eigenvector \mathbf{q}^* ; 3. Obtaining the optimal label vector via the sign of $\mathbf{X} \mathbf{X}^\top \mathbf{q}^*$.

Table 1: Summary of the proposed three MSPC algorithms.

ID	Data	N	d	KM	Iter-SVR	LG-MMC	MVC	LDA-KM	Dis-KM	MSPC-MPM	MSPC-GEP	MSPC-EIG
1	<i>ionosphere</i>	351	34	28.77	22.51	24.70	15.67	28.77	17.95	28.77	29.63	29.63
2	<i>breast</i>	683	10	3.81	3.22	3.51	2.93	3.81	3.81	2.93	2.63	2.78
3	<i>australian</i>	690	14	14.49	14.06	14.06	17.25	14.49	14.49	14.49	14.06	16.67
4	<i>diabetes</i>	768	8	33.07	30.86	32.68	29.63	32.94	28.26	32.55	31.51	31.77
5	<i>letter</i>	1555	16	6.30	5.53	0.00	5.66	5.53	5.59	5.59	5.53	8.75
6	<i>satellite</i>	2236	36	4.25	6.17	0.76	0.85	4.07	4.20	0.63	3.80	1.70
7	<i>spam</i>	4601	57	20.04	20.98	18.30	N/A	20.07	19.17	13.76	17.19	21.75
8	<i>mnist3vs8</i>	13966	784	20.04	20.06	18.12	N/A	18.78	18.78	17.85	18.54	23.19
9	<i>mnist1vs7</i>	15170	784	4.23	4.23	2.30	N/A	2.52	1.54	0.93	1.10	3.70
10	<i>mnist8vs9</i>	13783	784	7.34	40.16	7.19	N/A	5.12	4.43	4.25	4.67	18.50

Table 2: Clustering error (%) on ten clustering tasks.

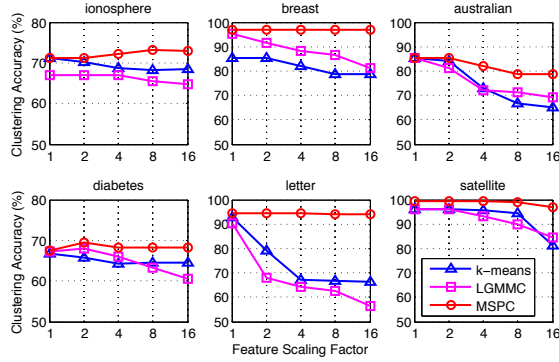


Figure 1: Clustering accuracy (%) with respect to anisotropic feature scaling.

stated. For comparisons to be fair, we normalize each feature of the UCI data sets to the range $[-1, 1]^d$.

Baselines and parameter setting The baseline clustering algorithms include two relatively efficient MMC algorithms, IterSVR (Zhang, Tsang, and Kwok 2007) and LGMMC (Li et al. 2009); MVC (Niu et al. 2013), and two related DC algorithms, LDA-KM (Ding and Li 2007) and Dis-KM (Ye, Zhao, and Wu 2007). For IterSVR, LGMMC and MVC, both the regularization parameter C and the Gaussian kernel parameter σ are tuned (linear kernel LGMMC is used for the three *mnist* data sets), following the settings in (Li et al. 2009) and (Niu et al. 2013), i.e., the hyperparameters correspond to highest clustering accuracy are directly selected. For LDA-KM, Dis-KM and the proposed three MSPC algorithms, the regularization coefficient, which is the only hyperparameter, is optimally chosen from the candidate set $[10^{-4}, 10^{-3}, \dots, 10^4]$ based on the clustering accuracy.

⁵ Actually, the proposed algorithms performed well on both normalized and unnormalized data, while some baselines degraded significantly without feature normalizing.

Clustering Accuracy and efficiency

We report the clustering error on seven UCI data sets and three binary clustering tasks in Table 2, where the best results are shown in bold. An "N/A" indicates an algorithm that fails to finish in a reasonable time (24 hours). It can be observed that MSPC^{MPM}, which directly optimizes the MSP, achieves very encouraging results. It yields the lowest clustering error on the five relatively larger data sets. This validates MSP as a proper clustering measure. In accordance with our analysis in previous sections, the results obtained by MSPC^{GEP} are similar to that of MSPC^{MPM} on the balanced data sets, e.g., *letter* and the three binary *mnist* tasks. However, the clustering error given by MSPC^{EIG} is relatively higher than that obtained by the other two MSPC algorithms.

In Table 3, we report the running time of the DC algorithms on six relatively larger data sets. From the results, one can observe that the proposed three algorithms are several orders of magnitude faster than MMC and MVC. MSPC^{MPM} is the most efficient algorithm on low dimensional data sets, while it is surpassed by MSPC^{GEP} on *mnist* tasks which have a relatively larger dimensionality. We attribute the high efficiency of MSPC^{MPM} and MSPC^{GEP} to two main reasons: 1) solving for \mathbf{w} is independent of the number of training data when the means and covariances are given, thus it is efficient when d is not large, and 2) the covariances can be updated efficiently by using the formula (2) in Appendix B (the means can be updated similarly). Although some baselines, such as Iter-SVR, also have a linear complexity in N , there is a large constant hidden in it. Note that for extremely high dimensional data with $d \gg N$, we can use MSPC^{EIG}, or use MSPC^{MPM} with kernel MPM introduced in (Lanckriet et al. 2003).

Sensitivity to anisotropic feature scaling

In order to verify that MSPC is insensitive to AFS, we run MSPC^{MPM} on six deliberately un-uniformly scaled data sets. At each time, we randomly select one of the features from

ID	Data	Iter-SVR	LG-MMC	MVC	LDA-KM	Dis-KM	MSPC ^{MPM}	MSPC ^{GEP}	MSPC ^{EIG}
5	<i>letter</i>	50.2	45.7	178	0.07	0.38	0.02	0.23	0.45
6	<i>satellite</i>	9.16	140	603	0.10	1.05	0.02	0.27	0.95
7	<i>spam</i>	42.6	1769	N/A	0.52	5.01	0.09	0.10	5.63
8	<i>mnist3vs8</i>	6372	1824	N/A	5.53	96.8	2.82	2.72	9.25
9	<i>mnist1vs7</i>	3602	2602	N/A	3.01	184	1.65	1.54	16.6
10	<i>mnist8vs9</i>	3504	2990	N/A	4.01	109	2.40	1.83	9.88

Table 3: Wall clock time (in seconds) of the clustering algorithms.




Figure 2: MSP (the circles) and CA (the bars) on three groups of data sets. From left to right: the 10 tasks presented in Table 2; the 21 binary clustering tasks created from *Satellite*; and the 45 binary clustering tasks created from *USPS* (the tasks are ordered by “0vs1”, “0vs2”, ..., “0vs9”, “1vs2”, ..., “8vs9”).




Figure 3: Correlation between CA and MSP.

a data set, and scale its values by a factor of a , where a is chosen from the set $[2, 4, 8, 16]$. Then we run k -means, LGMMC with linear kernel, and MSPC^{MPM} on it. This experiment is repeated 10 times (we ensure 10 different features are scaled, and for *diabetes* all the eight features are selected and scaled once) on each data set, and the average results obtained by the three algorithms are shown in Fig.1. It is obvious that MSPC^{MPM} is robust to such feature scaling, while the other two algorithms degrade significantly with a large scaling factor. Note that MSPC^{MPM} is a local search algorithm and it is initialized with k -means, thus it is not fully immune to AFS. If we can afford to find the global optimum, then MSPC solutions will totally be unaffected by AFS, as guaranteed by Theorem 1.

MSP and clustering accuracy

Estimating clustering accuracy (CA) without ground truth labels is a very challenging task. Since MSP is a problem-independent clustering quality measure, and it has the same value range as CA $([0,1])$, it is interesting to see how they correlate with each other empirically. Fig.2 and Fig.3 show the CA and corresponding MSP obtained by MSPC^{MPM} on the ten tasks presented in Table 2, the 21 binary clustering tasks created from *Satellite*, and the 45 binary clustering tasks created from *USPS*. One can observe that these two variables do correlate with each other to some extent. For example, the correlation coefficient on the 10 UCI and

MNIST data sets is 0.72, indicating a moderate positive linear relationship between them. If we only consider the later five larger data sets, then the correlation coefficient reaches 0.97, which means an almost perfect linear correlation. Interestingly, the MSP indicates the separability of the three *MNIST* tasks can be ordered as “1vs7” > “8vs9” > “3vs8”, which exactly matches the order given by CA. The results on the binary clustering tasks of *Satellite* and *USPS* also show a high correlation between MSP and CA, with coefficients of 0.86 and 0.81 respectively (shown in Fig.3).

Conclusion

In this paper, we proposed a new clustering metric, minimum separation probability (MSP), which is a lower bound of the generalization accuracy of a classifier learnt from the clustering labeling. Then we proposed three clustering approaches by approximately maximizing MSP. We provided a detailed analysis of the relations between MSPC and several existing discriminative clustering approaches. Empirical results demonstrated that MSP is insensitive to anisotropic feature scaling, and it showed a positive correlation with clustering accuracy on real data sets. The proposed algorithms were impressively faster than MMC and MVC algorithms, and compared favorably to state-of-the-art clustering algorithms in terms of accuracy. Future work could focus on further studying the relations between MSP and clustering accuracy, and extending MSPC for multi-class clustering.

Acknowledgements. Gao Huang and Shiji Song are partially supported by the National Natural Science Foundation of China under Grant Nos. 41427806 and 61273233, the Research Fund for the Doctoral Program of Higher Education under Grant Nos. 20120002110035 and 20130002130010, the Project of China Ocean Association under Grant DY125-25-02, and Tsinghua University Initiative Scientific Research Program under Grant No. 20131089300.

References

- Ackerman, M., and Ben-David, S. 2008. Measures of clustering quality: A working set of axioms for clustering. In *NIPS*.
- Balasubramanian, K.; Donmez, P.; and Lebanon, G. 2011. Unsupervised supervised learning II: Margin-based classification without labels. *The Journal of Machine Learning Research* 12:3119–3145.
- De la Torre, F., and Kanade, T. 2006. Discriminative cluster analysis. In *ICML*.
- Ding, C., and Li, T. 2007. Adaptive dimension reduction using discriminant analysis and k -means clustering. In *ICML*.
- Huang, G.; Song, S.; Xu, Z.; and Weinberger, K. 2014. Transductive minimax probability machine. In *ECML/PKDD*, 579–594.
- Lanckriet, G. R.; Ghaoui, L. E.; Bhattacharyya, C.; and Jordan, M. I. 2003. A robust minimax approach to classification. *The Journal of Machine Learning Research* 3:555–582.
- Li, Y.-F.; Tsang, I. W.; Kwok, J. T.; and Zhou, Z.-H. 2009. Tighter and convex maximum margin clustering. In *AIS-TATS*.
- Marshall, A. W.; Olkin, I.; et al. 1960. Multivariate chebyshev inequalities. *The Annals of Mathematical Statistics* 31(4):1001–1014.
- Milligan, G. W. 1981. A review of monte carlo tests of cluster analysis. *Multivariate Behavioral Research* 16(3):379–407.
- Niu, G.; Dai, B.; Shang, L.; and Sugiyama, M. 2013. Maximum volume clustering: A new discriminative clustering approach. *The Journal of Machine Learning Research* 14(1):2641–2687.
- Valizadegan, H., and Jin, R. 2007. Generalized maximum margin clustering and unsupervised kernel learning. In *NIPS*.
- Xu, L.; Neufeld, J.; Larson, B.; and Schuurmans, D. 2004. Maximum margin clustering. In *NIPS*.
- Ye, J.; Zhao, Z.; and Wu, M. 2007. Discriminative k -means for clustering. In *NIPS*.
- Zhang, K.; Tsang, I. W.; and Kwok, J. T. 2007. Maximum margin clustering made practical. In *ICML*.
- Zhao, B.; Wang, F.; and Zhang, C. 2008. Efficient multiclass maximum margin clustering. In *ICML*.
- Zhou, G.-T.; Lan, T.; Vahdat, A.; and Mori, G. Latent maximum margin clustering. In *NIPS*.