

# Toward Optimal Feature Selection in Naive Bayes for Text Categorization

Bo Tang, *Student Member, IEEE*, Steven Kay, *Fellow, IEEE*, and Haibo He, *Senior Member, IEEE*

**Abstract**—Automated feature selection is important for text categorization to reduce the feature size and to speed up the learning process of classifiers. In this paper, we present a novel and efficient feature selection framework based on the Information Theory, which aims to rank the features with their discriminative capacity for classification. We first revisit two information measures: Kullback-Leibler divergence and Jeffreys divergence for binary hypothesis testing, and analyze their asymptotic properties relating to type I and type II errors of a Bayesian classifier. We then introduce a new divergence measure, called Jeffreys-Multi-Hypothesis (JMH) divergence, to measure multi-distribution divergence for multi-class classification. Based on the JMH-divergence, we develop two efficient feature selection methods, termed maximum discrimination ( $MD$ ) and  $MD - \chi^2$  methods, for text categorization. The promising results of extensive experiments demonstrate the effectiveness of the proposed approaches.

**Index Terms**—Feature selection, text categorization, Kullback-Leibler divergence, Jeffreys divergence, Jeffreys-Multi-Hypothesis divergence



## 1 INTRODUCTION

WITH the increasing availability of text documents in electronic form, it is of great importance to label the contents with a predefined set of thematic categories in an automatic way, what is also known as automated *Text Categorization*. In last decades, a growing number of advanced machine learning algorithms have been developed to address this challenging task by formulating it as a classification problem [1] [2] [3] [4] [5]. Commonly, an automatic text classifier is built with a learning process from a set of prelabeled documents.

Documents need to be represented in a way that is suitable for a general learning process. The most widely used representation is “the bag of words”: a document is represented by a vector of features, each of which corresponds to a term or a phrase in a vocabulary collected from a particular data set. The value of each feature element represents the importance of the term in the document, according to a specific feature measurement.

A big challenge in text categorization is the learning from high dimensional data. On one hand, tens and hundreds of thousands terms in a document may lead to a high computational burden for the learning process. On the other hand, some irrelevant and redundant features may hurt predictive performance of classifiers for text categorization. To avoid the issue of the “curse of dimensionality” and to speed up the learning process, it is necessary to perform feature reduction to reduce feature size.

A common feature reduction approach for text categorization is feature selection that this paper concentrates on, where only a subset of original features are selected as input to the learning algorithms. In last decades, a number of feature selection methods have been proposed, which can be usually categorized into the following two types of approach: the *filter* approach and the *wrapper* approach [6]. The filter approach selects feature subsets based on the general characteristics of the data without involving the learning algorithms that will use the selected features. A score indicating the “importance” of the term is assigned to each individual feature based on an independent evaluation criterion, such as distance measure, entropy measure, dependency measure and consistency measure. Hence, the filter approach only selects a number of top ranked features and ignores the rest. Alternatively, the wrapper approach greedily searches for better features with an evaluation criterion based on the same learning algorithm. Although it has been shown that the wrapper approach usually performs better than the filter approach, it has much more computational cost than the filter approach, which sometimes makes it impractical.

Typically, the filter approach is predominantly used in text categorization because of its simplicity and efficiency. However, the filter approach evaluates the goodness of a feature by only exploiting the intrinsic characteristics of the training data without considering the learning algorithm for discrimination, which may lead to an undesired classification performance. Given a specific learning algorithm, it is hard to determine which filter feature selection approach is the best for discrimination.

In this paper, we present a feature selection method

• Bo Tang, Steven Kay, and Haibo He are with the Department of Electrical, Computer and Biomedical Engineering at the University of Rhode Island, Kingston, RI, 02881. E-mail: {btang, kay, he}@ele.uri.edu

which ranks the original features, aiming to maximize the discriminative performance for text categorization, when naive Bayes classifiers are used as learning algorithms. Unlike the existing filter approaches, our method evaluates the goodness of a feature without training a classifier explicitly, and selects these features that offer maximum discrimination in terms of a new divergence measure. Specifically, the contributions of this paper are as follows:

1. We introduce a new divergence measure for multi-class classification by extending the J-divergence measure, termed Jeffreys-Multi-Hypothesis divergence (JMH-divergence).
2. We propose an efficient approach to rank the order of features to approximately produce the maximum JMH divergence. The theoretical analysis shows that the JMH divergence is monotonically increasing when more features are selected.
3. We analyze the asymptotic distribution of the proposed test statistic, which leads to the  $\chi^2$  distribution. By doing so, we introduce another simple and effective feature ranking approach by maximizing the noncentrality measurement of the noncentral  $\chi^2$  distribution.

The rest of the paper is organized as follows: In Section 2, we introduce the previous work on naive Bayes classifiers and feature selection techniques for automatic text categorization. In Section 3, we present the theoretical framework of feature selection using the information measures. In Section 4, we introduce two efficient feature selection approaches in naive Bayes classifiers for text categorization. Experimental results are given in Section 5 along with the performance analysis compared with the state-of-the-art methods. A conclusion and future work discussion are given in Section 6.

## 2 PREVIOUS WORK

### 2.1 Document Representation

In text categorization, each document is commonly represented by the model of the “bag-of-words” with a  $M \times 1$  feature vector  $D : \mathbf{d} = [x_1, x_2, \dots, x_M]^T$ . The  $i$ -th element  $x_i$  corresponds to the measure of the  $i$ -th term (“word”) in a vocabulary or a dictionary. For a given data set, we first generate the vocabulary with a set of  $M$  unique terms from all documents. Then, for each document, a feature vector can be formed by using various feature models. Typically, the value of a feature represents the information about this particular term in a document. Two feature models have been widely used. The first one is the binary feature model in which the feature takes value either 0 or 1 corresponding to the presence or the absence of a particular term in the document. The distribution of such binary feature for each class can be usually modeled by a Bernoulli distribution. The other one is multi-value feature model in which the

feature takes values in  $\{0, 1, \dots\}$  corresponding to the number of occurrences of a particular term in the document, and thus it is also called term frequency (TF). The distribution of TF for each class can be usually modeled by a multinomial distribution model. We note here that several other feature models also exist in literature, such as normalized term frequency and inverse document frequency (tf-idf) [7] and probabilistic structure representation [8]. Recent work in learning vector representations of words using neural networks have shown superior performance in classification and clustering [9] [10] [11] [12], where both the ordering and semantics of the words are considered.

### 2.2 Naive Bayes

Naive Bayes classifier has been widely used for text categorization due to its simplicity and efficiency [13] [14] [15]. It is a model-based classification method and offers competitive classification performance for text categorization compared with other data-driven classification methods [16] [17] [18] [19], such as neural network, support vector machine (SVM), logistic regression, and k-nearest neighbors. The naive Bayes applies the Bayes’ theorem with the “naive” assumption that any pair of features are independent for a given class. The classification decision is made based upon the *maximum-a-posteriori* (MAP) rule. Usually, three distribution models, including Bernoulli model, multinomial model and Poisson model, have commonly been incorporated into the Bayesian framework and have resulted in classifiers of Bernoulli naive Bayes (BNB), multinomial naive Bayes (MNB) and Poisson naive Bayes (PNB), respectively.

Extensive experiments on real-life benchmarks have shown that the MNB usually outperforms the BNB at large vocabulary size [20]. Similar conclusions have been drawn in [3] [21] [22]. It has also been shown that the PNB is equivalent to the MNB, if the document length and document class are assumed to be independent [21]. For this reason, the naive Bayes usually refers to the MNB classifier. In this paper, we concentrate on the formulation of the proposed feature selection method for the MNB classifier. The methods can be easily extended to BNB and PNB classifiers.

MNB would be one of the best-known naive Bayes classification approaches using the term frequency to represent the document. Considering a text categorization problem with  $N$  classes (“topics”), let  $c$  be the discrete variable of class label taking values in  $\{1, 2, \dots, N\}$ , and  $\mathbf{x}$  be the integer-valued feature vector corresponding to the term frequency. The MNB classifier assumes that the number of times that each term appears in the document satisfies a multinomial distribution [13] [20]. In other words, a document with  $l$  terms is considered as  $l$  independent trials, and each

term is the result of a trial exactly falling into the vocabulary. Let the vocabulary size be  $M$ , and then each documents is represented by a  $M \times 1$  feature vector.

Hence, given a document  $D$ , we first count the number of times that each term appears and generate a feature vector  $\mathbf{d} = [x_1, x_2, \dots, x_M]^T$ . According to the multinomial distribution, the likelihood of observing  $\mathbf{d}$  conditioned on the class label  $c$  and the document length  $l$  can be calculated as follows:

$$p(\mathbf{d}|c, l; \theta_c^m) = \frac{l}{x_1! x_2! \dots x_M!} \prod_{i=1}^M p_{ic}^{x_i} \quad (1)$$

where  $l = \sum_{m=1}^M x_m$ , and  $\theta_c^m = [p_{1c}, p_{2c}, \dots, p_{Mc}]^T$  is a parameter vector, each of which denotes the probability that the term of a trial falls into one of  $M$  categories, so that  $p_{ic} \geq 0$  for  $i = 1, 2, \dots, M$  and  $p_{1c} + p_{2c} + \dots + p_{Mc} = 1$ . Using the MLE method, the estimate of each term probability  $\hat{p}_{ic}$  is given by

$$\hat{p}_{ic} = \frac{l_{ic}}{l_c} \quad (2)$$

where  $l_{ic}$  is the number of times the  $i$ -th term appears among documents in class  $c$ , and  $l_c$  is the total number of terms in class  $c$ . To avoid the *zero probability* issue, the technique of “Laplace smoothing” or the prior information for probability  $p_{ic}$  is usually applied [23], which leads to

$$\hat{p}_{ic} = \frac{l_{ic} + \beta_1}{l_c + \beta_2} \quad (3)$$

where  $\beta_1$  and  $\beta_2$  are the constant smoothing parameters. Using the Laplacean prior, we prime each word’s count with a count of one and have  $\beta_1 = 1$  and  $\beta_2 = M$  [20] [21].

Note that the document length  $l$  is commonly assumed to be independent on the document class to simplify the naive Bayes classification rule, that is, we have  $p(\mathbf{d}|c, l) = p(\mathbf{d}|c)$  and  $p(c|\mathbf{d}, l) = p(c|\mathbf{d})$  for likelihood and posterior probability, respectively. Otherwise, it leads to a more general formulation of the posterior probability for classification, which is given by

$$\begin{aligned} p(c|\mathbf{d}, l) &= p(\mathbf{d}, l|c)p(c) \\ &= p(\mathbf{d}|c, l)p(l|c)p(c) \end{aligned} \quad (4)$$

where the class-wise document length information  $p(l|c)$  is incorporated in a Bayesian fashion. The document length information sometimes may be useful for making classification decisions, e.g., when the class-wise document length distributions are different, but it requires the estimate of  $p(l|c)$  for a given data set. We will follow the common assumption that the document length is constant in our experiments, i.e.,  $p(\mathbf{d}|c) = p(\mathbf{d}|c, l)$ , and note that the solutions of our feature selection methods do not change with and without this assumption.

In naive Bayes classifiers, given a new document  $D_t$  to be classified, we first generate its feature vector  $\mathbf{d}_t$  and apply the following decision rule to make a classification:

$$\begin{aligned} c^* &= \arg \max_{c \in \{1, 2, \dots, N\}} p(c|\mathbf{d}_t; \theta_c) \\ &\propto \arg \max_{c \in \{1, 2, \dots, N\}} p(\mathbf{d}_t|c; \theta_c)p(c) \\ &\propto \arg \max_{c \in \{1, 2, \dots, N\}} \log p(\mathbf{d}_t|c; \theta_c) + \log p(c) \end{aligned} \quad (5)$$

The likelihood probability  $p(\mathbf{d}_t|c; \theta_c)$  could be a specific model-based distribution. Here, for MNB classifier, the probability distributions  $p(\mathbf{d}_t|c; \theta_c)$  for  $i = 1, 2, \dots, N$ , are the multinomial distributions given by Eq. (1).

### 2.3 Feature Selection

Feature selection, also called term selection, is a widely adopted approach for dimensionality reduction in text categorization. Given a predetermined integer  $r$ , the number of terms to be selected, the feature selection approach attempts to select  $r$  out of  $M$  terms in the original vocabulary. Yang and Pedersen [17] have shown that the employment of feature selection can remove 98% unique terms without hurting the classification performance too much, and thus feature selection can greatly reduce the computational burden for classification.

In [24], Kohavi and John have used a so-called *wrapper* feature selection approach in which a feature is either added or removed at one step towards the optimal feature subset selection. When a new feature set is generated, the classifier is re-trained with new features and it is further tested on a validation set. This approach has the advantage of searching the feature space in a greedy way and is always able to find a better feature subset in a sense of an improved classification performance. However, the high computational cost makes it prohibitive for practical text categorization applications.

The alternative and popular one is the *filter* approach, in which each feature is assigned with a score based on its importance measure and only the top ranked features with the highest scores are kept. The advantage of this approach is its easy implementation with low computational cost. In the rest of this section, we describe several state-of-the-art measures that are widely used in text categorization.

*Document frequency*, the number of documents in which a term occurs is a simple and effective feature selection approach. It removes from the original feature space the rare terms that are considered as non-informative for classification. The effectiveness of this approach also inspired researchers in their experiments to remove all the terms that occur no more than  $a$  times ( $a$  usually ranges from 1 to 3) in the training set as a preprocessing stage. By doing

so, tens of hundreds rare features can be removed before the step of feature selection. *TF-IDF* measure considers both term frequency and inverse document frequency to calculate the importance of features [25]. More recently, a generalized TF-IDF measure [26] is proposed by considering different level of hierarchies among words to effectively analyze tweet user behaviors.

Many other filter approaches are based on the information theory measures, including mutual information, information gain, relevancy score, chi-square statistic, odds ratio, expected cross entropy for text, GSS coefficient, to name a few. We describe some of these measures below.

*Mutual information* (MI) measures the mutual dependency of two variables. For a term  $t_k$  and a category  $c_i$ , the MI measure between  $t_k$  and  $c_i$  is defined to be

$$MI(t_k, c_i) = \log \frac{p(t_k, c_i)}{p(t_k)p(c_i)} \quad (6)$$

where  $p(t_k, c_i)$  denotes the probability that the term  $t_k$  appears in a document and this document belongs to the category  $c_i$ ,  $p(t_k)$  is the probability that the term  $t_k$  appears in a document, and  $p(c_i)$  is the probability that a document belongs to the category  $c_i$ . One can see that  $MI(t_k, c_i)$  is zero if  $t_k$  and  $c_i$  is independent, i.e., the term  $t_k$  is useless for discriminating the documents belonging to the category  $c_i$ . In [27], an expected cross entropy for text (CET) is proposed as

$$CET(t_k, c_i) = p(t_k, c_i) \log \frac{p(t_k, c_i)}{p(t_k)p(c_i)} \quad (7)$$

*Information gain* (IG) [17] [28] measures the information if one knows the presence or absence of a term in a document, which is defined as

$$IG(t_k, c_i) = p(t_k, c_i) \log \frac{p(t_k, c)}{p(t_k)p(c)} + p(\bar{t}_k, c_i) \log \frac{p(\bar{t}_k, c)}{p(\bar{t}_k)p(c)} \quad (8)$$

where  $p(\bar{t}_k, c_i)$  denotes the probability that the term  $t_k$  does not appear in a document and this document belongs to the category  $c_i$ ,  $p(\bar{t}_k)$  is the probability that the term  $t_k$  does not appear in a document. Unlike the MI criterion, the IG criterion is less influenced by the low frequency terms and usually performs much better than the MI criterion [17].

*Chi-square* statistic is proposed in [17] to measure the lack of independence between the term  $t_k$  and the category  $c_i$ , which is modeled by a Chi-square ( $\chi^2$ ) distribution. By considering the negative evidence of term in a document, a general  $\chi^2$  statistic measure is defined as

$$Chi(t_k, c_i) = \frac{[p(t_k, c_i)p(\bar{t}_k, \bar{c}_i) - p(t_k, \bar{c}_i)p(\bar{t}_k, c_i)]^2}{p(t_k, c_i)p(t_k, \bar{c}_i)p(\bar{t}_k, c_i)p(\bar{t}_k, \bar{c}_i)} \quad (9)$$

where the document space is divided into two categories,  $c_i$  and its complement  $\bar{c}_i$  that pools all the remaining categories,  $p(\bar{t}_k, \bar{c}_i)$  denotes the probability

that the term  $t_k$  does not appear in a document and also this document does not belong to the category  $c_i$ , and  $p(t_k, \bar{c}_i)$  denotes the probability that the term  $t_k$  appears in a document but this document does not belong to the category  $c_i$ .

A modified measure termed *GSS coefficient* using negative evidence is proposed by Galavotti et al. in [29], which is defined as

$$GSS(t_k, c_i) = p(t_k, c_i)p(\bar{t}_k, \bar{c}_i) - p(t_k, \bar{c}_i)p(\bar{t}_k, c_i) \quad (10)$$

It has been shown that this measure outperforms the original chi-square measure on several data sets [29].

Notice that almost all of these filter approaches based on the information theory measures use binary variables, e.g., the presence ( $t_k$ ) or the absence ( $\bar{t}_k$ ) of a term in a document, and a document belonging to a category ( $c_i$ ) or not ( $\bar{c}_i$ ). Unlike these existing filter approaches, our proposed approaches make use of the term occurrence to measure the term importance in the document, and hence more richer information is contained. Meanwhile, these existing filter approaches rank the features by only exploring the intrinsic characteristics of data based on the feature relevancy without considering their discriminative information in classifiers. It is difficult to select an optimal feature subset for discrimination in a theoretical way. In this paper, our new feature selection approaches are able to involve the learning algorithm by maximizing its discriminative capacity.

### 3 THEORETICAL FRAMEWORK OF FEATURE SELECTION

We follow the Information Theory to select feature subsets that have maximum discriminative capacity for distinguishing the samples among two or more classes. We first introduce some concepts on information measures for binary hypothesis testing (also known as “two-class” classification) and present a new divergence measure for multiple hypothesis testing (i.e., for “multi-class” classification).

#### 3.1 Divergence Measures for Binary Hypothesis Testing

Considering a two-class classification problem first, each class is represented by a particular distribution, saying  $P_1 = p(x|c_1; \theta_1)$  for class  $c_1$  and  $P_2 = p(x|c_2; \theta_2)$  for class  $c_2$ . A test procedure for classification can be considered as a binary hypothesis testing such that if a sample is drawn from  $P_1$  we accept the hypothesis  $H_1$  (reject the hypothesis  $H_2$ ), and if a sample is drawn from  $P_2$  we accept  $H_2$  (reject  $H_1$ ). In other words, we have  $p(x|c_1) = p(x|H_1)$  and  $p(x|c_2) = p(x|H_2)$ , and we also denote  $p(x|H_i)$  as the class conditional probability distribution in the rest of paper.

According to the Information Theory [30] [31], we define the KL-divergence  $\mathcal{KL}(P_1, P_2)$  between two probability distributions (from  $P_1$  to  $P_2$ ) as

$$\begin{aligned}\mathcal{KL}(P_1, P_2) &= \int_x p(x|H_1) \log \frac{p(x|H_1)}{p(x|H_2)} dx \\ &= E_{P_1} \left[ \log \frac{p(x|H_1)}{p(x|H_2)} \right]\end{aligned}\quad (11)$$

where  $E_{P_1}[x]$  denotes the expectation of  $x$  with respect to the probability distribution  $P_1$ . Specifically, it is easy to obtain the KL-divergence measure between two discrete distributions that are commonly used in text categorization. According to Eq. (1) and Eq. (11), the KL-divergence between two multinomial distributions  $P_1^m$  and  $P_2^m$  is given by

$$\mathcal{KL}(P_1^m, P_2^m) = \sum_{i=1}^M p_{i1} \log \frac{p_{i1}}{p_{i2}} \quad (12)$$

where  $0 \leq p_{ic} \leq 1$  and  $\sum_{i=1}^M p_{ic} = 1$  for  $i = 1, 2, \dots, M, c = 1, 2$ .

Under the MAP rule in Eq. (5), we would classify the sample  $x$  into class  $c_1$ , i.e., accept  $H_1$ , if

$$\begin{aligned}\log p(x|H_1) + \log p(H_1) &> \log p(x|H_2) + \log p(H_2) \\ \Rightarrow \log \frac{p(x|H_1)}{p(x|H_2)} &> -\log \frac{p(H_1)}{p(H_2)} = \gamma\end{aligned}\quad (13)$$

The logarithm of the likelihood ratio,  $\log \frac{p(x|H_1)}{p(x|H_2)}$ , measures the information of the observation  $x$  for discrimination in favor of  $H_1$  against  $H_2$  [30]. From the definition of the KL-divergence measure in Eq. (11), then  $\mathcal{KL}(P_1, P_2)$  indicates the mean information for discrimination in favor of  $H_1$  against  $H_2$ .

Taking the expectation with respect to the distribution  $P_1$  in Eq. (13) (i.e., considering all possible observations drawn from  $P_1$ ), we have

$$\mathcal{KL}(P_1, P_2) > \gamma \quad (14)$$

which also illustrates that  $\mathcal{KL}(P_1, P_2)$  is an indicator of a Bayesian classifier's discriminative capacity for discriminating the observation from class  $c_1$  in favor of  $H_1$  against  $H_2$ . With the extension of the central limit theorem, Chernoff in [32] showed that, for a large number of observations, the type I error  $\alpha^*$ , the probability of incorrectly accepting  $H_1$ , asymptotically has

$$\lim_{n \rightarrow \infty} \log \frac{1}{\alpha^*} = \lim_{n \rightarrow \infty} \mathcal{KL}(P_1, P_2; O_n) \quad (15)$$

where  $O_n$  denotes  $n$  independent observations. It has been said that a larger value of KL-divergence indicates a lower type I error when there are infinite number of observations.

Note that the KL-divergence measure is not symmetric. Alternatively, the KL-divergence  $\mathcal{KL}(P_2, P_1)$  indicates a Bayesian classifier's discriminative capacity for discriminating the observation from class  $c_2$  in favor of  $H_2$  against  $H_1$ . Similarly, under the MAP rule

in Eq. (5), considering all the observation drawn from class  $c_2$ , we have

$$\mathcal{KL}(P_2, P_1) > -\gamma \quad (16)$$

For a large number of observations, the type II error  $\beta^*$ , the probability of incorrectly accepting  $H_2$ , can be given by

$$\lim_{n \rightarrow \infty} \log \frac{1}{\beta^*} = \lim_{n \rightarrow \infty} \mathcal{KL}(P_2, P_1; O_n) \quad (17)$$

To minimize both type I error and type II errors in an asymptotic way, the Jeffreys divergence (J-divergence) can be used, which is defined by [33]

$$\mathcal{J}(P_1, P_2) = \mathcal{KL}(P_1, P_2) + \mathcal{KL}(P_2, P_1) \quad (18)$$

By combining Eq. (14) and Eq. (16), we have

$$\mathcal{KL}(P_1, P_2) > \gamma > -\mathcal{KL}(P_2, P_1) \quad (19)$$

Because

$$\mathcal{KL}(P_1, P_2) \geq 0, \quad \mathcal{KL}(P_2, P_1) \geq 0 \quad (20)$$

a larger  $\mathcal{J}(P_1, P_2)$  may lead to a smaller type I and type II error asymptotically. It is known that the J-divergence is able to measure the difficulty and capacity of discriminating between  $H_1$  and  $H_2$  [30] [32] [33].

The purpose of feature selection is to determine the most informative features which lead to the best prediction performance. Hence, it is natural to select those features that have the maximum discriminative capacity for classification, by minimizing the classification error (i.e., maximizing the KL-divergence or the J-divergence). However, the J-divergence is only defined for binary hypothesis. We next extend the J-divergence for multiple hypothesis testing (i.e., multi-class classification).

### 3.2 Jeffreys-Multi-Hypothesis Divergence

The Jensen-Shannon (JS) divergence [34] is the one that can be used to measure multi-distribution divergence, in which the divergences of each individual distribution with a reference distribution are calculated and summed together. Unlike the J-divergence, the measure of discrimination capacity may not hold. In [35], Sawyer presents a variant of J-divergence with a variance-covariance matrix for multiple comparisons of separate hypotheses. Here, we first generalize the J-divergence to a multi-distribution using the scheme of "one-vs-all" [36], which is defined as follows:

**Definition 1:** Let  $\mathcal{P} = \{P_1, P_2, \dots, P_N\}$  be the set of  $N$  distributions. The Jeffreys-Multi-Hypothesis (JMH) divergence, denoted by  $\mathcal{JMH}(P_1, P_2, \dots, P_N)$ , is defined to be

$$\mathcal{JMH}(P_1, P_2, \dots, P_N) = \sum_{i=1}^N \mathcal{KL}(P_i, \bar{P}_i) \quad (21)$$

where  $\bar{P}_i$  is the combination of all remaining  $N - 1$  distributions  $\bar{P}_i = \sum_{k=1, k \neq i}^N \pi_{ki} P_k$ , and  $\pi_{ki}$  are the prior coefficients.

Similar to the “one-vs-all” strategy for multi-class classification problem, we build  $N$  binary hypothesis testing detectors, each of which discriminates the samples in favor of  $P_i$  against  $\bar{P}_i$  which is the complement of  $P_i$ . In each detector,  $\bar{P}_i$  is represented by a mixture distribution over all the remaining  $N - 1$  classes with the coefficients  $\pi_{ki}$  which are given by

$$\pi_{ki} = \frac{p_{c_k}}{\sum_{m=1, m \neq i}^N p_{c_m}} \quad (22)$$

where  $p_{c_m}$  is the prior probability of class  $c_m$ . Since the KL-divergence of each detector is the measure of its discriminative capacity for discrimination, the new multi-distribution divergence is able to measure the discrimination capacity over all classes. Specifically, when  $N = 2$ , we have  $\mathcal{JMH}(P_1, P_2) = \mathcal{J}(P_1, P_2)$ .

Note that, since the JMH divergence is the sum of multiple J-divergences, it holds most properties of J-divergence. For example, JMH divergence is almost positive definite, i.e.,  $\mathcal{JMH}(P_1, P_2, \dots, P_N) \geq 0$ , with equality if and only if  $p_{1m} = p_{2m} = \dots = p_{Nm}$ ,  $m = 1, 2, \dots, M$ . Also it is a symmetric measure, that is,  $\mathcal{JMH}(\dots, P_i, \dots, P_k, \dots) = \mathcal{JMH}(\dots, P_k, \dots, P_i, \dots)$ .

## 4 SELECTING THE MAXIMUM DISCRIMINATIVE FEATURES

### 4.1 A Greedy Feature Selection Approach

We consider a binary (two-class) classification problem first and extend our feature selection method to a general multi-class classification problem later. Unlike those existing feature selection methods which compute the score (“importance”) of features based on the feature relevance to class, our goal is to select the features that offer the maximum discrimination for classification. By doing so, one can expect an improved classification performance for text categorization.

For a two-class classification problem, we know that the J-divergence indicates the discriminative capacity of discriminating two classes data under the MAP rule. Hence, we formulate the feature selection problem as follows: given a set of  $M$  features  $\mathcal{B}$  where  $|\mathcal{B}| = M$  and a predetermined integer  $r$ , the number of features to be selected, we aim to find the most  $r$  discriminative features  $\mathcal{B}_r^* \subset \mathcal{B}$  where  $|\mathcal{B}^*| = r$ , such that,

$$\begin{aligned} \mathcal{B}_r^* &= \arg \max_{\mathcal{B}_r \subset \mathcal{B}} \mathcal{J}(P_1, P_2 | \mathcal{B}_r) \\ &= \arg \max_{\mathcal{B}_r \subset \mathcal{B}} \mathcal{KL}(P_1, P_2 | \mathcal{B}_r) + \mathcal{KL}(P_2, P_1 | \mathcal{B}_r) \end{aligned} \quad (23)$$

where  $\mathcal{J}(\cdot, \cdot | \mathcal{B}_r)$  and  $\mathcal{KL}(\cdot, \cdot | \mathcal{B}_r)$  are the J-divergence and KL-divergence, respectively, when a subset of

features  $\mathcal{B}_r \subset \mathcal{B}$  are considered. This problem is also known as a NP-hard problem, if one examines each of  $\binom{M}{r}$  combinations, which is intractable particularly for a high dimensional data set. Meanwhile, in practice, we need to examine various  $r$  values to evaluate the classification performance using those selected  $r$  features. Hence, it is necessary to assign an importance score to each feature and rank the features.

Here, we start to propose a greedy approach to rank the features according to their discriminative capacity for naive Bayes. This approach starts to determine which feature of the  $M$  features produces the maximum JMH-divergence if only one single feature is used for classification. To determine the most discriminative feature, for each feature  $x_i$ ,  $i = 1, 2, \dots, M$ , we build two variables:  $x_i$  and  $\bar{x}_i$ , where  $x_i$  is the original  $i$ -th feature and  $\bar{x}_i$  is the pool of all the remaining  $M - 1$  features with parameters  $\bar{p}_{ic}^k$ . We use the superscript number  $k$  to indicate the  $k$ -th step of our greedy approach (e.g.,  $k = 1$  here). We have

$$\begin{aligned} \bar{p}_{i1}^1 &= \sum_{m=1, m \neq i}^M p_{m1} = 1 - p_{i1} \\ \bar{p}_{i2}^1 &= \sum_{m=1, m \neq i}^M p_{m2} = 1 - p_{i2} \end{aligned} \quad (24)$$

We denote the distributions of these two variables for class 1 and class 2 by  $P_{i1}^1$  and  $P_{i2}^1$ , respectively. Note that  $P_{i1}^1$  and  $P_{i2}^1$  also satisfy multinomial distribution but with different parameters. Then, we calculate the J-divergence between  $P_{i1}^1$  and  $P_{i2}^1$  as  $\mathcal{J}_i^1(P_{i1}^1, P_{i2}^1) = \mathcal{KL}(P_{i1}^1, P_{i2}^1) + \mathcal{KL}(P_{i2}^1, P_{i1}^1)$  with Eq. (12). At the end of the 1st step, we obtain  $M$  J-divergences  $\mathcal{J}_i^1(P_{i1}^1, P_{i2}^1)$  for  $i = 1, 2, \dots, M$ , and choose the first feature indexed by  $s_1$  that leads to the maximum J-divergence, that is,

$$s_1 = \arg \max_{i=1, 2, \dots, M} \mathcal{J}_i^1(P_{i1}^1, P_{i2}^1) \quad (25)$$

Then, we fix the first feature  $s_1$  and repeat the process over the remaining features. Specifically, at the  $k$ -th step to select the  $k$ -th feature, let  $\mathcal{S}_{k-1} = \{s_1, s_2, \dots, s_{k-1}\}$  be the feature index set which are selected from the previous  $k - 1$  steps. Again, for each individual feature  $x_i$ ,  $i = 1, 2, \dots, M$ ,  $i \notin \mathcal{S}_{k-1}$ , we form  $k + 1$  variables:  $x_{s_1}, x_{s_2}, \dots, x_{s_{k-1}}, x_i, \bar{x}_i$ , where the first  $k$  variables are the original features and the last variable  $\bar{x}_i$  is the pool of all remaining  $M - k$  features with parameters  $\bar{p}_{ic}^k$  for two classes ( $c = 1, 2$ ). We have

$$\begin{aligned} \bar{p}_{i1}^k &= 1 - p_{i1} - \sum_{m \in \mathcal{S}_{k-1}} p_{m1} \\ \bar{p}_{i2}^k &= 1 - p_{i2} - \sum_{m \in \mathcal{S}_{k-1}} p_{m2} \end{aligned} \quad (26)$$

---

**Algorithm 1:** A Greedy Feature Selection Algorithm Based on the Maximum J-Divergence for Two-class Classification

---

**INPUT:**

- The estimated probabilities of  $M$  features:  
 $\theta_i = [p_{1c}, p_{2c}, \dots, p_{Mc}]$ ,  $c = 1, 2$  for two classes;
- The ranked feature index set:  $S = \emptyset$ .
- The full index set:  $S_a = \{1, 2, \dots, M\}$

**ALGORITHM:**

for  $k = 1 : M$  do

  foreach  $i \in S_a \cap \bar{S}$  do

1. Form  $k + 1$  variables:  $x_{s_1}, x_{s_2}, \dots, x_{s_{k-1}}, x_i$  and  $\bar{x}_i$ , and denote  $P_{i1}^k$  and  $P_{i2}^k$  for class 1 and 2 distribution, respectively;
2. Calculate J-divergence  $J_i^k(P_{i1}^k, P_{i2}^k)$  between  $P_{i1}^k$  and  $P_{i2}^k$ ;

  end

3.  $s_k$ : select the feature using Eq. (27);

4.  $S = S \cup s_k$ ;

end

**OUTPUT:**

- A ranked feature index set:  $S = \{s_1, s_2, \dots, s_M\}$ .
- 

Denote the distributions of these variables for class 1 and class 2 at the  $k$ -th step by  $P_{i1}^k$  and  $P_{i2}^k$ , respectively. At the  $k$ -th step, we choose the feature indexed by  $s_k$  with the following maximum J-divergences:

$$s_k = \arg \max_{i=1,2,\dots,M, i \notin S_{k-1}} J_i^k(P_{i1}^k, P_{i2}^k) \quad (27)$$

Hence, at the end of the  $M$ -th step, a ranked feature index set  $S_M = \{s_1, s_2, \dots, s_M\}$  is produced. The implementation of this greedy feature selection approach based on the maximum J-divergence for two-class classification is given in Algorithm 1.

**Theorem 1:** The maximum J-divergences  $\mathcal{J}_{s_k}$  for  $k = 1, 2, \dots, M$  in Algorithm 1 monotonically increases, i.e.,

$$\mathcal{J}_{s_{k+1}}^{k+1}(P_{s_{k+1}1}^{k+1}, P_{s_{k+1}2}^{k+1}) \geq \mathcal{J}_{s_k}^k(P_{s_k1}^k, P_{s_k2}^k) \quad (28)$$

The proof of Theorem 1 is provided in our Supplemental Material. This theorem indicates that the discriminative capacity increases when more features are used for classification, under the assumption that the term occurrence of a document satisfies a particular multinomial distribution. Note that the proposed greedy feature selection algorithm makes a locally optimal choice at each step to approximate the global optimal solution of Eq. (23), by selecting a feature with the maximum discriminative capacity for classification. The significance of this algorithm is that it starts at the best first feature and towards the optimal solution.

This greedy approach can be considered as a wrapper approach. However, unlike those existing wrapper approaches, this greedy approach does not need to

evaluate the classification performance on a validation data set through retraining the classifier when a new feature is generated, because a closed form of KL-divergence is given in Eq. (12) to measure the discriminative capacity of MNB classifiers.

However, this greedy approach still has the computational complexity of  $\mathcal{O}(M^2/2)$ , which leads to a heavy computational load for a high-dimensional data set. Next, we provide a more efficient feature selection approach for text categorization.

## 4.2 An Efficient Feature Selection Approach

In Algorithm 1, the best one single feature is selected at the first step, providing an optimal starting point to approximate the optimal solution. At this step, we rank the J-divergences over all features, which are given by

$$J_{e_1}^1(P_{e_11}^1, P_{e_12}^1) \geq J_{e_2}^1(P_{e_21}^1, P_{e_22}^1) \geq \dots \geq J_{e_M}^1(P_{e_M1}^1, P_{e_M2}^1) \quad (29)$$

where  $e_i$  denotes the feature index and we know  $e_1$  is  $s_1$  which is given by Eq. (25) in Algorithm 1. Looking at the first two J-divergences in Eq. (29), we have

$$\begin{aligned} p_{e_11} \log \frac{p_{e_11}}{p_{e_12}} + (1 - p_{e_11}) \log \frac{1 - p_{e_11}}{1 - p_{e_12}} \\ \geq p_{e_21} \log \frac{p_{e_21}}{p_{e_22}} + (1 - p_{e_21}) \log \frac{1 - p_{e_21}}{1 - p_{e_22}} \end{aligned} \quad (30)$$

Since one single term probability is very small in practice, i.e.,  $p_{i1} \ll 1$  for  $i = 1, 2, \dots, M$ , the term  $\log \frac{1 - p_{i1}}{1 - p_{i2}}$  would be very close to zeros, and then Eq. (30) may lead to

$$p_{e_11} \log \frac{p_{e_11}}{p_{e_12}} \geq p_{e_21} \log \frac{p_{e_21}}{p_{e_22}} \quad (31)$$

At the second step in Algorithm 1, the feature  $e_2$  is usually selected due to the fact that  $\log \frac{1 - p_{e_11} - p_{e_21}}{1 - p_{e_12} - p_{e_22}} \approx 0$ . That is because, according to Eq. (31), we have

$$p_{e_11} \log \frac{p_{e_11}}{p_{e_12}} + p_{e_21} \log \frac{p_{e_21}}{p_{e_22}} \geq p_{e_11} \log \frac{p_{e_11}}{p_{e_12}} + p_{j1} \log \frac{p_{j1}}{p_{j2}} \quad (32)$$

for  $j = 1, 2, \dots, M$ ,  $j \neq e_1, e_2$ .

Therefore, instead of doing a greedy search, an efficient way is to use the ranked feature index set  $\mathcal{E} = \{e_1, e_2, \dots, e_M\}$  in Eq. (29). We summarize this efficient feature selection algorithm in Algorithm 2. Compared to the Algorithm 1, the proposed approach is much more *efficient* as each feature score is calculated once, and it has the computational complexity of  $\mathcal{O}(M)$ . This efficient approach evaluates the “importance” of each individual feature by measuring its discriminative capacity, when only one single feature is used for classification. We note that the Theorem 1 is also satisfied for this efficient approach, i.e., the J-divergence measure increases as more features are selected. Meanwhile, we also note that the features

---

**Algorithm 2:** An Efficient Feature Selection Algorithm Based on the Maximum J-Divergence for Two-class Classification

---

**INPUT:**

- The estimated probabilities of  $M$  features:  
 $\theta_i = [p_{1c}, p_{2c}, \dots, p_{Mc}]$ ,  $c = 1, 2$  for two classes;

**ALGORITHM:**

**for**  $i = 1 : M$  **do**

1. Form two variables:  $x_i$  and  $\bar{x}_i$ , and denote  $P_{i1}$  and  $P_{i2}$  for class 1 and 2 distribution, respectively;
2. Calculate J-divergence  $J_i(P_{i1}, P_{i2})$  between  $P_{i1}$  and  $P_{i2}$  using Eq. (12) as score for the  $i$ -th feature;

**end**

3. Sort the feature scores in a descend way:

$$\mathcal{J}_{e_1} > \mathcal{J}_{e_2} > \dots > \mathcal{J}_{e_M}.$$

**OUTPUT:**

- A ranked feature index set:  $\mathcal{E} = \{e_1, e_2, \dots, e_M\}$ .
- 

selected in Algorithm 2, are not necessarily the same as ones selected in Algorithm 1, as one can see, for example, that Eq. (31) and Eq. (32) approximately hold for the first and second feature selection. Considering the first feature selection as an example, this approximation depends on the value of  $\Delta$  defined as

$$\Delta = (1 - p_{e_21}) \log \frac{1 - p_{e_21}}{1 - p_{e_22}} - (1 - p_{e_11}) \log \frac{1 - p_{e_11}}{1 - p_{e_12}} \quad (33)$$

More precisely, given  $e_1$  is selected as the first feature in Algorithm 2,  $e_1$  also ranks first in Algorithm 1 if and only if the following condition holds:

$$\Delta \leq p_{e_11} \log \frac{p_{e_11}}{p_{e_12}} - p_{e_21} \log \frac{p_{e_21}}{p_{e_22}} \quad (34)$$

### 4.3 Multi-class Classification

In this section, we extend the above efficient feature selection method for multi-class classification problems. Considering an  $N$ -class classification problem, the  $r$  most discriminative features  $\mathcal{B}_r^*$  are selected by maximizing the JMH-divergence, which are given by

$$\begin{aligned} \mathcal{B}_r^* &= \arg \max_{\mathcal{S}_r \subset \mathcal{B}, |\mathcal{S}_r|=r} \mathcal{JMH}(P_1, P_2, \dots, P_N | \mathcal{S}_r) \\ &= \arg \max_{\mathcal{S}_r \subset \mathcal{B}, |\mathcal{S}_r|=r} \sum_{i=1}^N \mathcal{KL}(P_i, \bar{P}_i | \mathcal{S}_r) \end{aligned} \quad (35)$$

where  $\mathcal{JMH}(P_1, P_2, \dots, P_N | \mathcal{S}_r)$  is the JMH-divergence defined in Eq. (21) with the feature subset  $\mathcal{S}_r$ . Note that each KL-divergence indicates the discriminative capacity of one binary classifier to distinguish the samples in one class  $P_i$  from the samples in all remaining classes  $\bar{P}_i$ , and thus the JMH-divergence is able to measure the difficulty and

---

**Algorithm 3:** An Efficient Feature Selection Algorithm Based on the Maximum JMH-Divergence for  $N$ -class Classification

---

**INPUT:**

- The estimated probabilities of  $M$  features:  
 $\theta_c = [p_{1c}, p_{2c}, \dots, p_{Mc}]$ ,  $c = 1, 2, \dots, N$ ;
- The prior probabilities of  $N$  classes:  
 $p_1, p_2, \dots, p_N$ ;

**ALGORITHM:**

**for**  $i = 1 : M$  **do**

1. Form two variables:  $x_i$  and  $\bar{x}_i$ , and denote  $P_{ic}$  as the class distribution,  $c = 1, 2, \dots, N$ ;

**for**  $c = 1 : N$  **do**

2. Form two distributions:  $P_{ic}$  and  $\bar{P}_{ic}$ , where  $\bar{P}_{ic}$  is the one grouping all remaining  $N - 1$  classes;
3. Calculate KL-divergence  $\mathcal{KL}(P_{ic}, \bar{P}_{ic})$  between  $P_{ic}$  and  $\bar{P}_{ic}$ ;

**end**

4. Calculate the JMH-divergence

$$\mathcal{JMH}_i = \sum_{c=1}^N \mathcal{KL}(P_{ic}, \bar{P}_{ic}) \text{ as the score for the } i\text{-th feature;}$$

**end**

5. Sort the feature scores in a descend way:

$$\mathcal{JMH}_{e_1} > \mathcal{JMH}_{e_2} > \dots > \mathcal{JMH}_{e_M}.$$

**OUTPUT:**

- A ranked feature index set:  $\mathcal{E} = \{e_1, e_2, \dots, e_M\}$ .
- 

capacity of discriminating the samples among all classes.

The efficient feature selection method for an  $N$ -class classification problem based on the maximum JMH-divergence is presented in Algorithm 3. The value of JMH-divergence is used as the score for each feature. Sorting the feature scores in a descend way, we output a ranked feature index set  $\mathcal{E} = \{e_1, e_2, \dots, e_M\}$  for multi-class classification. The computational complexity of this algorithm is  $\mathcal{O}(MN)$ .

### 4.4 Feature Selection Based on $\chi^2$ Statistics

The KL-divergence measure  $\mathcal{KL}(P_1, P_2)$  is also known as the minimum discrimination information between the probability distributions  $P_1$  and  $P_2$ . Suppose that we have a random sample of  $l$  observations (e.g., a document with length  $l$ ), and we try to test a null hypothesis  $H_2$ , the observation is drawn from class 2 with the distribution  $P_2$ , against an alternative hypothesis  $H_1$ , the observation is drawn from class 1 with the distribution  $P_1$ . The minimum discrimination statistic [30] in favor of  $H_1$  against  $H_2$  is defined as

$$\mathcal{D}(\hat{P}_1, P_2) = \hat{P}_1 \log \frac{\hat{P}_1}{P_2} \quad (36)$$

where  $\hat{P}_1$  is the estimate of  $P_1$  from the given observations and  $P_2$  is assumed to be known. We may reject

the null hypothesis  $H_2$  and accept the alternative hypothesis  $H_1$  if the value of statistic  $\mathcal{D}(\hat{P}_1, P_2)$  exceeds a predetermined threshold. Asymptotically, the statistic  $2\mathcal{D}(\hat{P}_1, P_2)$  under the null hypothesis  $H_2$  satisfies a central Chi-squared distribution  $\chi^2_{M-1}$  with  $M-1$  degrees of freedom, and satisfies a non-central Chi-squared distribution  $\chi^2_{M-1}(\nu)$  under the alternative hypothesis  $H_1$ . The noncentrality parameter  $\nu_D$  is given by

$$\nu_D(\hat{P}_1, P_2) = l \sum_{i=1}^M \frac{(\hat{p}_{i1} - p_{i2})^2}{p_{i2}} \quad (37)$$

where  $\hat{p}_{i1}$  is the estimate of  $p_{i1}$  under  $H_1$  using the MLE method from the given observations. Asymptotically, the J-divergence  $\mathcal{J}(\hat{P}_1, P_2)$  is the sum of two  $\chi^2$  distributions. We have

$$\nu_J(\hat{P}_1, P_2) = \frac{l}{2} \sum_{i=1}^M \frac{(\hat{p}_{i1} - p_{i2})^2}{p_{i2}} + \frac{l}{2} \sum_{i=1}^M \frac{(\hat{p}_{i1} - p_{i2})^2}{\hat{p}_{i1}} \quad (38)$$

where the former one is known as the Pearson's  $\chi^2$  and the latter one is also known as the Neyman's  $\chi^2$  [30].

Thus, one can see that the noncentrality parameter  $\nu_J$  would also be a good sign to indicate the discriminative capacity of discriminating between  $H_1$  and  $H_2$ . Therefore, we can further select the features by maximizing the noncentrality parameter  $\nu_J$  for binary classification. Under the assumption that the number of samples in training data set goes infinity, we use the estimation of  $p_{i1}$  and  $p_{i2}$  for  $i = 1, 2, \dots, M$  from training data in Eq. (38). For a multi-class classification problem, unlike Algorithm 3, each feature is assigned with a score  $\mathcal{CHI}_i = \sum_{c=1}^N \nu(P_{ic}, \bar{P}_{ic})$  which is the sum of  $N$  noncentrality parameters in their  $\chi^2$  distributions, and the ranked feature index set  $\mathcal{S} = \{e_1, e_2, \dots, e_M\}$  is produced by sorting the scores:

$$\mathcal{CHI}_{e_1} > \mathcal{CHI}_{e_2} > \dots > \mathcal{CHI}_{e_M} \quad (39)$$

We need to note here that the feature selection approach based upon the Chi-squared statistic in Eq. (39) should be equivalent to the approach in Algorithm 3 if there are infinite training documents. When the assumption of large numbers is not satisfied, the features selected by Eq. (39) may lose some discriminative capacity. As we demonstrated through extensive experiments, the discrimination performance of the Chi-squared statistic is usually bounded by the approach in Algorithm 3. However, we sometimes also observe that the Chi-squared statistic performs better for some real-life text data sets. Therefore, in practice, we would like to recommend the use of both approaches because of their simplicity, efficiency and improved discrimination performance.

## 5 EXPERIMENTAL RESULTS AND ANALYSIS

### 5.1 Real-Life Data Sets

For our experiments,<sup>1</sup> we test our proposed feature selection approaches on three benchmarks that have been prepared by Deng et al. [37] [38] for text categorization: 20-NEWSGROUPS, REUTERS, and TOPIC DETECTION AND TRACKING (TDT2). These three benchmarks have been widely used in literature for performance evaluation. The 20-NEWSGROUPS benchmark consists of about 20,000 documents collected from the postings of 20 different online newsgroups or topics.

The REUTERS originally contains 21,578 documents with 135 topics, but some documents belong to multiple topics. For our experiments, we use the ModApte version of the Reuters by removing those documents with multiple labels. This version consists of 8,293 documents in 65 topics. Following the work in [20], we form three data sets, named REUTERS-10, REUTERS-20 and REUTERS-30, consisting of the documents of the first 10, 20 and 30 topics, respectively.

The TDT2 benchmark consists of 11,201 documents taken from two newswires (AP WorldStream and New York Times Newservice), two radio programs (PRI The World and VOA World News) and two television programs (CNN Headline News and ABC World News Tonight). Also, those documents that belong two or more topics have been removed. Because of the extremely imbalanced data for some categories, we only use the first 10 topics with the largest data size as our data set.

For all these data sets used in our experiments, we ignore those words in a stoplist, and discard those words that appear in less than 2 documents or messages in our preprocessing stage. For all data sets except TDT2, we perform classification on the officially split training and testing data sets. For the TDT2, we use 10-fold cross validation for performance evaluation, and the reported results are averaged over 10 runs.

### 5.2 Performance Evaluation Metrics

We use the following metrics to evaluate the classification performance: *accuracy*, *precision*, *recall*, and *F1 measure*. The accuracy metric is widely used in machine learning fields, which indicates the overall classification performance. The precision is the percentage of documents that are correctly classified as positive out of all the documents that are classified as positive, and the recall is the percentage of documents that are correctly classified as positive out of all the documents that are actually positive. The metrics of

1. We also verify the effectiveness of the proposed approaches for synthetic data sets. The detailed simulation results and analysis are given in our Supplemental Material.

precision and recall are defined as

$$\begin{aligned} \text{Precision} &= \frac{TP}{TP + FP} \\ \text{Recall} &= \frac{TP}{TP + FN} \end{aligned} \quad (40)$$

where  $TP$  denotes the number of true positive,  $FP$  denotes the number of false positive, and  $FN$  denotes the number of false negative. These two metrics have an inverse relationship between each other. In other words, increasing the precision is at the cost of reducing the recall, and vice versa. Among those measures that attempt to combine precision and recall as one single measure, the F1 measure is one of the most popular, which is defined by

$$F1 = \frac{2 \times \text{Precision} \times \text{Recall}}{\text{Precision} + \text{Recall}} \quad (41)$$

The metrics of precision, recall and F1 measure are originally defined for binary class. For multi-class classification, we follow several other studies [1] [20] [39] [40], in which binary classifiers are built for each individual class and a global F1 measure is obtained by averaging the F1 measure of each class weighted by the class prior.

### 5.3 Results

We compare our two efficient feature selection approaches: the maximum discrimination termed  $MD$  and its asymptotic  $\chi^2$  statistic termed  $MD-\chi^2$ , with the state-of-the-art feature ranking methods, including document frequency ( $DF$ ), expected cross entropy for text ( $CET$ ),  $\chi^2$  statistic and  $GSS$ . We carry out experiments on these three benchmarks when naive Bayes and SVM are used as classifiers. To compare the performance of these feature selection methods, we evaluate the classification accuracy and the F1 measure metric of these classifiers with different number of features ranging from 10 to 2,000.

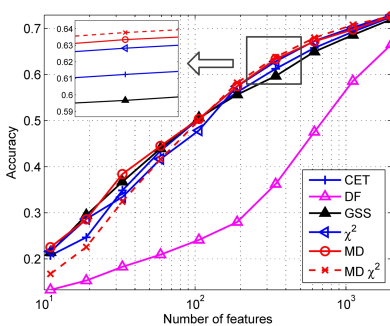
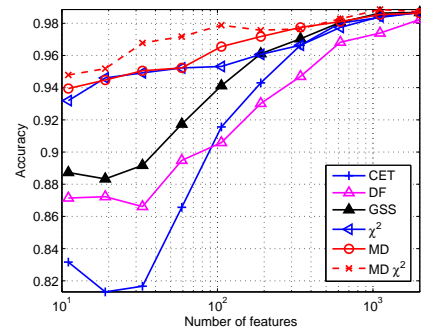
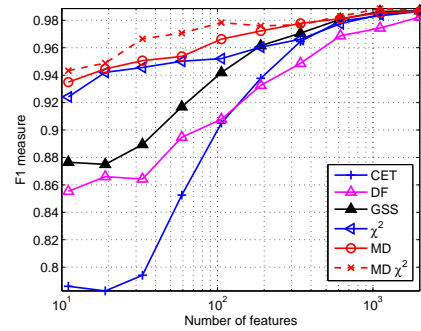


Fig. 1: An accuracy comparison of feature selection methods on the 20-NEWSGROUPS with 20 topics for naive Bayes.

We first test these feature selection approaches when naive Bayes is used as the classifier. Fig. 1 shows



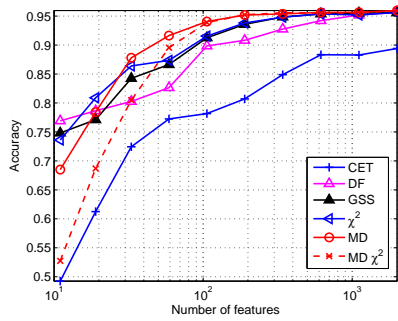
(a). Accuracy for ALT-COMP



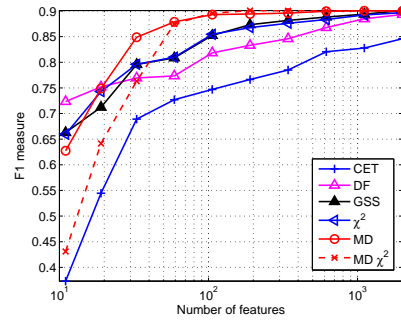
(b). F1 measure for ALT-COMP

Fig. 2: Performance comparisons of feature selection methods on the ALT-COMP data set: (a) accuracy, (b) F1 measure, when naive Bayes is used as the classifier.

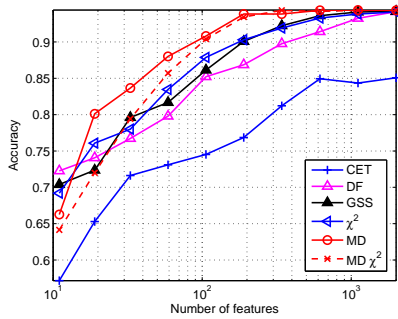
the results on the 20-NEWSGROUPS data set. It can be shown that the performance is improved when more features are selected. The proposed two approaches commonly perform better than others. As one can see in the zoomed-in figure, the proposed  $MD$  method performs better than the others. The  $DF$  method is the worst one for this data set. We further show the results in Fig. 2 on the data set of ALT-COMP that is a subset of 20-NEWSGROUPS with the categories *alt.\** and the categories *comp.\**. For this data set, the  $MD-\chi^2$  is the best one among all others. The comparison of the F1 measure on ALT-COMP is given in Fig. 2(b) and shows that the proposed  $MD$  and its asymptotic  $\chi^2$  statistic are the best two approaches. Comparing the results in 1 and Fig. 2, there is a significant difference on the performance behavior, although the ALT-COMP is a subset of 20-NEWSGROUPS. For example, the accuracy of the 20-NEWSGROUPS data set is even lower than 20% when a small subset of features are selected, while the accuracy of the ALT-COMP data set is higher than 80%. This might indicate a diverse feature characteristic of the 20-NEWSGROUPS data set. In the 20-NEWSGROUPS data set, some topics belong to the same category and are very closely related to each other, e.g., *rec.sport.baseball* and *rec.sport.hockey*, *comp.sys.mac.hardware* and *comp.sys.ibm.hardware*, etc. We also notice that the  $MD-\chi^2$  method performs better than  $MD$  method in Fig. 2, which might not



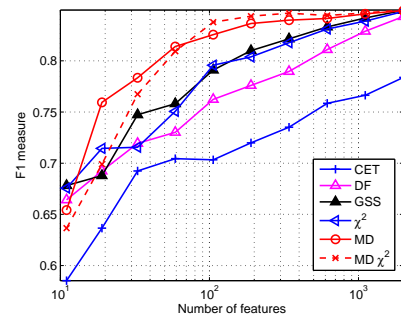
(a-1). Accuracy for REUTERS-10



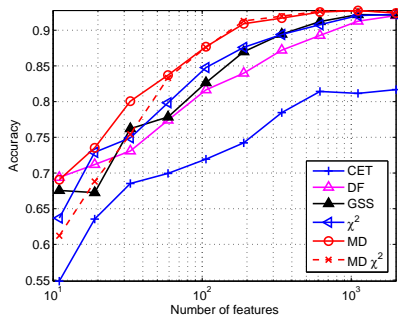
(a-2). F1 measure for REUTERS-10



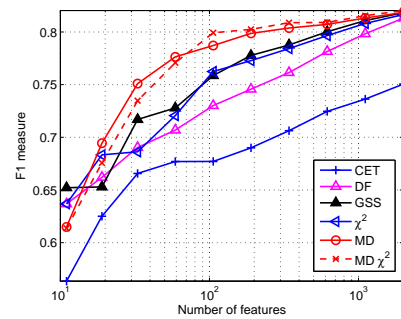
(b-1). Accuracy for REUTERS-20



(b-2). F1 measure for REUTERS-20



(c-1). Accuracy for REUTERS-30



(c-2). F1 measure for REUTERS-30

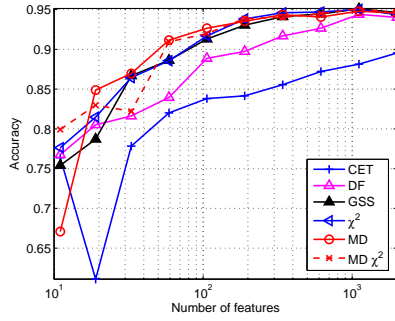
Fig. 3: The results of accuracy and F1 measure on the data sets of (a). REUTERS-10, (b). REUTERS-20, and (c). REUTERS-30, when naive Bayes is used as the classifier.

be always true for other data sets since the  $MD-\chi^2$  method is based on the asymptotic distribution of the statistic used in  $MD$  method. One possible explanation is that the correlation among some words holds more discriminative information, while the  $MD$  method assumes words are independent to each other. However, the asymptotic distribution may still hold because of the law of large number. The theoretical supports to determine which method performs better need further study.

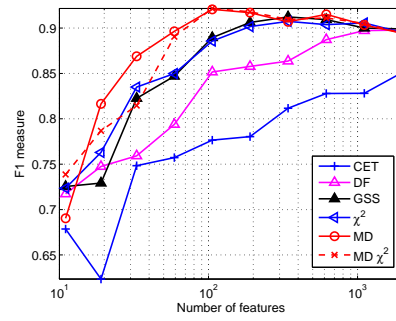
Fig. 3 shows the comparison results on three data sets in REUTERS: REUTERS-10, REUTERS-20, and REUTERS-30. It can be shown that our proposed two approaches with the first 200 selected features can achieve the similar performance as other four existing approaches with the first 1000 selected features. Moreover, as seen in Fig. 3 (a), (b) and (c), the performance improvement of the proposed two approaches

is increased in comparison with other methods, when more categories are considered.

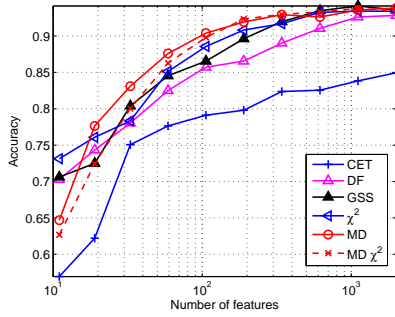
Fig. 4 shows results of the data set of TDT2. As the classification tasks on this data set are performed with the scheme of 10-fold cross validation, all the results in Fig. 4 are averaged across 10 runs. As shown in Fig. 4(a), the proposed  $MD$  and  $MD-\chi^2$  outperform all the others with respect to the metric of accuracy. It is interesting to notice that, when the first 100 features are selected, the  $MD$  obtains the accuracy of 96.46%, and the  $MD-\chi^2$  has 95.76%. For other methods, such as  $GSS$  and  $\chi^2$ , the first 1000 features need to be selected to achieve the same classification accuracy.



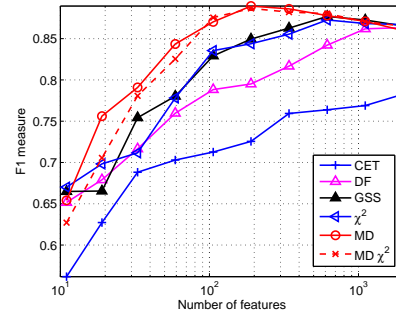
(a-1). Accuracy for REUTERS-10



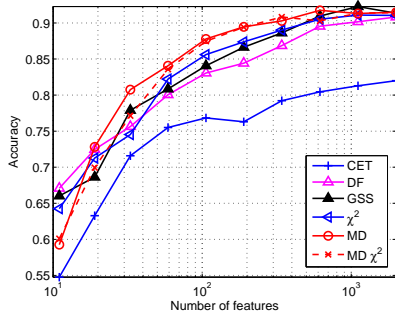
(a-2). F1 measure for REUTERS-10



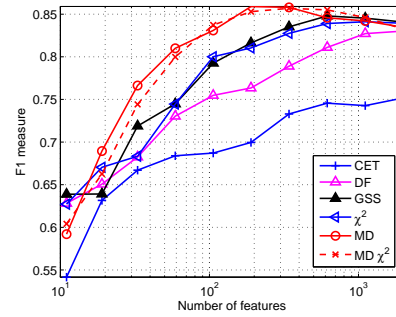
(b-1). Accuracy for REUTERS-20



(b-2). F1 measure for REUTERS-20



(c-1). Accuracy for REUTERS-30



(c-2). F1 measure for REUTERS-30

Fig. 5: The results of accuracy and F1 measure on the data sets of (a). REUTERS-10, (b). REUTERS-20, and (c). REUTERS-30, when SVM is used as the classifier.

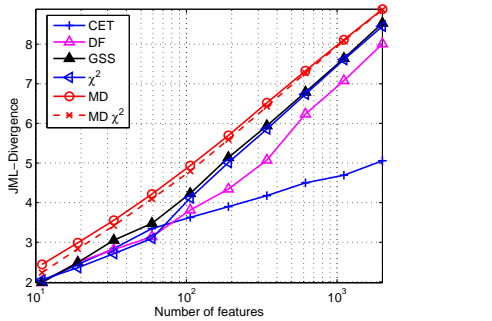
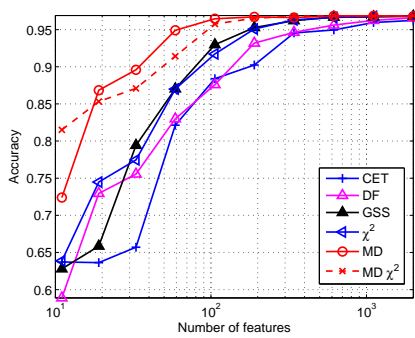


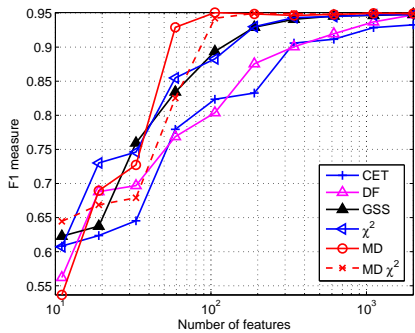
Fig. 6: JMH-divergence with different number of features selected by *CET*, *DF*, *GSS*,  $\chi^2$ , and our proposed *MD* and *MD* $\chi^2$  for the data set of REUTERS-20, when SVM is used as the classifier.

We also test our proposed feature selection approaches compared with the previously existing fea-

ture selection methods, when SVM is used as classifier for text categorization. Fig. 5 shows the classification results on three REUTERS data sets for SVM, in which the performance improvement of our proposed two approaches can be also seen. It shows that our proposed two approaches perform at least as well as previously existing methods at a very small feature size, and are consistently better when the feature size increases. Particularly, in Fig. 6, we show the JMH-divergence with different number of features that are selected by different feature selection methods for the data set of REUTERS-20, when SVM is used as the classifier. Notice that, since the JMH-divergence reflects the discriminative information of features for classification, the results in Fig. 6 can somehow explain the classification performance in accuracy as shown in Fig. 5 (b-1). Moreover, in Fig. 6, the JMH-divergence of our proposed *MD* method



(a). Accuracy for TDT2-10



(b). F1 measure for TDT2-10

Fig. 4: Performance comparisons of feature selection methods on the TDT2 data set: (a) accuracy, (b) F1 measure, when naive Bayes is used as the classifier.

is monotonically increasing when more features are selected, which can be guaranteed by Theorem 1.

## 6 CONCLUSIONS AND FUTURE WORKS

We have introduced new feature selection approaches based on the information measures for naive Bayes classifiers, aiming to select the features that offer the maximum discriminative capacity for text classification. We have also derived the asymptotic distributions of these measures, which leads to the other version of the Chi-square statistic approach for feature selection. Compared with the existing feature selection approaches that rank the features by only exploring the intrinsic characteristics of data without considering the learning algorithm for classification, our proposed approaches involve the learning model in the feature filtering process, which provides us a theoretical way to analyze the optimality of the selected features. The experiments we have conducted on several benchmarks have demonstrated their promising performance improvement compared with the previously existing feature selection approaches.

For future work, we will analyze the feature dependence and develop feature selection algorithms by weighting each individual features [41] [42], aiming to maximize the discriminative capacity. Furthermore, we will incorporate our feature selection approaches into other advanced machine learning algorithms

such as imbalanced learning [43] [44] and partial learning model [45] [46] to enhance the learning for rare categories.

## ACKNOWLEDGMENT

This research was partially supported by National Science Foundation (NSF) under grant ECCS 1053717 and CCF 1439011, and the Army Research Office under grant W911NF-12-1-0378.

## REFERENCES

- [1] T. Joachims, "Text categorization with support vector machines: Learning with many relevant features," in *European Conference on Machine Learning*. Springer, 1998.
- [2] W. Lam, M. Ruiz, and P. Srinivasan, "Automatic text categorization and its application to text retrieval," *IEEE Transactions on Knowledge and Data Engineering*, vol. 11, no. 6, pp. 865–879, 1999.
- [3] F. Sebastiani, "Machine learning in automated text categorization," *ACM computing surveys (CSUR)*, vol. 34, no. 1, pp. 1–47, 2002.
- [4] H. Al-Mubaid and S. Umair, "A new text categorization technique using distributional clustering and learning logic," *IEEE Transactions on Knowledge and Data Engineering*, vol. 18, no. 9, pp. 1156–1165, 2006.
- [5] Y. Aphinyanaphongs, L. D. Fu, Z. Li, E. R. Peskin, E. Efstathiadis, C. F. Aliferis, and A. Statnikov, "A comprehensive empirical comparison of modern supervised classification and feature selection methods for text categorization," *Journal of the Association for Information Science and Technology*, vol. 65, no. 10, pp. 1964–1987, 2014.
- [6] H. Liu and L. Yu, "Toward integrating feature selection algorithms for classification and clustering," *IEEE Transactions on Knowledge and Data Engineering*, vol. 17, no. 4, pp. 491–502, 2005.
- [7] G. Salton and C. Buckley, "Term-weighting approaches in automatic text retrieval," *Information Processing and Management*, vol. 24, no. 5, pp. 513–523, 1988.
- [8] N. Gövert, M. Lalmas, and N. Fuhr, "A probabilistic description-oriented approach for categorizing web documents," in *International Conference on Information and Knowledge Management*, 1999, pp. 475–482.
- [9] A. Mnih and G. E. Hinton, "A scalable hierarchical distributed language model," in *Advances in Neural Information Processing Systems*, 2009, pp. 1081–1088.
- [10] J. Turian, L. Ratinov, and Y. Bengio, "Word representations: a simple and general method for semi-supervised learning," in *Association for Computational Linguistics*, 2010, pp. 384–394.
- [11] T. Mikolov, K. Chen, G. Corrado, and J. Dean, "Efficient estimation of word representations in vector space," *arXiv preprint arXiv:1301.3781*, 2013.
- [12] Q. V. Le and T. Mikolov, "Distributed representations of sentences and documents," *arXiv preprint arXiv:1405.4053*, 2014.
- [13] D. D. Lewis, "Naive (Bayes) at forty: The independence assumption in information retrieval," in *European Conference on Machine Learning*, 1998, pp. 4–15.
- [14] D. Koller and M. Sahami, "Hierarchically classifying documents using very few words," *International Conference on Machine Learning*, pp. 170–178, 1997.
- [15] Y. H. Li and A. K. Jain, "Classification of text documents," *The Computer Journal*, vol. 41, no. 8, pp. 537–546, 1998.
- [16] G. Forman, "An extensive empirical study of feature selection metrics for text classification," *The Journal of Machine Learning Research*, vol. 3, pp. 1289–1305, 2003.
- [17] Y. Yang and J. O. Pedersen, "A comparative study on feature selection in text categorization," in *International Conference on Machine Learning*, vol. 97, 1997, pp. 412–420.
- [18] A. Genkin, D. D. Lewis, and D. Madigan, "Large-scale bayesian logistic regression for text categorization," *Technometrics*, vol. 49, no. 3, pp. 291–304, 2007.

[19] B. Tang and H. He, "ENN: Extended nearest neighbor method for pattern recognition [research frontier]," *IEEE Computational Intelligence Magazine*, vol. 10, no. 3, pp. 52–60, 2015.

[20] A. McCallum and K. Nigam, "A comparison of event models for naive bayes text classification," in *AAAI-98 Workshop on Learning for Text Categorization*, vol. 752, 1998, pp. 41–48.

[21] S. Eyheramendy, D. D. Lewis, and D. Madigan, "On the Naive Bayes model for text categorization," *Proceedings of the Ninth International Workshop on Artificial Intelligence and Statistics*, pp. 332–339, 2003.

[22] V. Metsis, I. Androutsopoulos, and G. Paliouras, "Spam filtering with naive bayes-which naive bayes?" in *Conference on Email and Anti-Spam*, 2006, pp. 27–28.

[23] V. N. Vapnik, *Estimation of dependences based on empirical data*. Springer-verlag New York, vol. 40.

[24] R. Kohavi and G. H. John, "Wrappers for feature subset selection," *Artificial intelligence*, vol. 97, no. 1, pp. 273–324, 1997.

[25] G. Salton, A. Wong, and C.-S. Yang, "A vector space model for automatic indexing," *Communications of the ACM*, vol. 18, no. 11, pp. 613–620, 1975.

[26] F. Bouillout, P. N. Hai, N. Béchet, S. Bringay, D. Ienco, S. Matwin, P. Poncelet, M. Roche, and M. Teisseire, "How to extract relevant knowledge from tweets?" in *Information Search, Integration and Personalization*, 2013, pp. 111–120.

[27] D. Mladenic and M. Grobelnik, "Feature selection for unbalanced class distribution and naive bayes," in *International Conference on Machine Learning*, vol. 99, 1999, pp. 258–267.

[28] M. F. Caropreso, S. Matwin, and F. Sebastiani, "A learner-independent evaluation of the usefulness of statistical phrases for automated text categorization," *Text Databases and Document Management: Theory and Practice*, pp. 78–102, 2001.

[29] L. Galavotti, F. Sebastiani, and M. Simi, "Experiments on the use of feature selection and negative evidence in automated text categorization," in *Research and Advanced Technology for Digital Libraries*, 2000, pp. 59–68.

[30] S. Kullback, *Information theory and statistics*. Courier Corporation, 1997.

[31] T. M. Cover and J. A. Thomas, *Elements of information theory*. John Wiley & Sons, 2012.

[32] H. Chernoff, "A measure of asymptotic efficiency for tests of a hypothesis based on the sum of observations," *The Annals of Mathematical Statistics*, pp. 493–507, 1952.

[33] H. Jeffreys, "An invariant form for the prior probability in estimation problems," *Proceedings of the Royal Society of London. Series A. Mathematical and Physical Sciences*, vol. 186, no. 1007, pp. 453–461, 1946.

[34] J. Lin, "Divergence measures based on the shannon entropy," *IEEE Transactions on Information Theory*, vol. 37, no. 1, pp. 145–151, 1991.

[35] K. Sawyer, "A multiple divergence criterion for testing between separate hypotheses," *Statistics & Probability Letters*, vol. 1, no. 1, pp. 26–30, 1982.

[36] R. Rifkin and A. Klautau, "In defense of one-vs-all classification," *The Journal of Machine Learning Research*, vol. 5, pp. 101–141, 2004.

[37] D. Cai, X. He, and J. Han, "Document clustering using locality preserving indexing," *IEEE Transactions on Knowledge and Data Engineering*, vol. 17, no. 12, pp. 1624–1637, 2005.

[38] D. Cai, Q. Mei, J. Han, and C. Zhai, "Modeling hidden topics on document manifold," in *ACM Conference on Information and Knowledge Management*, 2008, pp. 911–920.

[39] R. Liere and P. Tadepalli, "Active learning with committees for text categorization," in *Proceedings of AAAI-97, 14th Conference of the American Association for Artificial Intelligence*, 1997, pp. 591–596.

[40] E. F. Combarro, E. Montanes, I. Diaz, J. Ranilla, and R. Mones, "Introducing a family of linear measures for feature selection in text categorization," *IEEE Transactions on Knowledge and Data Engineering*, vol. 17, no. 9, pp. 1223–1232, 2005.

[41] X.-B. Xue and Z.-H. Zhou, "Distributional features for text categorization," *IEEE Transactions on Knowledge and Data Engineering*, vol. 21, no. 3, pp. 428–442, 2009.

[42] S.-B. Kim, K.-S. Han, H.-C. Rim, and S. H. Myaeng, "Some effective techniques for naive bayes text classification," *IEEE Transactions on Knowledge and Data Engineering*, vol. 18, no. 11, pp. 1457–1466, 2006.

[43] B. Tang and H. He, "KernelADASYN: Kernel based adaptive synthetic data generation for imbalanced learning," in *IEEE Congress on Evolutionary Computation (CEC)*, 2015, pp. 664–671.

[44] H. He and E. Garcia, "Learning from imbalanced data," *IEEE Transactions on Knowledge and Data Engineering*, vol. 21, no. 9, pp. 1263–1284, 2009.

[45] B. Tang, Q. Ding, H. He, and S. Kay, "Hybrid classification with partial models," in *International Joint Conference on Neural Networks*, 2014, pp. 3726–3731.

[46] B. Tang, H. He, Q. Ding, and S. Kay, "A parametric classification rule based on the exponentially embedded family," *IEEE Transactions on Neural Networks and Learning Systems*, vol. 26, no. 2, pp. 367–377, 2015.



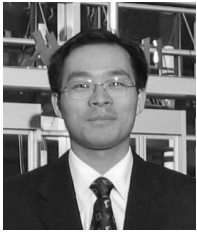
**Bo Tang** received the B.S. degree from the Central South University (CSU), Changsha, China, in 2007, and the M.S. degree from the Institute of Electronics, Chinese Academy of Science (IECAS), Beijing, China, in 2010. Currently, he is pursuing the Ph.D. degree in the Department of Electrical, Computer and Bio-medical Engineering, University of Rhode Island, Kingston, RI, USA. His current research interests include statistical machine learning, data mining, computational intelligence, computer vision, and Cloud computing.



**Steven Kay** was born in Newark, NJ, on April 5, 1951. He received the B.E. degree from Stevens Institute of Technology, Hoboken, NJ in 1972, the M.S. degree from Columbia University, New York, NY, in 1973, and the Ph.D. degree from Georgia Institute of Technology, Atlanta, GA, in 1980, all in electrical engineering.

From 1972 to 1975, he was with Bell Laboratories, Holmdel, NJ, where he was involved with transmission planning for speech communications and simulation and subjective testing of speech processing algorithms. From 1975 to 1977, he attended Georgia Institute of Technology to study communication theory and digital signal processing. From 1977 to 1980, he was with the Submarine Signal Division, Portsmouth, RI, where he engaged in research on autoregressive spectral estimation and the design of sonar systems. He is presently a Professor of Electrical Engineering at the University of Rhode Island, Kingston, and a consultant to numerous industrial concerns, the Air Force, the Army, and the Navy. As a leading expert in statistical signal processing, he has been invited to teach short courses to scientists and engineers at government laboratories, including NASA and the CIA. He has written numerous journal and conference papers and is a contributor to several edited books. He is the author of the textbooks *Modern Spectral Estimation* (Prentice-Hall, 1988), *Fundamentals of Statistical Signal Processing, Vol. I: Estimation Theory* (Prentice-Hall, 1993), *Fundamentals of Statistical Signal Processing, Vol. II: Detection Theory* (Prentice-Hall, 1998), and *Intuitive Probability and Random Processes using MATLAB* (Springer, 2006). His current interests are spectrum analysis, detection and estimation theory, and statistical signal processing.

Dr. Kay is a Fellow of the IEEE, and a member of Tau Beta Pi and Sigma Xi. He has been a distinguished lecturer for the IEEE signal processing society. He has been an associate editor for the IEEE Signal Processing Letters and the IEEE Transactions on Signal Processing. He has received



**Haibo He (SM'11)** received the B.S. and M.S. degrees in electrical engineering from Huazhong University of Science and Technology, China, in 1999 and 2002, respectively, and the Ph.D. degree in electrical engineering from Ohio University in 2006. He is currently the Robert Haas Endowed Professor and Professor in Electrical Engineering at the University of Rhode Island. From 2006 to 2009, he was an Assistant Professor at the Department of Electrical and Computer

Engineering, Stevens Institute of Technology.

His research interests include machine learning, cyber-physical systems, computational intelligence, hardware design for machine intelligence, and various applications such as smart grid and renewable energy systems. He has published one research book (Wiley), edited 1 research book (Wiley-IEEE) and 6 conference proceedings (Springer), and authored and co-authored over 150 peer-reviewed journal and conference papers. His researches have been covered by national and international media such as IEEE Smart Grid Newsletter, The Wall Street Journal, and Providence Business News. Currently, he is an Associate Editor of the IEEE Transactions on Neural Networks and Learning Systems, and IEEE Transactions on Smart Grid. He was a recipient of the IEEE International Conference on Communications Best Paper award (2014), IEEE Computational Intelligence Society (CIS) Outstanding Early Career Award (2014), National Science Foundation (NSF) CAREER Award (2011), and Providence Business News (PBN) Rising Star Innovator Award (2011).