

# Feature Augmentation via Nonparametrics and Selection (FANS) in High-Dimensional Classification

Jianqing FAN, Yang FENG, Jiancheng JIANG, and Xin TONG

We propose a high-dimensional classification method that involves nonparametric feature augmentation. Knowing that marginal density ratios are the most powerful univariate classifiers, we use the ratio estimates to transform the original feature measurements. Subsequently, penalized logistic regression is invoked, taking as input the newly transformed or augmented features. This procedure trains models equipped with local complexity and global simplicity, thereby avoiding the curse of dimensionality while creating a flexible nonlinear decision boundary. The resulting method is called feature augmentation via nonparametrics and selection (FANS). We motivate FANS by generalizing the naive Bayes model, writing the log ratio of joint densities as a linear combination of those of marginal densities. It is related to generalized additive models, but has better interpretability and computability. Risk bounds are developed for FANS. In numerical analysis, FANS is compared with competing methods, so as to provide a guideline on its best application domain. Real data analysis demonstrates that FANS performs very competitively on benchmark email spam and gene expression datasets. Moreover, FANS is implemented by an extremely fast algorithm through parallel computing.

**KEY WORDS:** Classification; Density estimation; Feature augmentation; Feature selection; High-dimensional space; Nonlinear decision boundary; Parallel computing.

## 1. INTRODUCTION

Classification aims to identify to which category a new observation belongs based on feature measurements. Numerous applications include spam detection, image recognition, and disease classification (using high-throughput data such as microarray gene expression and SNPs). Well-known classification methods include Fisher's linear discriminant analysis (LDA), logistic regression, naive Bayes (NB),  $k$ -nearest neighbors, neural networks, and many others. All these methods can perform well in the classical low-dimensional settings, in which the number of features is much smaller than the sample size. However, in many contemporary applications, the number of features  $p$  is large compared to the sample size  $n$ . For instance, the dimensionality  $p$  in microarray data is frequently in thousands or beyond, while the sample size  $n$  is typically in the order of tens. Besides computational issues, the central conflict in high-dimensional setup is that the model complexity is not supported by limited access to data. In other words, the "variance" of conventional models is high in such new settings, and even simple models such as LDA need to be regularized. We refer to Hastie, Tibshirani, and Friedman (2009) and Bühlmann and van de Geer (2011) for overviews of statistical challenges associated with high dimensionality.

In this article, we propose a classification procedure feature augmentation via nonparametrics and selection (FANS). Before

introducing the algorithm, we first detail its motivation. Suppose feature measurements and responses are coded by a pair of random variables  $(X, Y)$ , where  $X \in \mathcal{X} \subset \mathbb{R}^p$  denotes the features and  $Y \in \{0, 1\}$  is the binary response. Recall that a classifier  $h$  is a data-dependent mapping from the feature space to the labels. Classifiers are usually constructed to minimize the risk  $P(h(X) \neq Y)$ .

Denote by  $g$  and  $f$  the class conditional densities, respectively, for class 0 and class 1, that is,  $(X|Y=0) \sim g$  and  $(X|Y=1) \sim f$ . It can be shown that the Bayes rule is  $\mathbb{I}(r(x) \geq 1/2)$ , where  $r(x) = E(Y|X=x)$ . Let  $\pi = P(Y=1)$ , then

$$r(x) = \frac{\pi f(x)}{\pi f(x) + (1-\pi)g(x)}.$$

Assume for simplicity that  $\pi = 1/2$ , then the oracle decision boundary is

$$\{x : f(x)/g(x) = 1\} = \{x : \log f(x) - \log g(x) = 0\}.$$

Denote by  $g_1, \dots, g_p$  the marginals of  $g$ , and by  $f_1, \dots, f_p$  those of  $f$ . NB models assume that the conditional distributions of each feature given the class labels are independent, that is,

$$\log \frac{f(x)}{g(x)} = \sum_{j=1}^p \log \frac{f_j(x_j)}{g_j(x_j)}. \quad (1.1)$$

NB is a simple approach, but it is useful in many high-dimensional settings. Taking a two-class Gaussian model with a common covariance matrix, Bickel and Levina (2004) showed that naively carrying out the Fisher's discriminant rule performs poorly due to diverging spectra. In addition, the authors argued that independence rule which ignores the covariance structure performs better than the Fisher's rule in some high-dimensional settings. However, correlation among features is usually an essential characteristic of data, and it can help classification under

Jianqing Fan is Frederick L. Moore Professor of Finance, Department of Operations Research and Financial Engineering, Princeton University, Princeton, NJ 08544 (E-mail: [jqfan@princeton.edu](mailto:jqfan@princeton.edu)). Yang Feng is Assistant Professor, Department of Statistics, Columbia University, New York, NY 10027 (E-mail: [yangfeng@stat.columbia.edu](mailto:yangfeng@stat.columbia.edu)). Jiancheng Jiang is Associate Professor, Department of Mathematics and Statistics, University of North Carolina at Charlotte, Charlotte, NC 28223 (E-mail: [jjiang1@uncc.edu](mailto:jjiang1@uncc.edu)). Xin Tong is Assistant Professor, Department of Data Sciences and Operations, University of Southern California, Los Angeles, CA 90089 (E-mail: [xint@marshall.usc.edu](mailto:xint@marshall.usc.edu)). The financial support from National Institutes of Health grants R01-GM072611 and R01GM100474-01 and National Science Foundation grants DMS-1206464 and DMS-1308566 is greatly acknowledged. The authors thank the editor, the associate editor, and referees for their constructive comments.

Color versions of one or more of the figures in the article can be found online at [www.tandfonline.com/r/jasa](http://www.tandfonline.com/r/jasa).

© 2016 American Statistical Association  
Journal of the American Statistical Association  
March 2016, Vol. 111, No. 513, Theory and Methods  
DOI: 10.1080/01621459.2015.1005212

suitable models and with relative abundance of the sample. Examples in bioinformatics study can be found in Ackermann and Strimmer (2009). Recently, Fan, Feng, and Tong (2012) showed that the independence assumption can lead to huge loss in classification power when correlation prevails and proposed a regularized optimal affine discriminant (ROAD). ROAD is a linear plug-in rule targeting directly on the classification error, and it takes advantages of the unregularized pooled sample covariance matrix.

Relaxing the two-class Gaussian assumption in parametric NB gives us a general NB formulation such as (1.1). However, this model also fails to capture the correlation or dependence among features in general. On the other hand, the marginal density ratios are the most powerful univariate classifiers, and using them as features in multivariate classifiers can yield very powerful procedures. This consideration motivates us to ask the following question: Are there advantages of combining these transformed features rather than untransformed feature? More precisely, we would like to learn a decision boundary from the following set

$$\mathcal{D} = \left\{ \mathbf{x} : \beta_0 + \beta_1 \log \frac{f_1(x_1)}{g_1(x_1)} + \dots + \beta_p \log \frac{f_p(x_p)}{g_p(x_p)} = 0, \beta_0, \dots, \beta_p \in \mathbb{R} \right\}. \quad (1.2)$$

(All coefficients are 1 in the NB model, so optimization is not necessary.) For univariate problems, properly thresholding the marginal density ratio delivers the best classifier. In this sense, the marginal density ratios can be regarded as the best transforms of future measurements, and (1.2) is an effort toward combining those most powerful univariate transforms to build more powerful classifiers.

This is in a similar spirit to the sure independence screening (SIS) in Fan and Lv (2008) where the best marginal predictors are used as probes for their utilities in the joint model. By combining these marginal density ratios and optimizing over their coefficients  $\beta_j$ 's, we wish to build a good classifier that takes into account feature dependence. Note that although our target boundary  $\mathcal{D}$  is not linear in the original features, it is linear in the parameters  $\beta_j$ 's. Therefore, any linear classifiers can be applied to the transformed variables. For example, we can use logistic regression, one of the most popular linear classification rules. Other choices, such as SVM (linear kernel), are good alternatives, but we choose logistic regression for the rest of discussion.

Recall that logistic regression models the log odds by

$$\log \frac{P(Y = 1 | \mathbf{X} = \mathbf{x})}{P(Y = 0 | \mathbf{X} = \mathbf{x})} = \beta_0 + \sum_{j=1}^p \beta_j x_j,$$

where the  $\beta_j$ 's are estimated by the maximum likelihood approach. We should note that without explicitly modeling correlations, logistic regression takes into account features' joint effects and levels a good linear combination of features as the decision boundary. Its performance is similar to LDA, but both models can only capture decision boundaries linear in original features.

On the other hand, logistic regression might serve as a building block for the more flexible FANS algorithm. Concretely,

if we know the marginal densities  $f_j$  and  $g_j$ , and run logistic regression on the transformed features  $\{\log(f_j(x_j)/g_j(x_j))\}$ , we create a decision boundary nonlinear in the original features. The use of these transformed features is easily interpretable: one naturally combines the "most powerful" univariate transforms (building blocks of univariate Bayes rules)  $\{\log(f_j(x_j)/g_j(x_j))\}$  rather than the original measurements. In special cases such as the two-class Gaussian model with a common covariance matrix, the transformed features are not different from the original ones. Some caution should be taken: if  $f_j = g_j$  for some  $j$ , that is, the marginal densities for feature  $j$  are exactly the same, this feature will not have any contribution in classification. Deletion like this might lose power, because features having no marginal contribution on their own might boost classification performance when they are used jointly with other features. In view of this defect, a variant of FANS augments the transformed features with the original ones.

Since marginal densities  $f_j$  and  $g_j$  are unknown, we need to first estimate them, and then run a penalized logistic regression (PLR) on the estimated transforms. Note that some regularization (e.g., penalization) is necessary to reduce model complexity in the high-dimensional paradigm. This two-step classification rule of feature augmentation via nonparametrics and selection will be called FANS for short. Precise algorithmic implementation of FANS is described in the next section. Numerical results show that our new method excels in many scenarios, in particular when no linear decision boundary can separate the data well.

To understand where FANS stands compared to NB, PLR, and the ROAD, we showcase a simple simulation example. In this example, the choice is between a multivariate Gaussian distribution and some componentwise mixture of two multivariate Gaussian distributions:

Class 0:  $N((5 \times \mathbf{1}_{10}^T, \mathbf{0}_{p-10}^T)^T, \Sigma)$ ,

Class 1:  $\mathbf{w} \circ N(\mathbf{0}_p, \mathbf{I}_p) + (1 - \mathbf{w}) \circ N((6 \times \mathbf{1}_{10}^T, \mathbf{0}_{p-10}^T)^T, \Sigma)$ , where  $p = 1000$ ,  $\circ$  is the element-wise product between matrices,  $\Sigma_{ii} = 1$  for all  $i = 1, \dots, p$ ,  $\Sigma_{ij} = 0.5$  for all  $i, j = 1, \dots, p$  and  $i \neq j$ , and  $\mathbf{w} = (w_1, \dots, w_p)^T$ , in which  $w_j \stackrel{\text{iid}}{\sim} \text{Bernoulli}(0.5)$ .

The median classification error with the standard error shown in the error bar for 100 repetitions as a function of training sample size  $n$  is rendered in Figure 1. This figure suggests that increasing the sample size does not help NB boost performance (in terms of the median classification error), because the NB model is severely biased in view of significant correlation presence. It is interesting to compare PLR with ROAD. ROAD is a more efficient approach when the sample size is small; however, PLR eventually performs better when the sample size becomes large enough. This is not surprising because the underlying true model is not two-class Gaussian with a common covariance matrix. So the less "biased" PLR beats ROAD on larger samples. Nevertheless, even if ROAD uses a misspecified model, it still benefits from a specific model assumption on small samples. Finally, since the oracle decision boundary in this example is nonlinear, the newly proposed FANS approach performs significantly better than others when the sample size is reasonably large. The above analysis seems to suggest that FANS does well as long as we have enough data to construct accurate marginal density estimates. Note also that ROAD is better than FANS




Figure 1. The median test errors for Gaussian versus mixture of Gaussian when the training data size varies. Standard errors shown in the error bars.

when the training sample size is extremely small. Figure 1 shows that even under the same data distribution, the best method in practice largely depends on the available sample abundance.

A popular extension of logistic regression and close relative to FANS is the additive logistic regression, which belongs to the generalized additive model (Hastie and Tibshirani 1990). Additive logistic regression allows (smooth) nonparametric feature transformations to appear in the decision boundary by modeling

$$\log \frac{P(Y = 1|X = \mathbf{x})}{P(Y = 0|X = \mathbf{x})} = \sum_{j=1}^p h_j(x_j), \quad (1.3)$$

where  $h_j$ 's are smooth functions. This kind of additive decision boundary is very general, in which FANS and logistic regression are special cases. It works well for small- $p$ -large- $n$  scenarios, while its penalized versions adapt to high-dimensional settings. We will compare FANS with penalized additive logistic regression in numerical studies. The major drawback of additive logistic regression (generalized additive model) is the heavy computational complexity (e.g., the backfitting algorithm) involved in searching the transformation functions  $h_j(\cdot)$ . Moreover, the available algorithms, for example, the algorithm for penGAM (Meier, Geer, and Bühlmann 2009), fail to give an estimate when the sample size is very small. Compared to FANS, the generalized additive model uses a factor of  $K_n$  more parameters, where

$K_n$  is the number of knots in the approximation of every additive components  $\{h_j(\cdot)\}_{j=1}^p$ . While this reduces possible biases in comparison with FANS, it increases variances in the estimation and results in more computation cost (see Table 2). Moreover, FANS admits a nice interpretation of optimal combination of optimal building blocks for univariate classifiers.

Besides the aforementioned references, there is a huge literature on high-dimensional classification. Examples include principal component analysis in Bair et al. (2006) and Zou, Hastie, and Tibshirani (2006), partial least squares in Nguyen and Rocke (2002), Huang (2003), and Boulesteix (2004), and sliced inverse regression in Li (1991) and Antoniadis, Lambert-Lacroix, and Leblanc (2003). Recently, there has been a surge of interest for extending the LDA to high-dimensional settings including Guo, Hastie, and Tibshirani (2007), Wu et al. (2009), Clemmensen et al. (2011), Shao et al. (2011), Cai and Liu (2011), Mai, Zou, and Yuan (2012), and Witten and Tibshirani (2012).

The rest of the article is organized as follows. Section 2 introduces the FANS algorithm. Section 3 is dedicated to simulation studies and real data analysis. Theoretical results are presented in Section 4. We conclude with a discussion in Section 5. Longer proofs and technical results are relegated to the Appendix.

## 2. ALGORITHM

In this section, an efficient algorithm (S1–S5) for FANS will be introduced. We will also describe a variant of FANS (FANS2), which uses the original features in addition to the transformed ones.

### 2.1 FANS and its Running Time Bound

- S1. Given  $n$  pairs of observations  $D = \{(X_i, Y_i), i = 1, \dots, n\}$ . Randomly split the data into two parts for  $L$  times:  $D_l = (D_{l1}, D_{l2}), l = 1, \dots, L$ .
- S2. On each  $D_{l1}, l \in \{1, \dots, L\}$ , apply kernel density estimation and denote the estimates by  $\hat{f} = (\hat{f}_1, \dots, \hat{f}_p)^T$  and  $\hat{g} = (\hat{g}_1, \dots, \hat{g}_p)^T$ .
- S3. Calculate the transformed observations  $\hat{Z}_i = \mathbf{Z}_{\hat{f}, \hat{g}}(X_i)$ , where  $\hat{Z}_{ij} = \log \hat{f}_j(X_{ij}) - \log \hat{g}_j(X_{ij})$ , for each  $i \in D_{l2}$  and  $j \in \{1, \dots, p\}$ .
- S4. Fit an  $L_1$ -PLR to the transformed data  $\{(\hat{Z}_i, Y_i), i \in D_{l2}\}$ , using cross-validation to get the best penalty parameter. For a new observation  $\mathbf{x}$ , we estimate transformed features by  $\log \hat{f}_j(x_j) - \log \hat{g}_j(x_j), j = 1, \dots, p$ , and plug them into the fitted logistic function to get the predicted probability  $p_l$ .

Table 1. Median test error (in percentage) for the simulation examples

Ex( $\rho$ )	FANS	FANS2	ROAD	PLR	penGAM	NB	FAIR	SVM
1(0)	6.8 (1.1)	6.2 (1.2)	6.0 (1.3)	6.5 (1.2)	6.6 (1.1)	11.2 (1.4)	5.7 (1.0)	13.2 (1.5)
1(0.5)	16.5 (1.7)	16.2 (1.8)	16.5 (5.3)	15.9 (1.7)	16.9 (1.6)	20.6 (1.7)	17.2 (1.6)	22.5 (1.8)
2(0.5)	4.2 (0.9)	2.0 (0.6)	2.0 (0.6)	2.5 (0.6)	3.7 (0.9)	43.5 (11.1)	25.3 (1.6)	5.3 (1.1)
2(0.9)	3.1 (1.1)	0.0 (0.0)	0.0 (0.0)	0.0 (0.0)	0.2 (1.4)	46.8 (8.8)	30.2 (1.9)	0.0 (0.1)
3(0)	0.0 (0.0)	0.0 (0.0)	49.6 (2.4)	50.0 (1.3)	0.0 (0.1)	50.4 (2.2)	50.2 (2.1)	31.8 (2.4)
3(0.5)	3.4 (0.7)	3.4 (0.7)	49.3 (2.4)	50.0 (1.3)	3.7 (0.8)	50.0 (2.1)	50.2 (2.0)	19.8 (2.4)
4	0.0 (0.0)	0.0 (0.0)	28.2 (1.8)	50.0 (10.7)	0.0 (0.0)	41.0 (1.1)	34.6 (1.4)	0.0 (0.0)

Note: Standard errors are in the parentheses.

Table 2. Computation time (in seconds) comparison for FANS, SVM, ROAD, and penGAM

Ex( $\rho$ )	FANS	FANS(para)	SVM	ROAD	penGAM
1(0)	12.0 (2.6)	3.8 (0.2)	59.4 (12.8)	99.1 (98.2)	243.7 (151.8)
1(0.5)	12.7 (2.1)	3.5 (0.2)	81.3 (19.2)	100.7 (89.3)	325.8 (194.3)
2(0.5)	16.0 (3.1)	4.0 (0.2)	77.6 (18.1)	106.8 (90.7)	978.0 (685.7)
2(0.9)	22.0 (4.6)	4.5 (0.3)	75.7 (17.8)	98.3 (83.9)	3451.1 (3040.2)
3(0)	12.1 (2.1)	3.4 (0.2)	152.1 (27.4)	96.3 (68.8)	254.6 (130.0)
3(0.5)	11.9 (2.0)	3.4 (0.2)	342.1 (58.0)	95.9 (74.8)	298.7 (167.4)
4	22.4 (3.9)	6.6 (0.4)	264.3 (45.0)	75.1 (54.0)	4811.9 (3991.7)

Note: The parallel computing technique is applied. Standard errors are in the parentheses.

S5. Repeat S2–S4 for  $l = 1, \dots, L$ , use the average predicted probability  $\text{prob} = L^{-1} \sum_{l=1}^L p_l$  as the final prediction, and assign the observation  $\mathbf{x}$  to class 1 if  $\text{prob} \geq 1/2$ , and 0 otherwise.

A few comments on the technical implementation are made as follows.

*Remark 1.*

- (i) In S2, if an estimated marginal density is less than some threshold  $\varepsilon$  (say  $10^{-2}$ ), we set it to be  $\varepsilon$ . This Winsorization increases the stability of the transformations because the estimated transformations  $\log \hat{f}_j$  and  $\log \hat{g}_j$  are unstable in regions where true densities are low.
- (ii) In S4, we take PLR, but any linear classifier can be used. For example, support vector machine (SVM) with linear kernel is also a good choice.
- (iii) In S4, the  $L_1$  penalty (Tibshirani 1996) was adopted since our primary interest is the classification error. We can also apply other penalty functions, such as SCAD (Fan and Li 2001), adaptive LASSO (Zou 2006), and MCP (Zhang 2010).
- (iv) In S5, the average predicted probability is taken as the final prediction. An alternative approach is to make a decision on each random split and listen to majority vote.

In S1, we split the data multiple times. The rationale behind multiple splitting lies in the two-step prototype nature of FANS, which uses the first part of the data for marginal nonparametric density estimates (in S2) and (transformation of) the second part for PLR (in S4). Multiple splitting and prediction averaging not only make our procedure more robust against arbitrary assignments of data usage, but also make more efficient use of limited data. This idea is related to random forest (Breiman 2001), where the final prediction is the average over results from multiple bootstrap samples. Other related literature includes Fu, Carroll, and Wang (2005) which considers estimation of misclassification error with small samples via bootstrap cross-validation. The number of splits is fixed at  $L = 20$  throughout all numerical studies. This choice reflects our cluster's node number. Interested readers can as well leverage their better computing resources for a larger  $L$ . However, we observed that further increasing  $L$  leads to similar performance for all simulation examples. Also, we recommend a balanced assignment by switching the role of data used for feature transformation and

for feature selection, that is,  $D_{2l} = (D_{(2l-1),2}, D_{(2l-1),1})$  when  $D_{2l-1} = (D_{(2l-1),1}, D_{(2l-1),2})$ .

It is straightforward to derive a running time bound for our algorithm. Suppose splitting has been done. In S2, we need to perform kernel density estimation for each variable, which costs  $O(n^2 p)^1$ . The transformations in S3 cost  $O(np)$ . In S4, we call the R package `glmnet` to implement PLR, which employs the coordinate descent algorithm for each penalty level. This step has a computational cost at most  $O(npT)$ , where  $T$  is the number of penalty levels, that is, the number of times the coordinate descent algorithm is run (see Friedman et al. 2007 for a detailed analysis). The default setting is  $T = 100$ , though we can set it to other constants. Therefore, a running time bound for the whole algorithm is  $O(L(n^2 p + np + npT)) = O(Lnp(n + T))$ .

The above bound does not look particularly interesting. However, smart implementation of the FANS procedure can fully unleash the potential of our algorithm. Indeed, not only the  $L$  repetitions, but also the marginal density estimates in S2 can be done via parallel computing. Suppose  $L$  is the number of available nodes, and the cpu core number in each node is  $N \geq n/T$ . This assumption is reasonable because  $T = 100$  by default,  $N = 8$  for our implementation, and sample sizes  $n$  for many applications are less than a multiple of  $TN$ . Under this assumption, the  $L$  predicted probabilities calculations can be carried out simultaneously and the results are combined later in S5. Moreover in S2, the running time bound becomes  $O(n^2 p/N)$ . Henceforth, a bound for the whole algorithm will be  $O(npT)$ , which is the same as that for PLR. The exciting message here is that, by leveraging modern computer architecture, we are able to implement our nonparametric classification rule FANS within running time at the order of a parametric method. The computation times for various simulation setups are reported in Table 2, where the first column reports results when only  $L$  repetitions are paralleled, and the second column reports the improvement when marginal density estimates in S2 are paralleled within each node.

## 2.2 Augmenting Linear Features

As we argued in the introduction, features with no marginal discrimination power do not make contribution in FANS. One remedy is to run (in S4) the PLR using both the transformed features and the original ones, which amounts to modeling the

<sup>1</sup>Approximate kernel density estimates can be computed faster, see, for example, Raykar, Duraiswami, and Zhao (2010).

log odds by

$$\beta_0 + \beta_1 \log \frac{f_1(x_1)}{g_1(x_1)} + \cdots + \beta_p \log \frac{f_p(x_p)}{g_p(x_p)} + \beta_{p+1} x_1 + \cdots + \beta_{2p} x_p.$$

This variant of FANS is named FANS2, and it allows features with no marginal power to enter the model in a linear fashion. FANS2 helps when a linear decision boundary separates data reasonably well.

### 3. NUMERICAL STUDIES

#### 3.1 Simulation

In simulation studies, FANS and FANS2 are compared with competing methods: PLR (Friedman, Hastie, and Tibshirani 2010), penalized additive logistic regression models (penGAM, Meier, Geer, and Bühlmann 2009), SVM, ROAD (Fan, Feng, and Tong 2012), LDA, NB, and feature annealed independence rule (FAIR, Fan and Fan 2008).

In all simulation settings, we set  $p = 1000$  and training and testing data sample sizes of each class to be 300. Five-fold cross-validation is conducted when needed, and we repeat 50 times for each setting (the relative small number of replications is due to the long computation time of penGAM, see Table 2). Table 1 summarizes median test errors for each method along with the corresponding standard errors. This table omits Fisher's classifier (using pseudo inverse for sample covariance matrix) because it gives a test error around 50%, equivalent to random guessing.

*Example 1.* We consider the two-class Gaussian settings where  $\Sigma_{ii} = 1$  for all  $i = 1, \dots, p$  and  $\Sigma_{ij} = \rho^{|i-j|}$ ,  $\boldsymbol{\mu}_1 = \mathbf{0}_{1000}$  and  $\boldsymbol{\mu}_2 = (\mathbf{1}_{10}^T, \mathbf{0}_{990}^T)^T$ , in which  $\mathbf{1}_d$  is a length  $d$  vector with all entries 1, and  $\mathbf{0}_d$  is a length  $d$  vector with all entries 0. Two different correlations  $\rho = 0$  and  $\rho = 0.5$  are investigated.

This is the classical LDA setting. In view of the linear optimal decision boundary, the nonparametric transformations in FANS is not necessary. Table 1 indicates some efficiency (not much) loss due to the more complex model FANS. However, by including the original features, FANS2 is comparable to the methods (e.g., PLR and ROAD) which learn boundaries linear in original features. In other words, the price to pay for using the more complex method FANS (FANS2) is small in terms of the classification error.

An interesting observation is that penGAM, which is based on a more general model class than FANS and FANS2, performs worse than our new methods. This is also expected as the complex parameter space considered by penGAM is unnecessary in view of a linear optimal decision boundary. Surprisingly, SVM performs poorly (even worse than NB), especially when all features are independent.

*Example 2.* The same settings as Example 1 except the common covariance matrix is an equal correlation matrix, with a common correlation  $\rho = 0.5$  and  $\rho = 0.9$ .

Same as in Example 1, FANS and FANS2 have performance comparable to PLR and ROAD. Although FAIR works very well in Example 1, where the features are independent (or nearly

independent), it fails badly when there is significant global pairwise correlation. Similar observations also hold for NB. This example shows that ignoring correlation among features could lead to significant loss of information and deterioration in the classification error.

*Example 3.* One class follows a multivariate Gaussian distribution, and the other a mixture of two multivariate Gaussian distributions. Precisely,

$$\text{Class 0: } N((3 \times \mathbf{1}_{10}^T, \mathbf{0}_{p-10}^T)^T, \boldsymbol{\Sigma}_p),$$

Class 1:  $0.5 \times N(\mathbf{0}_p, \mathbf{I}_p) + 0.5 \times N((6 \times \mathbf{1}_{10}^T, \mathbf{0}_{p-10}^T)^T, \boldsymbol{\Sigma}_p)$ , where  $\Sigma_{ii} = 1$ ,  $\Sigma_{ij} = \rho$  for  $i \neq j$ . Correlations  $\rho = 0$  and  $\rho = 0.5$  are considered.

In this example, Class 0 and Class 1 have the same mean, but have different marginal densities for the first ten dimensions. Table 1 shows that all methods based on linear boundary perform like random guessing, because the optimal decision boundary is highly nonlinear. penGAM is comparable to FANS and FANS2, but SVM cannot capture the oracle decision boundary well even if a nonlinear kernel is applied.

*Example 4.* Two classes follow uniform distributions,

Class 0: Unif ( $A$ ), and

Class 1: Unif ( $B \setminus A$ ), where  $A = \{\mathbf{x} \in \mathbb{R}^p : \|\mathbf{x}\|_2 \leq 1\}$  and  $B = [-1, 1]^p$ .

Clearly, the oracle decision boundary is  $\{\mathbf{x} \in \mathbb{R}^p : \|\mathbf{x}\|_2 = 1\}$ . Again, FANS and FANS2 capture this simple boundary well while the linear-boundary based methods fail to do so.

Computation times (in seconds) for various classification algorithms are reported in Table 2. FANS is extremely fast thanks to parallel computing. While penGAM performs similarly to FANS in the simulation examples, its computation cost is much higher. The similarity in performance is due to the abundance in training examples. We will demonstrate with an email spam classification example that penGAM fails to deliver satisfactory results on small samples.

#### 3.2 Real Data Analysis

We study two real examples and compare FANS (FANS2) with competing methods.

*3.2.1 E-mail Spam Classification.* First, we investigate a benchmark e-mail spam dataset. This dataset has been studied by Hastie, Tibshirani, and Friedman (2009) among others to demonstrate the power of additive logistic regression models. There are a total of  $n = 4601$  observations with  $p = 57$  numeric attributes. The attributes are, for instance, the percentage of specific words or characters in an e-mail, the average and maximum run lengths of upper case letters, and the total number of such letters. To show suitable application domains of FANS and FANS2, we vary the training proportion, from 5%, 10%, 20%, ..., to 80% of the data while assigning the rest as test set. Splits are repeated for 100 times, and we report the median classification errors.

Figure 2 and Table 3 summarize the results. First, we notice that FANS and FANS2 are very competitive when training sample sizes are small. As the training sample size increases, SVM becomes comparable to FANS2 and slightly better than FANS.




Figure 2. The median test classification error for the spam dataset using various proportions of the data as training sets for different classification methods.

In general, these three methods dominate throughout different training proportions. The more complex model penGAM failed to yield classifiers when training data proportion is less than 30% due to the difficulty of matrix inversion with the splines basis functions. For larger training samples, penGAM performs better than linear decision rules; however, it is not as competitive as either FANS or FANS2. Also interestingly, when the training sample size is 5%, NB performs as well as the sophisticated method FANS2 in terms of median classification error, but NB has a larger standard error. Moreover, the median classification error of NB remains almost unchanged when the sample size increases. In other words, NB’s independence assumption allows good training given very few data points, but it cannot benefit from larger samples due to severe model bias.

**3.2.2 Lung Cancer Classification.** We now evaluate the newly proposed classifiers on a popular gene expression dataset “Lung Cancer” (Gordon et al. 2002), which comes with predetermined, separate training and test sets. It contains  $p = 12533$  genes for  $n_0 = 16$  adenocarcinoma (ADCA) and  $n_1 = 16$  mesothelioma training vectors, along with 134 ADCA and 15 mesothelioma test vectors.

Following Dudoit, Fridlyand, and Speed (2002), Fan and Fan (2008), and Fan, Feng, and Tong (2012), we standardized each

sample to zero mean and unit variance. The classification results for FANS, FANS2, ROAD, penGAM, NB, FAIR, and SVM are summarized in Table 4. FANS and FANS2 achieve 0 test classification error, while the other methods fail to do so.

#### 4. THEORETICAL RESULTS

In this section, an oracle inequality regarding the excess risk is derived for FANS. Denote by  $\mathbf{f} = (f_1, \dots, f_p)^T$  and  $\mathbf{g} = (g_1, \dots, g_p)^T$  vectors of marginal densities of each class with  $\mathbf{f}_0 = (f_{0,1}, \dots, f_{0,p})^T$  and  $\mathbf{g}_0 = (g_{0,1}, \dots, g_{0,p})^T$  being the true densities. Let  $\{(X_i, Y_i)\}_{i=1}^n$  be iid copies of  $(X, Y)$ , and the regression function be modeled by

$$P(Y_1 = 1|X_1) = \frac{1}{1 + \exp(-m(\mathbf{Z}_1))},$$

where  $\mathbf{Z}_1 = (Z_{11}, \dots, Z_{1p})^T$ , each  $Z_{1j} = Z_{1j}(X_1) = \log f_j(X_{1j}) - \log g_j(X_{1j})$ , and  $m(\cdot)$  is a generic function in some function class  $\mathcal{M}$  that includes the linear functions. Now, let  $\mathcal{Q} = \{q = (m, \mathbf{f}, \mathbf{g})\}$  be the parameter space of interest with constraints on  $m, \mathbf{f}$  and  $\mathbf{g}$  be specified later. The loss function we consider is

$$\begin{aligned} \rho(q) &= \rho(m, \mathbf{f}, \mathbf{g}) = \rho_q(X_1, Y_1) \\ &= -Y_1 m(\mathbf{Z}_1) + \log(1 + \exp[m(\mathbf{Z}_1)]). \end{aligned}$$

Let  $m_0 = \arg \min_{m \in \mathcal{M}} P\rho(m, \mathbf{f}_0, \mathbf{g}_0)$ . Then, the target parameter is  $q^* = (m_0, \mathbf{f}_0, \mathbf{g}_0)$ . We use a working model with  $m_\beta(\mathbf{Z}_1) = \beta^T \mathbf{Z}_1$  to approximate  $m_0$ . Under this working model, for a given parameter  $q = (m_\beta, \mathbf{f}, \mathbf{g})$ , let

$$\pi_q(X_1) = P(Y_1 = 1|X_1) = \frac{1}{1 + \exp(-\beta^T \mathbf{Z}_1)}. \tag{4.4}$$

With this linear approximation, the loss function is the logistic loss

$$\rho(q) = \rho_q(X_1, Y_1) = -Y_1 \beta^T \mathbf{Z}_1 + \log(1 + \exp[\beta^T \mathbf{Z}_1]).$$

Denote the empirical loss by  $P_n \rho(q) = \sum_{i=1}^n \rho_q(X_i, Y_i)/n$ , and the expected loss by  $P\rho(q) = E\rho_q(X, Y)$ . In the following, we take  $\mathcal{M}$  as linear combinations of the transformed features so that  $m_0 = m_{\beta_0}$ , where

$$\beta_0 = \arg \min_{\beta \in \mathbb{R}^p} P\rho(m_\beta, \mathbf{f}_0, \mathbf{g}_0).$$

Table 3. Median classification error (in percentage) on e-mail spam data when the size of the training data varies

%	FANS	FANS2	ROAD	PLR	penGAM	LDA	NB	FAIR	SVM
5	11.1 (2.6)	10.5 (1.1)	13.6 (0.9)	13.5 (1.7)	—	13.6 (1.1)	10.5 (5.0)	15.6 (1.7)	11.2 (0.8)
10	8.7 (2.4)	8.5 (0.9)	11.3 (0.8)	10.5 (1.1)	—	11.3 (0.9)	10.7 (4.2)	13.5 (0.9)	9.4 (0.7)
20	8.0 (2.1)	7.7 (0.7)	10.6 (0.6)	9.0 (0.8)	—	10.3 (0.6)	10.7 (5.3)	12.4 (0.7)	8.1 (0.7)
30	7.8 (1.7)	7.4 (0.5)	10.3 (0.4)	8.9 (0.6)	9.2 (0.6)	10.1 (0.5)	10.7 (4.0)	11.7 (0.4)	7.4 (0.6)
40	7.2 (2.2)	6.9 (0.5)	10.1 (0.5)	9.0 (0.6)	8.6 (0.5)	10.0 (0.4)	10.5 (5.1)	11.5 (0.6)	7.0 (0.5)
50	7.4 (2.2)	7.0 (0.5)	9.9 (0.5)	8.5 (0.6)	8.3 (0.5)	9.9 (0.4)	10.7 (4.1)	11.8 (0.6)	6.9 (0.5)
60	7.4 (2.2)	6.8 (0.5)	9.8 (0.6)	9.3 (0.6)	7.8 (0.6)	9.5 (0.5)	10.6 (4.8)	11.8 (0.7)	6.5 (0.6)
70	7.2 (1.6)	6.4 (0.6)	9.5 (0.7)	9.2 (0.7)	7.4 (0.7)	9.4 (0.6)	10.5 (4.6)	11.4 (0.7)	6.4 (0.7)
80	6.9 (1.6)	6.3 (0.7)	9.4 (0.6)	9.3 (0.9)	7.4 (0.8)	9.2 (0.6)	10.4 (4.7)	11.4 (0.8)	6.3 (0.9)

Note: Standard errors are in the parentheses.

Table 4. Classification error and number of selected genes on lung cancer data

	FANS	FANS2	ROAD	PLR	penGAM	NB	FAIR	SVM
Training Error	0	0	1	0	0	6	0	0
Testing Error	0	0	1	6	2	36	7	4
No. of selected genes	52	52	52	15	16	12533	54	12533

In other words,  $q_0 = (m_{\beta_0}, f_0, g_0) = q^*$ . Hence, the excess risk for a parameter  $q$  is

$$\mathcal{E}(q) = P[\rho(q) - \rho(q^*)] = P[\rho(q) - \rho(q_0)]. \quad (4.5)$$

As described in Section 2, densities  $f_0$  and  $g_0$  are unavailable and must be estimated. Theorem 2 will establish the excess risk bound for the  $L = 1$  base procedure, which implies that the logistic regression coefficient and density estimates are close to the corresponding true values. Therefore, we expect for each  $l = 1, \dots, L$ , the predicted probability  $p_l$  is close to the oracle  $\pi_{q^*}(\cdot)$ . This further implies that  $p = 1/L \sum_{l=1}^L p_l$  is close to  $\pi_{q^*}(\cdot)$ . Given the above analysis, we fix  $L = 1$  in the FANS algorithm (i.e., only one random splitting is conducted) throughout the theoretical development.

Suppose we have labeled samples  $\{X_1^+, \dots, X_{n_1}^+\}$  (used to learn  $f_0$ ) and  $\{X_1^-, \dots, X_{n_1}^-\}$  (used to learn  $g_0$ ; theory carries over for different sample sizes), in addition to an iid sample  $\{(X_1, Y_1), \dots, (X_n, Y_n)\}$  (used to conduct PLR). Moreover, suppose  $\{(X_1, Y_1), \dots, (X_n, Y_n)\}$  is independent of  $\{X_1^+, \dots, X_{n_1}^+\}$  and  $\{X_1^-, \dots, X_{n_1}^-\}$ . A simple way to comprehend the above theoretical set up is that the sample size of  $2n_1 + n$  has been split into three groups. The notations  $P$  and  $E$  are regarding the random couple  $(X, Y)$ . We use the notation  $P^n$  to denote the probability measure induced by the sample  $\{(X_1, Y_1), \dots, (X_n, Y_n)\}$ , and notations  $P^+$  and  $P^-$  for the probability measures induced by the samples  $\{X_1^+, \dots, X_{n_1}^+\}$  and  $\{X_1^-, \dots, X_{n_1}^-\}$ .

The density estimates  $\hat{f} = (\hat{f}_1, \dots, \hat{f}_p)^T$  and  $\hat{g} = (\hat{g}_1, \dots, \hat{g}_p)^T$  are based on samples  $\{X_1^+, \dots, X_{n_1}^+\}$  and  $\{X_1^-, \dots, X_{n_1}^-\}$ :

$$\hat{f}_j(x) = \frac{1}{n_1 h} \sum_{i=1}^{n_1} K\left(\frac{X_{ij}^+ - x}{h}\right)$$

and

$$\hat{g}_j(x) = \frac{1}{n_1 h} \sum_{i=1}^{n_1} K\left(\frac{X_{ij}^- - x}{h}\right) \text{ for } j = 1, \dots, p,$$

in which  $K(\cdot)$  is a kernel function and  $h$  is the bandwidth. Then with these estimated marginal densities, we have an ‘‘oracle estimate’’  $q_1 = (\beta_1, \hat{f}, \hat{g})$ , where

$$\beta_1 = \arg \min_{\beta \in \mathbb{R}^p} P \rho(m_\beta, \hat{f}, \hat{g}).$$

It is the oracle given marginal density estimates  $\hat{f}$  and  $\hat{g}$  and is estimated in FANS by

$$\hat{\beta}_1 = \arg \min_{\beta \in \mathbb{R}^p} P_n \rho(m_\beta, \hat{f}, \hat{g}) + \lambda \|\beta\|_1.$$

Let  $\hat{q}_1 = (m_{\hat{\beta}_1}, \hat{f}, \hat{g})$ . Our goal is to control the excess risk  $\mathcal{E}(\hat{q}_1)$ , where  $\mathcal{E}$  is defined by (4.5). In the following, we introduce technical conditions for this task.

Let  $Z^0$  be the  $n \times p$  design matrix consisting of transformed covariates based on the true densities  $f_0$  and  $g_0$ . That is  $Z_{ij}^0 = \log f_{0,j}(X_{ij}) - \log g_{0,j}(X_{ij})$ , for  $i = 1, \dots, n$  and  $j = 1, \dots, p$ . In addition, let  $Z^0 = (Z_1^0, Z_2^0, \dots, Z_n^0)^T$ . Also, denote by  $|S|$  the cardinality of the set  $S$ , and by  $\|D\|_{\max} = \max_{ij} |D_{ij}|$  for any matrix  $D$  with elements  $D_{ij}$ .

*Assumption 1 (compatibility condition).* The matrix  $Z^0$  satisfies compatibility condition with a compatibility constant  $\phi(\cdot)$ , if for every subset  $S \subset \{1, \dots, p\}$ , there exists a constant  $\phi(S)$ , such that for all  $\beta \in \mathbb{R}^p$  that satisfy  $\|\beta_{S^c}\|_1 \leq 3\|\beta_S\|_1$ , it holds that

$$\|\beta_S\|_1^2 \leq \frac{1}{n\phi^2(S)} \|Z^0 \beta\|^2 |S|.$$

A direct application of Corollary 6.8 in Bühlmann and van de Geer (2011) leads to a compatibility condition on the estimated transform matrix  $\hat{Z}$ , in which  $\hat{Z}_{ij} = \log \hat{f}_j(X_{ij}) - \log \hat{g}_j(X_{ij})$ .

*Lemma 1.* Denote by  $E = \hat{Z} - Z^0$  the estimation error matrix of  $Z^0$ . If the compatibility condition is satisfied for  $Z^0$  with a compatibility constant  $\phi(\cdot)$ , and the following inequalities hold

$$\frac{32\|E\|_{\max}|S|}{\phi(S)^2} \leq 1, \text{ for every } S \subset \{1, \dots, p\}, \quad (4.6)$$

the compatibility condition holds for  $\hat{Z}$  with a new compatibility constant  $\phi_1(\cdot) \geq \phi(\cdot)/\sqrt{2}$ .

The compatibility condition can be interpreted as a condition that bounds the restricted eigenvalues. The irrerepresentable condition (Zhao and Yu 2006) and the sparse Riesz condition (SRC) (Zhang and Huang 2008) are in similar spirits. Essentially, these conditions avoid high correlation among subsets where signals are concentrated; such high correlation may cause difficulty in parameter estimation and risk prediction.

To help theoretical derivation, we introduce two intermediate  $L_0$ -penalized estimates. Given the true densities  $f_0$  and  $g_0$ , consider a penalized theoretical solution  $q_0^* = (\beta_0^*, f_0, g_0)$ , where

$$\beta_0^* = \arg \min_{\beta \in \mathbb{R}^p} 3P\rho(m_\beta, f_0, g_0) + 2H\left(\frac{4\lambda\sqrt{s_\beta}}{\phi(S_\beta)}\right), \quad (4.7)$$

in which  $H(\cdot)$  is a strictly convex function on  $[0, \infty)$  with  $H(0) = 0$ ,  $s_\beta = |S_\beta|$  is the cardinality of  $S_\beta = \{j : \beta_j \neq 0\}$ , and  $\phi(\cdot)$  is the compatibility constant for  $Z^0$ . Throughout the article, we consider a specific quadratic function<sup>2</sup>  $H(v) = v^2/(4c)$  whose convex conjugate is  $G(u) = \sup_v \{uv - H(v)\} = cu^2$ .

<sup>2</sup>The following theoretical results can be derived for a generic strictly convex function  $H(\cdot)$  along the same lines.

Then, Equation (4.7) defines an  $L_0$ -penalized oracle:

$$\beta_0^* = \arg \min_{\beta \in \mathbb{R}^p} 3P\rho(m_\beta, f_0, g_0) + \frac{8\lambda^2 s_\beta}{c\phi^2(S_\beta)}. \quad (4.8)$$

Similarly, with density estimate vectors  $\hat{f}$  and  $\hat{g}$ , we define an  $L_0$ -penalized oracle estimate  $q_1^* = (m_{\beta_1^*}, \hat{f}, \hat{g})$ , where

$$\beta_1^* = \arg \min_{\beta \in \mathbb{R}^p} 3P\rho(m_\beta, \hat{f}, \hat{g}) + \frac{8\lambda^2 s_\beta}{c\phi_1^2(S_\beta)}. \quad (4.9)$$

To study the excess risk  $\mathcal{E}(\hat{q}_1)$ , we consider its relationship with  $\mathcal{E}(q_1^*)$  and  $\mathcal{E}(q_0^*)$ .

*Assumption 2 (Uniform Margin Condition).* There exists  $\eta > 0$  such that for all  $(m_\beta, f, g)$  satisfying  $\|\beta - \beta_0\|_\infty + \max_{1 \leq j \leq p} \|f_j - f_{0,j}\|_\infty + \max_{1 \leq j \leq p} \|g_j - g_{0,j}\|_\infty \leq 2\eta$ , we have

$$\mathcal{E}(m_\beta, f, g) \geq c\|\beta - \beta_0\|_2^2, \quad (4.10)$$

where  $c$  is the positive constant in (4.8).

The uniform margin condition is related to the one defined in Tsybakov (2004) and van de Geer (2008). It is a type of “identifiability” condition. Basically, near the target parameter  $q_0 = (m_{\beta_0}, f_0, g_0)$ , the functional value needs to be sufficiently different from the value on  $q_0$  to enable enough separability of parameters. Note that we impose the uniform margin condition in both the neighborhood of the parametric component  $\beta_0$  and the nonparametric components  $f_0$  and  $g_0$ , because we need to estimate the densities, in addition to the parametric part. A related concept in binary classification is called “margin assumption,” which was first introduced by Polonik (1995) for densities.

To study the relationship between  $\mathcal{E}(\hat{q}_1)$  and  $\mathcal{E}(q_1^*)$ , we define

$$\begin{aligned} v_n(\beta) &= (P_n - P)\rho(m_\beta, \hat{f}, \hat{g}) \text{ and } W_M \\ &= \sup_{\|\beta - \beta_1^*\| \leq M} |v_n(\beta) - v_n(\beta_1^*)|. \end{aligned}$$

Denote by

$$2\epsilon^* = 3\mathcal{E}(m_{\beta_1^*}, \hat{f}, \hat{g}) + \frac{8\lambda^2 s_{\beta_1^*}}{c\phi_1^2(S_{\beta_1^*})}.$$

Set  $M^* = \epsilon^*/\lambda_0$  ( $\lambda_0$  to specified in Theorem 1) and

$$\mathcal{J}_1 = \{W_{M^*} \leq \lambda_0 M^*\} = \{W_{M^*} \leq \epsilon^*\}.$$

The idea here is to choose  $\lambda_0$  such that the event  $\mathcal{J}_1$  has high probability.

A few more notations are introduced to facilitate the discussion. Let  $\tau > 0$ . Denote by  $\lfloor \tau \rfloor$  the largest integer strictly less than  $\tau$ . For any  $x, x' \in \mathbb{R}$  and any  $\lfloor \tau \rfloor$  times continuously differentiable real valued function  $u$  on  $\mathbb{R}$ , we denote by  $u_x$  its Taylor polynomial of degree  $\lfloor \tau \rfloor$  at point  $x$ :

$$u_x(x') = \sum_{|s| \leq \lfloor \tau \rfloor} \frac{(x' - x)^s}{s!} D^s u(x).$$

For  $L > 0$ , the  $(\tau, L, [-1, 1])$ -Hölder class of functions, denoted by  $\Sigma(\tau, L, [-1, 1])$ , is the set of functions  $u : \mathbb{R} \rightarrow \mathbb{R}$  that are  $\lfloor \tau \rfloor$  times continuously differentiable and satisfy, for any  $x, x' \in [-1, 1]$ , the inequality:

$$|u(x') - u_x(x')| \leq L|x - x'|^\tau.$$

The  $(\tau, L, [-1, 1])$ -Hölder class of density is defined as

$$\begin{aligned} \mathcal{P}_\Sigma(\tau, L, [-1, 1]) \\ = \left\{ p : p \geq 0, \int p = 1, p \in \Sigma(\tau, L, [-1, 1]) \right\}. \end{aligned}$$

*Assumption 3.* Assume that  $\beta_1$  is in the interior of some compact set  $C_p$ . There exists an  $\epsilon_0 \in (0, 1)$  such that for all  $\beta \in C_p$  and  $f_j, g_j \in \mathcal{P}_\Sigma(2, L, [-1, 1])$ ,  $j = 1, \dots, p$ ,  $\epsilon_0 < \pi_{(m_\beta, f, g)}(\cdot) < 1 - \epsilon_0$ .

*Assumption 4.*  $\|Z^0\|_{\max} \leq K$  for some absolute constant  $K > 0$ , and  $\|\beta_0\|_\infty \leq C_1$  for some absolute constant  $C_1 > 0$ .

*Assumption 5.* The penalty level  $\lambda$  is in the range of  $(8\lambda_0, L\lambda_0)$  for some  $L > 8$ . Moreover, the following holds

$$\frac{8KL^2(e^\eta/\epsilon_0 + 1)^2}{\eta} \frac{\lambda_0 s_{\beta_1^*}}{\phi_1^2(S_{\beta_1^*})} \leq 1,$$

where  $\eta$  is as in the uniform margin condition.

Assumption 3 is a regularity condition on the probability of the event that the observation belongs to class 1. Since the FANS estimator is based on the estimated densities, we impose the constraints in a neighborhood of the oracle estimate  $\beta_1$  (when using  $\hat{f}$  and  $\hat{g}$ ). Assumption 4 bounds the maximum absolute entry of the design matrix as well as the maximum absolute true regression coefficient. Assumption 5 posits a proper range of the penalty parameter  $\lambda$  to guarantee that the penalized estimator mimics the un-penalized oracle.

*Assumption 6.* Suppose the feature measurement  $X$  has a compact support  $[-1, 1]^p$ , and  $f_{0,j}, g_{0,j} \in \mathcal{P}_\Sigma(2, L, [-1, 1])$  for all  $j = 1, \dots, p$ , where  $\mathcal{P}_\Sigma$  denotes a Hölder class of densities.

*Assumption 7.* Suppose there exists  $\epsilon_l > 0$  such that for all  $j = 1, \dots, p$ ,  $\epsilon_l \leq f_{0,j}, g_{0,j} \leq \epsilon_l^{-1}$ . Also we truncate estimates  $\hat{f}_j$  and  $\hat{g}_j$  at  $\epsilon_l$  and  $\epsilon_l^{-1}$ .

*Assumption 8.*

$$n_1^{\frac{7}{10} - \frac{3}{4}\alpha} (\log(3p))^{\frac{3}{4}} (\log n_1)^{\frac{1}{10}} = o(1),$$

and

$$n_1^{1 - \alpha} (\log(3p))^{\frac{1}{2}} (\log n_1)^{\frac{2}{5}} = o(1),$$

for some constant  $\alpha > 7/15$ .

Assumption 6 imposes constraints on the support of  $X$  and smoothness condition on the true densities  $f_0$  and  $g_0$ , which help control the estimation error incurred by the nonparametric density estimates. Assumption 7 assumes that the marginal densities and the kernel are strictly positive on  $[-1, 1]^p$ . Assumption 8 puts a restriction on the growth of the dimensionality  $p$  in terms of sample size  $n_1$ .

We now provide a lemma to bound the uniform deviation between  $\hat{f}_j$  and  $f_{0,j}$  for  $j = 1, \dots, p$ .

*Lemma 2.* Under Assumptions 6–8, taking the bandwidth  $h = (\frac{\log n_1}{n_1})^{1/5}$ , for any  $\delta_1 > 0$ , there exists  $N_1^*$  such that if  $n_1 \geq$



$N_1^*$ ,

$$P^{+-} \left( \max_{1 \leq j \leq p} \|\hat{f}_j - f_{0,j}\|_\infty \geq m \right) \leq \delta_1, \quad \text{and}$$

$$P^{+-} \left( \max_{1 \leq j \leq p} \|\hat{g}_j - g_{0,j}\|_\infty \geq m \right) \leq \delta_1,$$

for  $m = C_2 \sqrt{\frac{2 \log(3p/\delta_1)}{n_1^{1-\alpha}}}$ , and  $C_2$  is an absolute constant.

Denote by

$$\mathcal{J}_2 = \left\{ \max_{1 \leq j \leq p} \|\hat{f}_j - f_{0,j}\|_\infty \leq \eta/2, \max_{1 \leq j \leq p} \|\hat{g}_j - g_{0,j}\|_\infty \leq \eta/2 \right\},$$

where  $\eta$  is the constant in the uniform margin condition. It is straightforward from Lemma 2 that

$$P^{+-}(\mathcal{J}_2) \geq 1 - \frac{6p}{\exp(\eta^2 n_1^{1-\alpha}/4C_2^2)}.$$

The next lemma can be similarly derived as Lemma 2, so its proof is omitted.

*Lemma 3.* Under Assumptions 6–8, taking the bandwidth  $h = (\frac{\log n_1}{n_1})^{1/5}$ , for any  $\delta > 0$ , there exists  $N_2^*$  such that if  $n_1 \geq N_2^*$ ,

$$P^{+-}(\|\mathbf{E}\|_{\max} \geq m) \leq \delta,$$

where  $\mathbf{E}$  is the estimation error matrix as defined in Lemma 1 and  $m = C_3 \sqrt{\frac{2 \log(3p/\delta)}{n_1^{1-\alpha}}}$  for some absolute constant  $C_3$ .

*Corollary 1.* Under Assumptions 6–8, take the bandwidth  $h = (\frac{\log n_1}{n_1})^{1/5}$ . On the event  $\mathcal{J}_3 = \{\|\mathbf{E}\|_{\max} \leq C_3 \sqrt{\frac{2 \log(3p/\delta)}{n_1^{1-\alpha}}}\}$  (regarding labeled samples) with  $P^{+-}(\mathcal{J}_3) > 1 - \delta$ , there exists  $N_2^* \in \mathbb{N}$  and  $C_4 > 0$  such that if  $n_1 \geq N_2^*$ ,  $|F_{kl}| = |\hat{Z}_{1k} - Z_{1k}^0| \cdot |\hat{Z}_{1l} - Z_{1l}^0| \leq C_4 b_{n_1}$  uniformly for  $k, l = 1, \dots, p$ , where  $b_{n_1} = 2 \log(3p/\delta)/n_1^{1-\alpha}$ .

Denote by

$$\mathcal{J}_4 = \left\{ 32 \|\mathbf{E}\|_{\max} \max_{S \subset \{1, \dots, p\}} \frac{|S|}{\phi(S)^2} \leq 1 \right\}.$$

On the event  $\mathcal{J}_4$ , the inequality (4.6) holds, and the compatibility condition is satisfied for  $\hat{\mathbf{Z}}$  if we assume Assumption 1 (by Lemma 1). Moreover, it can be derived from Lemma 3 by taking a specific  $\delta$ ,

$$P^{+-}(\mathcal{J}_4) \geq 1 - 3p \exp\{-n_1^{1-\alpha}/(2048C_3^2 A_p^2)\},$$

where  $A_p = \max_{S \subset \{1, \dots, p\}} |S|/\phi(S)^2$ . Combining Lemma 2 and the uniform margin condition, we see that for given estimators  $\hat{\mathbf{f}}$  and  $\hat{\mathbf{g}}$ , the margin condition holds for the estimated transformed matrix  $\hat{\mathbf{Z}}$  involved in the FANS estimator  $\hat{\beta}_1$ . Following similar lines as in van de Geer (2008) delivers the following theorem, so a formal proof is omitted.

*Theorem 1* (Oracle Inequality). In addition to Assumptions 1–8, assume  $\|m_{\beta_1^*} - m_{\beta_0}\|_\infty \leq \eta/2$  and  $\mathcal{E}(m_{\beta_1^*}, \hat{\mathbf{f}}, \hat{\mathbf{g}})/\lambda_0 \leq \eta/4$ . Then on the event  $\mathcal{J}_1 \cap \mathcal{J}_2 \cap \mathcal{J}_3 \cap \mathcal{J}_4$ , we have

$$\mathcal{E}(m_{\hat{\beta}_1}, \hat{\mathbf{f}}, \hat{\mathbf{g}}) + \lambda \|\hat{\beta}_1 - \beta_1^*\|_1 \leq 6\mathcal{E}(m_{\beta_1^*}, \hat{\mathbf{f}}, \hat{\mathbf{g}}) + \frac{16\lambda^2 s_{\beta_1^*}^2 (e^\eta/\epsilon_0 + 1)^2}{c\phi_1^2(S_{\beta_1^*})}.$$

Moreover, when  $n_1 \geq \max(N_1^*, N_2^*)$  and under the normalization condition that  $\|Z_{1j}\|_\infty \leq 1$  for all  $j = 1, \dots, p$ , it holds that

$$\begin{aligned} & \mathbb{P}(\mathcal{J}_1 \cap \mathcal{J}_2 \cap \mathcal{J}_3 \cap \mathcal{J}_4) \\ & \geq 1 - \exp(-t) - 6p \exp\{-\eta^2 n_1^{1-\alpha}/(4C_2^2)\} - \delta \\ & \quad - 3p \exp\{-n_1^{1-\alpha}/(2048C_3^2 A_p^2)\}, \end{aligned}$$

for

$$\lambda_0 := 4\lambda^* + \frac{tK}{3n} + \sqrt{\frac{2t}{n}(1 + 8\lambda^*)},$$

where  $\mathbb{P}$  is the probability with regards to all the samples and

$$\lambda^* = \sqrt{\frac{2 \log(2p)}{n}} + \frac{K \log(2p)}{3n}.$$

Theorem 1 shows that with high probability, the excess risk of the FANS estimator can be controlled in terms of the excess risk of  $q_1^*$  when using the estimated density functions  $\hat{\mathbf{f}}$  and  $\hat{\mathbf{g}}$  plus a term of explicit order. Next, we will study the excess risk of  $q_1^*$ .

*Assumption 9.* Let  $\mathbf{Z}_1^0(\beta_1)$  be the subvector of  $\mathbf{Z}_1^0$  corresponding to the nonzero components of  $\beta_1$ , and  $b_{n_1} = \log(3p/\delta_1)/n_1^{1-\alpha}$ . Assume  $s_{\beta_1} \leq a_{n_1}$  for some deterministic sequence  $\{a_{n_1}\}$ , and  $a_{n_1} \cdot b_{n_1} = o(1)$ . In addition,  $0 < C_5 \leq \lambda_{\min}(P\{\mathbf{Z}_1^0(\beta_1)\mathbf{Z}_1^0(\beta_1)^T\})$ , for some absolute constant  $C_5$ .

Assumption 9 allows the number of nonzero elements of  $\beta_1$  to diverge at a slow rate with  $n_1$ . Also, it demands a lower bound of the restricted eigenvalue of the submatrix of  $\mathbf{Z}^0$  corresponding to the nonzero components of  $\beta_1$ .

*Lemma 4.* Let  $Q(\beta) = P\rho(m_\beta, \hat{\mathbf{f}}, \hat{\mathbf{g}}) + \lambda\|\beta\|_0$ , and  $\bar{\beta}_1 = \min\{|\beta_{1,j}| : j \in S_{\beta_1}\}$ . Under Assumptions 3, 6–9, on the event  $\mathcal{J}_3$ , there exists a constant  $N_3^*$  such that, if  $n_1 \geq N_3^*$  and the penalty parameter  $\lambda < 0.5C_5\epsilon_0(1 - \epsilon_0)\bar{\beta}_1^2$ , the  $L_0$  penalized solution coincides with the unpenalized version; that is  $\beta_1^* = \bar{\beta}_1$ .

*Theorem 2* (Oracle Inequality). In addition to Assumptions 1–9, suppose  $4C_1C_4s_{\beta_0}^2 b_{n_1} \leq \lambda_0\eta$ , the penalty parameter  $\lambda \in (8\lambda_0, \min(L\lambda_0, 0.5C_5\epsilon_0(1 - \epsilon_0) \cdot \min_{j:\beta_{1,j} \neq 0}(|\beta_{1,j}|)))$ , where  $C_5$  is defined in Assumption 9,  $\|m_{\beta_1^*} - m_{\beta_0}\|_\infty \leq \eta/2$  and  $n_1 \geq \max(N_1^*, N_2^*, N_3^*)$ . Taking the bandwidth  $h = (\frac{\log n_1}{n_1})^{1/5}$ , on the event  $\mathcal{J}_1 \cap \mathcal{J}_2 \cap \mathcal{J}_3 \cap \mathcal{J}_4$  as in Theorem 1, we have

$$\mathcal{E}(m_{\beta_1^*}, \hat{\mathbf{f}}, \hat{\mathbf{g}}) \leq C_1 C_4 s_{\beta_0}^2 b_{n_1}.$$

Then in view of Theorem 1, we have

$$\mathcal{E}(m_{\hat{\beta}_1}, \hat{\mathbf{f}}, \hat{\mathbf{g}}) \leq \frac{16\lambda^2 s_{\beta_1^*}^2 (e^\eta/\epsilon_0 + 1)^2}{c\phi_1^2(S_{\beta_1^*})} + 6C_1 C_4 s_{\beta_0}^2 b_{n_1}.$$

This theorem finale requires quite some conditions. We now deconvolute them by providing a high level description of the motivations behind these conditions. Because FANS is essentially a two-step procedure, we need both steps to do well to have the theoretical performance guarantee. The first step is to estimate the transformed features. In this step, we need regularity conditions on the class conditional densities  $f_0$  and  $g_0$ , and regularity conditions on the kernel density estimate components, such as the kernel  $K$ . Also, the sample size needs to

Table 5. Median test error (in percentages) for Example 5

Ex	FANS	FANS2	ROAD	PLR	penGAM	NB	FAIR	SVM
5	6.7 (1.1)	6.9 (1.1)	50.2 (2.1)	50.0 (1.4)	8.0 (2.3)	50.2 (2.3)	49.7 (2.1)	50.0 (2.0)

Note: Standard errors are in the parentheses.

be big enough so that the kernel density estimate is close to the truth. The second step is PLR using the estimated transformed features. In this step, usual conditions on the penalty level, design matrix, and signal strength are needed. Moreover, some conditions that link nonparametric and parametric components, that is, the first and second steps, such as the uniform margin condition should be in place.

From Theorem 2, it is clear that the excess risk of the FANS estimator is naturally decomposed into two parts. One part is due to the nonparametric density estimation while the other part is due to the regularized logistic regression on the estimated transformed covariates. When both the penalty parameter  $\lambda$  and the bandwidth  $h$  of the nonparametric density estimates  $\hat{f}$  and  $\hat{g}$  are chosen appropriately, the FANS estimator will have a diminishing excess risk with high probability. Note that one can make explicit  $\lambda$  to obtain a bound on the excess risk in terms of the sample sizes  $n$  and  $n_1$ , and the dimensionality  $p$ . Also, it is worth noting that the development of oracle inequality of the FANS procedure  $\hat{\beta}_1$  is accomplished via an important bridge of the  $L_0$ -regularized estimator  $\beta_1^*$ .

The oracle inequality for FANS2 can be developed along similar lines. In particular, the parameter under the working model will be changed to  $q_2 = (m_{(\beta, \gamma)}, \mathbf{f}, \mathbf{g})$  and the success probability given  $\mathbf{X}_1$  will be modeled by a modified logistic function

$$\pi_{q_2}(\mathbf{X}_1) = P(Y_1 = 1 | \mathbf{X}_1) = \frac{1}{1 + \exp(-\boldsymbol{\beta}^T \mathbf{Z}_1 - \boldsymbol{\gamma}^T \mathbf{X}_1)}, \quad (4.11)$$

where we note that in addition to the transformed features, the original features are also included. We would like to emphasize that  $\mathbf{X}_1$  is observed, and therefore, there is no need to control its estimation error as we did for  $\mathbf{Z}_1$ . The conditions for the theory of FANS can be adapted to establish an oracle inequality for FANS2. We omit the details to avoid duplication of similar conditions and arguments.

## 5. DISCUSSION

We propose a new two-step nonlinear rule FANS (and its variant FANS2) to tackle binary classification problems in high-dimensional settings. FANS first augments the original feature space by leveraging flexibility of nonparametric estimators and then achieves feature selection through regularization (penalization). It combines linearly the best univariate transforms that essentially augment the original features for classification. Since nonparametric techniques are only performed on each dimension, we enjoy a flexible decision boundary without suffering from the curse of dimensionality. An array of simulation and real data examples, supported by an efficient parallelized al-

gorithm, demonstrate the competitive performance of the new procedures.

To verify our methods' performance against model misspecification, we evaluate different classifiers on the following example that has nonadditive optimal decision boundary, with the results shown in Table 5. Similar to Example 4, FANS and FANS2 perform the best among all competing methods (penGAM performs slightly worse with a larger standard error).

*Example 5.* Nonadditive decision boundary. In particular, for  $\mathbf{x} \sim N(\mathbf{0}_p, I_p)$ , let  $y = \mathbb{I}\{x_1^2 \sqrt{x_2^2 + x_3^4 + 1} \geq 0.75\}$ .

One problem in applications we are faced with is whether we should use FANS or FANS2. While we do not have a universal rule, a rule of thumb might shed some insight. From the simulation examples, we see when the sample size is small and/or decision boundary is highly nonlinear, FANS is recommended over FANS2. Otherwise, FANS2 is recommended. Admittedly, in real data applications, it is often impossible to know a priori how the oracle decision boundary looks like. Data abundance can be a rough guideline in these scenarios.

A few extensions are worth further investigation. For example, an extension to multi-class classification is an interesting future work. Beyond a specific procedure, FANS establishes a general two-step classification framework. For the first step, one can use other types of marginal density estimators, for example, local polynomial density estimates. For the second step, one might rely on other classification algorithms, for example, the SVM,  $k$ -nearest neighbors, etc. Searching for the best two-step combination is an important but difficult task, and we believe that the answer mainly depends on the specific applications.

We can further augment the features by adding pairwise bivariate density ratios. These bivariate densities can be approximated by the bivariate kernel density estimates. Alternatively, we can restrict our attention to bivariate ratios of features selected by FANS. The latter has significantly fewer features.

Dimensions of datasets (e.g., SNPs) in many contemporary applications could be in millions. In such ultra-high dimensional scenarios, directly applying the FANS (FANS2) approach could cause problems due to high computational complexity and instability of the estimation. It will be beneficial to have a prior step to reduce the dimensionality in the original data. Notable works toward this effort on the theoretical front include Fan and Lv (2008), which introduced the SIS to screen out the marginally unimportant variables. Subsequently, Fan, Feng, and Song (2011) proposed nonparametric independence screening, an extension of SIS to the additive models.

APPENDIX

The appendix contains technical proofs and Lemma A.1.

*Proof of Lemma 2.* For any  $r, m > 0$ ,

$$\begin{aligned} & P^{+-} \left( \max_{1 \leq j \leq p} \|\hat{f}_j - f_{0,j}\|_\infty \geq m \right) \\ & \leq e^{-rm} E^{+-} \exp \left( \max_{1 \leq j \leq p} r \|\hat{f}_j - f_{0,j}\|_\infty \right) \\ & = e^{-rm} E^{+-} \left( \max_{1 \leq j \leq p} \exp r \|\hat{f}_j - f_{0,j}\|_\infty \right) \\ & \leq e^{-rm} \sum_{j=1}^p E^{+-} \left( \exp r \|\hat{f}_j - f_{0,j}\|_\infty \right). \end{aligned}$$

Since we assumed that all  $\hat{f}_j$  and  $f_{0,j}$  are uniformly bounded by  $\epsilon_l^{-1}$ ,  $\|\hat{f}_j - f_{0,j}\|_\infty$  is bounded by  $\epsilon_l^{-1}$  for all  $j \in \{1, \dots, p\}$ . This coupled with Lemma 1 in Tong (2013), provides a high probability bound for  $\|\hat{f}_j - f_{0,j}\|_\infty$ , gives rise to the following inequality,

$$E^{+-} \exp \left( r \|\hat{f}_j - f_{0,j}\|_\infty \right) \leq \exp \left( r \sqrt{\frac{\log(n_1/\delta_2)}{n_1 h}} \right) + \exp(r\epsilon_l^{-1}) \cdot \delta_2,$$

where  $\delta_2$  plays the role of  $\epsilon$  in Lemma 1 of Tong (2013)(taking constant  $C = 1$  for simplicity).

Finding the optimal order for  $r$  does not seem to be feasible. So we plug in  $r = n_1^{1-\alpha} m$  and  $\delta_2 = \exp(-r\epsilon_l^{-1})$ , then

$$\begin{aligned} & P^{+-} \left( \max_{1 \leq j \leq p} \|\hat{f}_j - f_{0,j}\|_\infty \geq m \right) \\ & \leq p \exp(-n_1^{1-\alpha} m^2) \left\{ 1 + \exp \left( n_1^{1-\alpha} m \sqrt{\frac{\log n_1}{n_1 h} + \frac{n_1^{1-\alpha} m \epsilon_l^{-1}}{n_1 h}} \right) \right\} \\ & \leq p \exp(-n_1^{1-\alpha} m^2) \left\{ 1 + \exp \left[ \sqrt{2} n_1^{1-\alpha} m \left( \sqrt{\frac{\log n_1}{n_1 h}} + \sqrt{\frac{m \epsilon_l^{-1}}{n_1 h}} \right) \right] \right\} \\ & \leq p \exp(-n_1^{1-\alpha} m^2) \left\{ 1 + \exp \left[ \sqrt{2} n_1^{1-\alpha} m \left( \frac{\log n_1}{n_1} \right)^{\frac{2}{5}} \right. \right. \\ & \quad \left. \left. + \sqrt{2} m^{\frac{3}{2}} \epsilon_l^{-\frac{1}{2}} n_1^{\frac{11}{10} - \frac{3}{2}\alpha} (\log n_1)^{\frac{1}{10}} \right] \right\}, \end{aligned}$$

where in the last inequality we have used the bandwidth  $h = (\frac{\log n_1}{n_1})^{1/5}$ .

The results are derived by taking  $m = \sqrt{\frac{2 \log(3p/\delta_1)}{n_1^{1-\alpha}}}$  (so  $\delta_1 = 3p \exp(-n_1^{1-\alpha} m^2)$ ), and by taking Assumption 8. Note that we need to introduce  $\alpha > 0$  because the consistency conditions do not hold for  $\alpha = 0$ . In fact, we need at least  $\alpha > 7/15$ . Under this assumption, there exists a positive integer  $N_1^*$  such that if  $n_1 \geq N_1^*$ ,

$$\begin{aligned} & 1 + \exp \left[ 2^{\frac{3}{5}} \epsilon_l^{-\frac{1}{2}} n_1^{\frac{7}{10} - \frac{3}{2}\alpha} (\log(3p/\delta_1))^{\frac{3}{5}} (\log n_1)^{\frac{1}{10}} \right. \\ & \quad \left. + 2 n_1^{\frac{1}{10} - \alpha} (\log(3p/\delta_1))^{\frac{1}{2}} (\log n_1)^{\frac{2}{5}} \right] \leq 3. \end{aligned}$$

Therefore, for  $n_1 \geq N_1^*$ ,

$$P^{+-} \left( \max_{1 \leq j \leq p} \|\hat{f}_j - f_{0,j}\|_\infty \geq m \right) \leq \delta_1, \text{ for } m = \sqrt{\frac{2 \log(3p/\delta_1)}{n_1^{1-\alpha}}}.$$

*Lemma A.1.* For any vector  $\theta_0 = (\theta_{0,1}, \dots, \theta_{0,p})^T$ , let  $S_{\theta_0} = \{j : \theta_{0,j} \neq 0\}$ , and let the minimum signal level be  $\theta_0 = \min\{|\theta_{0,j}| : j \in S_{\theta_0}\}$ . Let  $g(\theta_j) = c_j(\theta_j - \theta_{0,j})^2 + \lambda \|\theta_j\|_0$ , where  $c_j > 0$ . If  $\lambda \leq c_j \bar{\theta}_0^2$ ,  $g(\theta_j)$  achieves the unique minimum at  $\theta_j = \theta_{0,j}$ .

*Proof of Lemma A.1.*

For  $\theta_{0,j} = 0$ , the result is obvious. For  $\theta_{0,j} \neq 0$ , we have  $j \in S_{\theta_0}$  and

$$\begin{aligned} g(\theta_j) & \geq \lambda \|\theta_j\|_0 I(\theta_j \neq 0) + c_j(\theta_j - \theta_{0,j})^2 I(\theta_j = 0) \\ & = \lambda \|\theta_j\|_0 I(\theta_j \neq 0) + c_j \theta_{0,j}^2 I(\theta_j = 0). \end{aligned}$$

If  $\lambda \|\theta_{0,j}\|_0 \leq c_j \bar{\theta}_0^2$ ,

$$g(\theta_j) \geq \lambda \|\theta_j\|_0 I(\theta_j \neq 0) + \lambda \|\theta_{0,j}\|_0 I(\theta_j = 0) = \lambda \|\theta_{0,j}\|_0.$$

Since  $g(\theta_{0,j}) = \lambda \|\theta_{0,j}\|_0$ , the lemma follows.

*Proof of Lemma 4.*

Denote  $Q_0(\beta) = P\rho(m\beta, \hat{f}, \hat{g})$ . Then, we have  $\beta_1 = \arg \min_{\beta \in \mathbb{R}^p} Q_0(\beta)$ . Since  $\nabla Q_0(\beta_1) = 0$  and

$$\begin{aligned} \nabla^2 Q_0(\beta) & = P\{\hat{Z}_1 \hat{Z}_1^T \exp(\hat{Z}_1^T \beta) (1 + \exp(\hat{Z}_1^T \beta))^{-2}\} \\ & \geq \epsilon_0 (1 - \epsilon_0) P\{\hat{Z}_1 \hat{Z}_1^T\} \geq \mathbf{0}. \end{aligned}$$

By Taylor's expansion of  $Q_0(\beta)$  at  $\beta_1$ ,

$$Q(\beta) = Q_0(\beta_1) + 0.5(\beta - \beta_1)^T \nabla^2 Q_0(\tilde{\beta})(\beta - \beta_1) + \lambda \|\beta\|_0, \quad (\text{A.1})$$

where  $\tilde{\beta}$  lies between  $\beta$  and  $\beta_1$ . Let  $\hat{M} = P\{\hat{Z}_1(\beta_1)\hat{Z}_1(\beta_1)^T\}$ , where  $\hat{Z}_1(\beta_1)$  is the subvector of  $\hat{Z}_1$  corresponding to the nonzero components of  $\beta_1$ , and  $M = P\{Z_1^0(\beta_1)Z_1^0(\beta_1)^T\}$ , where  $Z_1^0(\beta_1)$  is the subvector of  $Z_1^0$  corresponding to the nonzero components of  $\beta_1$ . Let  $F = \hat{M} - M$  (a symmetric matrix). From the uniform deviance result of Lemma 3, with probability  $1 - \delta$  regarding the labeled samples, there exists a constant  $C_4 > 0$  such that  $|F_{kl}| \leq C_4 b_{n_1}$  uniformly for  $k, l = 1, \dots, s_{\beta_1}$ , where  $b_{n_1} = 2 \log(3p/\delta)/n_1^{1-\alpha}$ .

Hence,  $\|\mathbf{F}\|_2 \leq \|\mathbf{F}\|_F \leq C_4 s_{\beta_1} b_{n_1} \leq C_4 a_{n_1} b_{n_1}$ . For any eigenvalue  $\lambda(\hat{M})$ , by the Bauer-Fike inequality (Bhatia 1997), we have  $\min_{1 \leq k \leq s_{\beta_1}} |\lambda(\hat{M}) - \lambda_k(M)| \leq \|\mathbf{F}\|_2 \leq C_4 a_{n_1} b_{n_1}$ , where  $\lambda_k(A)$  denotes the  $k$ th largest eigenvalue of  $A$ . In addition, in view of Assumption 9, there exists  $k \in S_{\beta_1}$  such that

$$\lambda_{\min}(\hat{M}) \geq \lambda_k(M) - C_4 a_{n_1} b_{n_1} \geq \lambda_{\min}(M) - C_4 a_{n_1} b_{n_1} \geq C_5 - C_4 a_{n_1} b_{n_1}.$$

Since  $a_{n_1} b_{n_1} = o(1)$ , there exists  $N_3^*(\delta)$  such that when  $n_1 > N_3^*(\delta)$ , we have  $\lambda_{\min}(\hat{M}) > 0$ .

Let  $\beta_1^{(1)}$  be the subvector of  $\beta_1$  consisting of the nonzero components. Then by (A.1) and Lemma A.1 for each  $j \in S_{\beta_1}$  with  $\lambda < 0.5C_5\epsilon_0(1 - \epsilon_0)\bar{\beta}_1^2$ , we have

$$\begin{aligned} Q(\beta) & \geq Q_0(\beta_1) + 0.5(C_5 - C_4 a_{n_1} b_{n_1})\epsilon_0(1 - \epsilon_0)\|\beta^{(1)} - \beta_1^{(1)}\|^2 \\ & \quad + \lambda \|\beta\|_0 \\ & \geq Q_0(\beta_1) + \sum_{j \in S_{\beta_1}} \{0.5(C_5 - C_4 a_{n_1} b_{n_1})\epsilon_0(1 - \epsilon_0)(\beta_j - \beta_{1,j})^2 \\ & \quad + \lambda \|\beta_j\|_0\}, \end{aligned} \quad (\text{A.2})$$

where  $\beta_j$  and  $\beta_{1,j}$  are the  $j$ th components of  $\beta$  and  $\beta_1$ , respectively. For  $n_1 \geq N_3^*(\delta)$ ,

$$Q(\beta) \geq Q_0(\beta_1) + \lambda \sum_{j \in S_{\beta_1}} \|\beta_{1,j}\|_0 = Q_0(\beta_1) + \lambda \|\beta_1\|_0.$$

By (A.1), we have

$$Q(\beta_1) = Q_0(\beta_1) + \lambda \|\beta_1\|_0.$$

Therefore,  $\beta_1$  is a local minimizer of  $Q(\beta)$ . It then follows from the convexity of  $Q(\beta)$  that  $\beta_1$  is the global minimizer  $\beta_1^*$  of  $Q(\beta)$ .

*Proof of Theorem 2.* For simplicity, denote by  $\rho(m(Z_1), Y_1)$  the loss function  $\rho_q(X_1, Y_1) = -Y_1 m(Z_1) + \log(1 + \exp(m(Z_1)))$ . Note that

$$\frac{\partial \rho(m(Z_1), Y_1)}{\partial m(Z_1)} = -Y_1 + \frac{\exp(m(Z_1))}{1 + \exp(m(Z_1))} = -Y_1 + \pi_{m, f_0, g_0}(X_1),$$

and

$$\frac{\partial^2 \rho(m(\mathbf{Z}_1), Y_1)}{[\partial m(\mathbf{Z}_1)]^2} = \frac{\exp(m(\mathbf{Z}_1))}{[1 + \exp(m(\mathbf{Z}_1))]^2}.$$

By the second-order Taylor expansion, we obtain that

$$\begin{aligned} \rho(m_{\beta}(\hat{\mathbf{Z}}_1), Y_1) &= \rho(m_{\beta_0}(\mathbf{Z}_1^0), Y_1) + [\partial \rho(m_{\beta_0}(\mathbf{Z}_1^0), Y_1) / \partial m_{\beta}(\mathbf{Z}_1)] \\ &\quad \times (m_{\beta}(\hat{\mathbf{Z}}_1) - m_{\beta_0}(\mathbf{Z}_1^0)) \\ &\quad + \frac{1}{2} \frac{\partial^2 \rho(m^*, Y_1)}{[\partial m_{\beta}(\mathbf{Z}_1)]^2} (m_{\beta}(\hat{\mathbf{Z}}_1) - m_{\beta_0}(\mathbf{Z}_1^0))^2, \end{aligned} \quad (\text{A.3})$$

where  $m^*$  lies between  $m_{\beta}(\hat{\mathbf{Z}}_1)$  and  $m_{\beta_0}(\mathbf{Z}_1^0)$ . Since

$$P \left[ \frac{\partial \rho(m_{\beta_0}(\mathbf{Z}_1^0), Y_1)}{\partial m_{\beta}(\mathbf{Z}_1)} \right] = 0 \quad (\text{A.4})$$

and  $0 < \partial^2 \rho(m^*, Y_1) / [\partial m_{\beta}(\mathbf{Z}_1)]^2 < 1$ , taking the expectation we obtain that

$$\begin{aligned} |P\rho(m_{\beta}(\hat{\mathbf{Z}}_1), Y_1) - P\rho(m_{\beta_0}(\mathbf{Z}_1^0), Y_1)| &< 0.5 P[(m_{\beta}(\hat{\mathbf{Z}}_1) - m_{\beta_0}(\mathbf{Z}_1^0))^2] \\ &= 0.5 P[(\hat{\mathbf{Z}}_1^T \boldsymbol{\beta} - \mathbf{Z}_1^0{}^T \boldsymbol{\beta}_0)^2]. \end{aligned}$$

Hence, from Corollary 1, on the event  $\mathcal{J}_3$ ,

$$\begin{aligned} |P\rho(m_{\beta_0}(\hat{\mathbf{Z}}_1), Y_1) - P\rho(m_{\beta_0}(\mathbf{Z}_1^0), Y_1)| \\ \leq 0.5 \boldsymbol{\beta}_0^T P[(\hat{\mathbf{Z}}_1 - \mathbf{Z}_1^0)(\hat{\mathbf{Z}}_1 - \mathbf{Z}_1^0)^T] \boldsymbol{\beta}_0 \\ \leq C_1 C_4 s_{\beta_0}^2 b_{n_1}, \end{aligned}$$

where  $s_{\beta} = |S_{\beta}|$  is the cardinality of  $S_{\beta} = \{j : \beta_j \neq 0\}$ . Naturally,  $P\rho(m_{\beta_0}(\hat{\mathbf{Z}}_1), Y_1) \leq P\rho(m_{\beta_0}(\mathbf{Z}_1^0), Y_1) + C_1 C_4 s_{\beta_0}^2 b_{n_1}$ .

In addition, by definition of  $\beta_1$ ,  $P\rho(m_{\beta_1}(\hat{\mathbf{Z}}_1), Y_1) = \min_{\beta} P\rho(m_{\beta}(\hat{\mathbf{Z}}_1), Y_1)$ . As a result,  $P\rho(m_{\beta_1}(\hat{\mathbf{Z}}_1), Y_1) \leq P\rho(m_{\beta_0}(\hat{\mathbf{Z}}_1), Y_1)$ . Thus, we have

$$P\rho(m_{\beta_1}(\hat{\mathbf{Z}}_1), Y_1) \leq P\rho(m_{\beta_0}(\mathbf{Z}_1^0), Y_1) + C_1 C_4 s_{\beta_0}^2 b_{n_1}. \quad (\text{A.5})$$

In addition, by (A.3) and (A.4), for any  $\beta$  we have  $P\rho(m_{\beta}(\hat{\mathbf{Z}}_1), Y_1) \geq P\rho(m_{\beta_0}(\mathbf{Z}_1^0), Y_1)$ . Then, setting  $\beta = \beta_1$  on the left side leads to

$$P\rho(m_{\beta_1}(\hat{\mathbf{Z}}_1), Y_1) \geq P\rho(m_{\beta_0}(\mathbf{Z}_1^0), Y_1). \quad (\text{A.6})$$

Combining (A.5) and (A.6) leads to

$$|P\rho(m_{\beta_1}(\hat{\mathbf{Z}}_1), Y_1) - P\rho(m_{\beta_0}(\mathbf{Z}_1^0), Y_1)| \leq C_1 C_4 s_{\beta_0}^2 b_{n_1}. \quad (\text{A.7})$$

As a result, we have

$$\mathcal{E}(m_{\beta_1}, \hat{\mathbf{f}}, \hat{\mathbf{g}}) \leq C_1 C_4 s_{\beta_0}^2 b_{n_1}. \quad (\text{A.8})$$

(A.8) combined with Lemma 4 ( $\beta_1^* = \beta_1$ ) leads to

$$\mathcal{E}(m_{\beta_1^*}, \hat{\mathbf{f}}, \hat{\mathbf{g}}) \leq C_1 C_4 s_{\beta_0}^2 b_{n_1}. \quad (\text{A.9})$$

Recall the oracle estimator

$$\beta_1^* = \arg \min_{\beta \in \mathcal{B}} \left\{ \mathcal{E}(m_{\beta}, \hat{\mathbf{f}}, \hat{\mathbf{g}}) + \frac{8\lambda s_{\beta}}{c\phi_1^2(S_{\beta})} \right\}.$$

Then by Theorem 1,

$$\mathcal{E}(\hat{q}_1) = \mathcal{E}(m_{\hat{\beta}_1}, \hat{\mathbf{f}}, \hat{\mathbf{g}}) \leq 6\mathcal{E}(m_{\beta_1^*}, \hat{\mathbf{f}}, \hat{\mathbf{g}}) + \frac{16\lambda^2 s_{\beta_1^*} (e^{\eta}/\epsilon_0 + 1)^2}{c\phi^2(S_{\beta_1^*})}. \quad (\text{A.10})$$

Therefore, by (A.9) and (A.10),

$$\mathcal{E}(\hat{q}_1) \leq \frac{16\lambda^2 s_{\beta_1^*} (e^{\eta}/\epsilon_0 + 1)^2}{c\phi^2(S_{\beta_1^*})} + C_1 C_4 s_{\beta_0}^2 b_{n_1}.$$

[Received December 2013. Revised October 2014.]

## REFERENCES

Ackermann, M., and Strimmer, K. (2009), "A General Modular Framework for Gene Set Enrichment Analysis," *BMC Bioinformatics*, 10, 1471–2105. [276]

Antoniadis, A., Lambert-Lacroix, S., and Leblanc, F. (2003), "Effective Dimension Reduction Methods for Tumor Classification Using Gene Expression Data," *Bioinformatics*, 19, 563–570. [277]

Bair, E., Hastie, T., Paul, D., and Tibshirani, R. (2006), "Prediction by Supervised Principal Components," *Journal of American Statistical Association*, 101, 119–137. [277]

Bickel, P. J., and Levina, E. (2004), "Some Theory for Fisher's Linear Discriminant Function, 'naive Bayes', and Some Alternatives When There are Many More Variables Than Observations," *Bernoulli*, 10, 989–1010. [275]

Boulesteix, A.-L. (2004), "PLS Dimension Reduction for Classification With Microarray Data," *Statistical Applications in Genetics and Molecular Biology*, 3, 1–30. [277]

Breiman, L. (2001), "Random Forests," *Machine Learning*, 45, 5–32. [278]

Bühlmann, P., and van de Geer, S. (2011), *Statistics for High-Dimensional Data*, New York: Springer. [275, 281]

Cai, T., and Liu, W. (2011), "A Direct Estimation Approach to Sparse Linear Discriminant Analysis," *Journal of American Statistical Association*, 106, 1566–1577. [277]

Clemmensen, L., Hastie, T., Witten, D., and Ersboll, B. (2011), "Sparse Discriminant Analysis," *Technometrics*, 53, 406–413. [277]

Dudoit, S., Fridlyand, J., and Speed, T. P. (2002), "Comparison of Discrimination Methods for the Classification of Tumors Using Gene Expression Data," *Journal of American Statistical Association*, 97, 77–87. [280]

Fan, J., and Fan, Y. (2008), "High-Dimensional Classification Using Features Annealed Independence Rules," *Annals of Statistics*, 36, 2605–2637. [279, 280]

Fan, J., Feng, Y., and Song, R. (2011), "Nonparametric Independence Screening in Sparse Ultra-High Dimensional Additive Models," *Journal of American Statistical Association*, 106, 544–557. [284]

Fan, J., Feng, Y., and Tong, X. (2012), "A Road to Classification in High Dimensional Space: The Regularized Optimal Affine Discriminant," *Journal of the Royal Statistical Society, Series B*, 74, 745–771. [276, 279, 280]

Fan, J., and Li, R. (2001), "Variable Selection via Nonconcave Penalized Likelihood and its Oracle Properties," *Journal of American Statistical Association*, 96, 1348–1360. [278]

Fan, J., and Lv, J. (2008), "Sure Independence Screening for Ultrahigh Dimensional Feature Space" (with discussion), *Journal of the Royal Statistical Society, Series B*, 70, 849–911. [276, 284]

Friedman, J., Hastie, T., Höfling, H., and Tibshirani, R. (2007), "Pathwise Coordinate Optimization," *Annals of Applied Statistics*, 1, 302–332. [278]

Friedman, J., Hastie, T., and Tibshirani, R. (2010), "Regularization Paths for Generalized Linear Models via Coordinate Descent," *Journal of Statistical Software*, 33, 1–22. [279]

Fu, W. J., Carroll, R. J., and Wang, S. (2005), "Estimating Misclassification Error With Small Samples via Bootstrap Cross-Validation," *Bioinformatics*, 21, 1979–1986. [278]

Gordon, G. J., Jensen, R. V., Hsiao, L.-L., Gullans, S. R., Blumenstock, J. E., Ramaswamy, S., Richards, W. G., Sugarbaker, D. J., and Bueno, R. (2002), "Translation of Microarray Data Into Clinically Relevant Cancer Diagnostic Tests Using Gene Expression Ratios in Lung Cancer and Mesothelioma," *Cancer Research*, 62, 4963–4967. [280]

Guo, Y., Hastie, T., and Tibshirani, R. (2007), "Regularized Linear Discriminant Analysis and its Application in Microarrays," *Biostatistics*, 8, 86–100. [277]

Hastie, T., and Tibshirani, R. (1990), *Generalized Additive Models*, London: Chapman & Hall/CRC. [277]

Hastie, T., Tibshirani, R., and Friedman, J. H. (2009), *The Elements of Statistical Learning: Data Mining, Inference, and Prediction* (2nd ed.), New York: Springer-Verlag Inc. [275, 279]

Huang, P. W. X. (2003), "Linear Regression and Two-Class Classification With Gene Expression Data," *Bioinformatics*, 19, 2072–2978. [277]

Li, K.-C. (1991), "Sliced Inverse Regression for Dimension Reduction" (with discussion), *Journal of American Statistical Association*, 86, 316–342. [277]

Mai, Q., Zou, H., and Yuan, M. (2012), "A Direct Approach to Sparse Discriminant Analysis in Ultra-High Dimensions," *Biometrika*, 99, 29–42. [277]

Meier, L., Geer, V., and Bühlmann, P. (2009), "High-Dimensional Additive Modeling," *Annals of Statistics*, 37, 3779–3821. [277, 279]

Nguyen, D. V., and Rocke, D. M. (2002), "Tumor Classification by Partial Least Squares Using Microarray Gene Expression Data," *Bioinformatics*, 18, 39–50. [277]

- Polonik, W. (1995), "Measuring Mass Concentrations and Estimating Density Contour Clusters-An Excess Mass Approach," *Annals of Statistics*, 23, 855–881. [282]
- Raykar, V., Duraiswami, R., and Zhao, L. (2010), "Fast Computation of Kernel Estimators," *Journal of Computational and Graphical Statistics*, 19, 205–220. [278]
- Shao, J., Wang, Y., Deng, X., and Wang, S. (2011), "Sparse Linear Discriminant Analysis by Thresholding for High Dimensional Data," *Annals of Statistics*, 39, 1241–1265. [277]
- Tibshirani, R. (1996), "Regression Shrinkage and Selection via the Lasso," *Journal of the Royal Statistical Society, Series B*, 58, 267–288. [278]
- Tong, X. (2013), "A Plug-in Approach to Anomaly Detection," *Journal of Machine Learning Research*, 14, 3011–3040. [285]
- Tsybakov, A. (2004), "Optimal Aggregation of Classifiers in Statistical Learning," *Annals of Statistics*, 32, 135–166. [282]
- van de Geer, S. (2008), "High-Dimensional Generalized Linear Models and the Lasso," *Annals of Statistics*, 36, 614–645. [282,283]
- Witten, D., and Tibshirani, R. (2012), "Penalized Classification Using Fisher's Linear Discriminant," *Journal of the Royal Statistical Society, Series B*, 73, 753–772. [277]
- Wu, M. C., Zhang, L., Wang, Z., Christiani, D. C., and Lin, X. (2009), "Sparse Linear Discriminant Analysis for Simultaneous Testing for the Significance of a Gene Set/Pathway and Gene Selection," *Bioinformatics*, 25, 1145–1151. [277]
- Zhang, C.-H. (2010), "Nearly Unbiased Variable Selection Under Minimax Concave Penalty," *Annals of Statistics*, 38, 894–942. [278]
- Zhang, C.-H., and Huang, J. (2008), "The Sparsity and Bias of the LASSO Selection in High-Dimensional Linear Regression," *Annals of Statistics*, 36, 1567–1594. [281]
- Zhao, P., and Yu, B. (2006), "On Model Selection Consistency of Lasso," *Journal of Machine Learning Research*, 7, 2541–2563. [281]
- Zou, H. (2006), "The Adaptive Lasso and its Oracle Properties," *Journal of American Statistical Association*, 101, 1418–1429. [278]
- Zou, H., Hastie, T., and Tibshirani, R. (2006), "Sparse Principal Component Analysis," *Journal of Computational and Graphical Statistics*, 15, 265–286. [277]