Published in final edited form as:

Stat Anal Data Min. 2015 February ; 8(1): 49–63. doi:10.1002/sam.11259.

Feature Import Vector Machine: A General Classifier with Flexible Feature Selection

Samiran Ghosh^{1,2} and Yazhen Wang³

Samiran Ghosh: sghos@med.wayne.edu

¹Department of Family Medicine & Public Health Sciences, Wayne State University

²Center of Molecular Medicine and Genetics, Wayne State University

³Department of Statistics, University of Wisconsin, Madison

Abstract

The support vector machine (SVM) and other reproducing kernel Hilbert space (RKHS) based classifier systems are drawing much attention recently due to its robustness and generalization capability. General theme here is to construct classifiers based on the training data in a high dimensional space by using all available dimensions. The SVM achieves huge data compression by selecting only few observations which lie close to the boundary of the classifier function. However when the number of observations are not very large (small n) but the number of dimensions/features are large (large p), then it is not necessary that all available features are of equal importance in the classification context. Possible selection of an useful fraction of the available features may result in huge data compression. In this paper we propose an algorithmic approach by means of which such an *optimal* set of features could be selected. In short, we reverse the traditional sequential observation selection strategy of SVM to that of sequential feature selection. To achieve this we have modified the solution proposed by Zhu and Hastie (2005) in the context of import vector machine (IVM), to select an *optimal* sub-dimensional model to build the final classifier with sufficient accuracy.

Keywords

Classification; Import Vector Machine; Radial Basis Function; Regularization; Reproducing Kernel Hilbert Space; Support Vector Machine

1 Introduction

Many machine learning as well as nonparametric function estimation methods can be recast into a unified regularization based modeling framework. Almost all the present kernel machine based approaches try to do function estimation by representing the original function as a linear combination of the basis functions in the higher dimensional space. The popular support vector machine (SVM) is a member of this framework. SVM achieves huge data compression by selecting only few support points (i.e. observations) that lies near the

boundary of the classifier function. However domains such as bioinformatics which generally produces thousands of dimensions/features/covariates (p) and only few observations (n) are well suited for kernel machine based function estimation if compression can be achieved in feature space (i.e. in p) rather than in terms of observations (i.e. in n). Another advantage of such feature selection is biomarker discovery, which generally corresponds to only few selected features. In this paper we propose a new kernel machine based on regularization principle which achieves function estimation with sufficient accuracy by selecting only few features in the original input space and thus named feature import vector machine (FIVM). The basic idea (see Lin et al. (2000); Smola and Scholkopf (2000)) of selecting a sub model in place of a complicated full model is not new. Recently Zhu and Hastie (2005) proposed import vector machine (IVM) as an alternative to the popular SVM, which is built on kernel logistic regression (KLR). The IVM has following advantages over standard SVM:

- Probabilistic interpretation of the classification result
- Selection of few observations, hence SVM like data compression
- Straightforward generalization of IVM for multi-class classification.

The most innovative feature of IVM is replacement of the SVM loss by the negative log-likelihood (NLL) of the binomial distribution. However in doing so it destroys the sparse observation selection property of SVM. To solve this problem Zhu and Hastie (2005) proposed a sequential selection (a variant of greedy search) strategy to select only few *important* observations named as import vectors. We pose the same problem if we exchange the role played by n and p in IVM. Hence the idea is to create a dimension/feature screening/selection methodology via a sequential search strategy over the original feature space. Our proposal has following features:

- It uses kernel machine for classifier construction
- It produces a nonlinear classification boundary in the original input space
- The feature selection is done in the original input space, not in the kernel transformed feature space
- Unlike SVM (both L_1 and L_2), the result has probabilistic interpretation.

Though the feature selection via SVM in the transformed kernel space or only via a linear kernel (see Nguyen and Torre (2010) and references therein) is a well studied problem, to the best of our knowledge except for the only one article (Weston et al., 2000), there exist very limited attempt where feature selection is done in SVM in the original space, when a non-linear kernel function is in use. However as reported in Nguyen and Torre (2010) optimization problem presented in Weston et al. (2000) is potentially non-convex due to the placement of feature scaling factor inside the kernel (see also Canu and Grandvalet (2002)), as a result computational and resulting convergence is questionable. Moreover except for empirical observations no theoretical justification is provided in Weston et al. (2000), neither about the optimality of the selected features nor about the classification accuracy under any setup. If classification can be done using non-linear kernel but feature selection in the original space, this leads to tremendous advantage as in one hand we are able to combine the

power of non-linear classification via kernel machine, on the other hand important features are selected in the original input space, thus making their physical interpretation straight forward.

The rest of the paper is organized as follows. A brief introduction of the RKHS is provided in Section 2. We also briefly describe some key properties of SVM and KLR in Section 3. Dimension screening and some key optimality properties are described in Section 4. Here we also formally propose FIVM algorithm. Section 5 illustrates our proposed methodology on some synesthetic as well as real life classification examples from bioinformatics domain. Multiclass extension of the proposed methodology is briefly discussed in Section 6. We conclude the article with a discussion in Section 7. For brevity all the proofs are provided in the appendix.

2 Reproducing Kernel Hilbert Space and Function Representation

Formally speaking, let $T = (\mathbf{x}_i, y_i)$, $1 \leq i \leq n$, are i.i.d. pairs distributed as $(\mathbf{x}, y)$ with probability measure $\mathbb{P}$, where $y_i \in C = \{-1, 1\}$ or $y_i \in \mathbb{R}$ and $\mathbf{x}_i = (x_1, \dots, x_p) \in \mathbb{R}^p$. A general class of regularization problems has the form

$$\operatorname{argmin}_{f \in \mathcal{H}} \left[\frac{1}{n} \sum_{i=1}^n L(y_i, f(\mathbf{x}_i)) + \lambda J(f) \right], \quad (2.1)$$

where $L(y_i, f(\mathbf{x}_i))$ is a convex loss function, $J(f)$ is a penalty functional, $\lambda > 0$ is the smoothing or regularization parameter, and $\mathcal{H}$ is a space of functions on which $J(f)$ is defined. An important subclass of the form (2.1) is generated by a positive definite kernel $K(\mathbf{x}, \mathbf{x}')$ and the corresponding space of functions $\mathcal{H}_k$ is called reproducing kernel Hilbert space (RKHS). In this article we will employ radial basis function (RBF) as kernel which is

given by, $K(\mathbf{x}_i, \mathbf{x}_j) = \exp \left\{ -\frac{\|\mathbf{x}_i - \mathbf{x}_j\|^2}{2\theta^2} \right\}$. For special case of squared norm regularizer, by the important ‘‘Representer Theorem’’ of Kimeldorf and Wahba (Kimeldorf and Wahba, 1971; Wahba, 1990), any solution to the above problem is finite-dimensional, and has the form

$$f(\mathbf{x}) = \sum_{i=1}^n \alpha_i K(\mathbf{x}, \mathbf{x}_i), \quad (2.2)$$

where $J(f) = \|f\|_{\mathcal{H}_k}^2$. In the light of (2.2), (2.1) can be rewritten as

$$\operatorname{argmin}_{\alpha \in \mathbb{R}^p} \left[\frac{1}{n} \sum_{i=1}^n L(y_i, f(\mathbf{x}_i)) + \lambda \|f\|_{\mathcal{H}_k}^2 \right]. \quad (2.3)$$

Smoothing spline ANOVA, Penalized polynomial regression, SVM, IVM, KLASO (Kernelized version of LASSO proposed by Roth (2004)) are all examples of RKHS based

methodology. For a more rigorous treatment in this matter please refer to Wahba (1990), Parzen (1970) and Aronszajn (1950).

3 SVM, Kernel Logistic Regression and Lasso

The standard SVM (Vapnik, 1995) belongs to a general class of regularization problems described earlier in (2.3), where $L(y_i, f(\mathbf{x}_i))$ is chosen as the hinge loss function. The SVM could be cast in the form of minimizing a loss function as a function of the margins $yf(\mathbf{x})$ (see Evgeniou et al. (1999); Rosset et al. (2004)). Fitting SVM is equivalent to finding

$$\operatorname{argmin}_{f \in \mathcal{H}_k} \frac{1}{n} \sum_{i=1}^n [1 - y_i f(\mathbf{x}_i)]_+ + \lambda \|f\|_{\mathcal{H}_k}^2. \quad (3.1)$$

The standard SVM produces a nonlinear classification boundary in the original input space via kernel transformation on the features. The dimension of the transformed feature space is often very high and may be infinite for some cases. However it employs *kernel trick* on a positive definite reproducing kernel to achieve this seemingly impossible computation. The optimal solution of the (3.1) is given by equation (2.2). It turns out that for most cases a sizeable number of α_i 's are zero, which seems an attractive property for data compression. Nonzero $\mathbf{x}_i$'s are known as support points and equation (2.2) is known as sparse representation of $f(\mathbf{x})$. Recently Lin (2002) has shown that the SVM implements Bayes' rule asymptotically.

However SVM has two major drawbacks. First, it works well for two-class problems. Recently different researchers (see Vapnik (1998); Weston and Watkins (1999); Lee et al. (2004)) have proposed methodologies for its possible multiclass extension. However the success of the proposed methodologies is questionable. Second, SVM can only estimate sign

$[p(\mathbf{x}) - \frac{1}{2}]$, while the quantity $p(\mathbf{x}) = P(Y=1|X=\mathbf{x})$ giving classification probability is often of interest by itself. Noting the similarity of the hinge loss of SVM and that of the NLL of the binomial distribution (plotted in Figure 1), Zhu and Hastie (2005) proposed to replace the hinge loss in equation 3.1. This essentially produces kernel logistic regression (KLR) given by:

$$\operatorname{argmin}_{f \in \mathcal{H}_k} \frac{1}{n} \sum_{i=1}^n \ln[1 + e^{-y_i f(\mathbf{x}_i)}] + \lambda \|f\|_{\mathcal{H}_k}^2. \quad (3.2)$$

The KLR eliminates two drawbacks of SVM stated above. However it no longer enjoys support point property of the SVM, thus making all of the α_i 's non-zero in equation (2.2). Zhu and Hastie (2005) suggested an idea based on selecting an approximate sub-model to overcome this difficulty. Their algorithm termed as import vector machine (IVM) uses both y_j and input $\mathbf{x}_j$ to select a subset of T to approximate the full model. However for feature selection we face a problem of different kind, where selection of dimension (p) is of prime interest. For example in large p small n domain, reduction of dimension (p) is more important than reduction of the number of observations (n). Simultaneous screening of p and

n is another possibility in the recently popular *Big data* context. However in this paper our focus is exclusively on dimension/feature selection. In the variable selection context Lasso proposed by Tibshirani (1996) is a very successful method for automatic feature selection. In penalized regression context, L_1 penalized methods offer automatic dimension selection owing to the nature of the penalty function. However though Lasso is successful in many situations, it has limitations. For example in the context of large p small n domain ($p \gg n$), Lasso can select at most n dimensions only. Also in situations where two (or more) dimensions have high correlation, Lasso tends to select only one dimension from the group. Park and Hastie (2008) considered NLL of the binomial distribution with L_1 -penalty for dimension selection in the classification framework, however they do not employ any kernel transformation. Feature selection in the transformed feature space is another possibility, however owing to the non-linearity of the kernel it is often impossible to map the direct relationship between transformed and original features. Moreover selection via L_1 -penalty will inherit all merits as well as drawbacks of Lasso discussed so far. The FIVM algorithm that we are proposing here does not have these drawbacks as it does not rely on a specific penalty function for feature selection. Depending upon the choice of the convergence parameter (ε introduced latter) as many/few dimensions can be selected, as desired.

4 Feature Selection in KLR Framework

Let us denote the index set $\mathcal{L} = \{1, 2, \dots, p\}$ as the collection of all available features/ dimensions in the original input space and $\mathcal{S} \subseteq \mathcal{L}$ is an arbitrary subset of these available features such that $q = \text{card}(\mathcal{S}) \leq p$. Given an index set $\mathcal{S}$, it will also induce a corresponding feature space which we denote by $\mathcal{F}_{\mathcal{S}}$. For the RBF, a kernel on any arbitrary collection of

dimensions such as $\mathcal{S}$ can be defined as $K_{\mathcal{S}}(\mathbf{x}_i, \mathbf{x}_j) = \exp \left\{ -\frac{\sum_{t \in \mathcal{S}} (x_{it} - x_{jt})^2}{2\theta^2} \right\}$. We can rephrase the original regularization problem on this arbitrary set $\mathcal{S}$ as to obtain

$$\underset{\alpha \in \mathbf{R}^q}{\text{argmin}} \left[\frac{1}{n} \sum_{i=1}^n L_{\mathcal{S}}(y_i, \alpha_i, \mathbf{x}_i) + \lambda \sum_{i,j=1}^n \alpha_i \alpha_j K_{\mathcal{S}}(\mathbf{x}_i, \mathbf{x}_j) \right]. \quad (4.1)$$

It is easy to note that when $\mathcal{S} = \mathcal{L}$, problem (4.1) reduces to (2.3). Now using Wahba's representer theorem, solution for (4.1) is given by

$$f_{\mathcal{S}}(\mathbf{x}) = \sum_{i=1}^n \alpha_i K_{\mathcal{S}}(\mathbf{x}, \mathbf{x}_i). \quad (4.2)$$

We want to study the behavior of the function $f_{\mathcal{S}}(\mathbf{x})$ as $\mathcal{S} \rightarrow \mathcal{L}$. More specifically speaking, whether the addition of more dimensions in $\mathcal{S}$ will increase the classifiers accuracy. Next we propose a theorem under the assumption that the transformed feature space is so rich that the training data are completely separable.

Theorem 4.1. If training data is separable in $\mathcal{S}$ then it will be separable in any higher dimensional space provided it is a superset of $\mathcal{S}$.

The advantage of selecting a submodel based on few features is multi-fold. Not only computational burden is greatly reduced but also it obeys the principle of parsimony without jeopardizing classification accuracy. Theorem 4.1 shows that separating hyperplane in $\mathcal{S}$ is also a separating hyperplane in $\mathcal{L}$. Though it does not guarantee that it is also a margin maximizing hyperplane in $\mathcal{L}$. Under complete separability of the training data set in $\mathcal{L}$, theoretical justification of using a submodel in $\mathcal{S}(\subseteq \mathcal{L})$ is provided next. Suppose the dictionary of the basis functions of the transformed feature space is

$$\{h_1(\mathbf{x}), h_2(\mathbf{x}), \dots, h_l(\mathbf{x})\},$$

where l denotes the dimension of the transformed feature space. The classification boundary, which is a hyperplane in the transformed feature space, is given by,

$$\{\mathbf{x}: f(\mathbf{x}) = \beta_0 + \mathbf{h}(\mathbf{x})^T \boldsymbol{\beta} = 0\}.$$

Consider the original full feature space $\mathcal{F}_{\mathcal{L}}$ and corresponding index set $\mathcal{L}$. For simplicity and without loss of generality, we assume the first q coordinates are true features (or signals) and $\mathbf{x}^* = (x_1, \dots, x_q)$, $q < p$, while remaining dimensions ($> q$) contains just noise. The classification boundary is given by a hyperplane in the true feature space $\mathcal{F}_{\mathcal{S}}$ as

$$\{\mathbf{x}^*: f(\mathbf{x}^*) = \beta_0 + \mathbf{h}(\mathbf{x}^*)^T \boldsymbol{\beta}^* = 0\}.$$

As in Theorem 4.1 suppose training data $\mathbf{x}^*$ in $\mathcal{F}_{\mathcal{S}}$ are separable, then the margin-maximizing SVM in $\mathcal{S}$ is:

$$Q(q, \mathbf{x}^*, \mathcal{S}) = \max_{\beta_0, \boldsymbol{\beta}^*, \|\boldsymbol{\beta}^*\|=1} D^* \text{ subject to } y_i(\beta_0 + \mathbf{h}(\mathbf{x}_i^*)^T \boldsymbol{\beta}^*) \geq D^*, i=1, \dots, n$$

where D^* is the shortest distance from the training data to the separating hyperplane in $\mathcal{F}_{\mathcal{S}}$. If training data $\mathbf{x}^*$ in $\mathcal{S}$ are separable, then the whole training data $\mathbf{x}$ in $\mathcal{L}$ are also separable by Theorem 4.1. The margin-maximizing SVM in $\mathcal{L}$ reduces to:

$$Q(p, \mathbf{x}, \mathcal{L}) = \max_{\beta_0, \boldsymbol{\beta}, \|\boldsymbol{\beta}\|=1} D, \text{ subject to } y_i(\beta_0 + \mathbf{h}(\mathbf{x}_i)^T \boldsymbol{\beta}) \geq D, i=1, \dots, n$$

where D is the shortest distance from the training data to the separating hyperplane in $\mathcal{F}_{\mathcal{L}}$. A KLR problem in $\mathcal{L}$ can be equivalently stated as obtaining

$$\min_{\beta_0, \beta} \frac{1}{n} \sum_{i=1}^n \ln \left(1 + e^{-y_i f(\mathbf{x}_i)} \right) \text{ subject to } \|\beta\|_2^2 \leq s, f(\mathbf{x}_i) = \beta_0 + \mathbf{h}(\mathbf{x}_i)^T \beta, i=1, \dots, n$$

where β_0 is the bias term and may be omitted for simplification purposes. Let us denote the solution of the above KLR problem as $\tilde{\beta}(s)$. Equivalently for any $\mathcal{S}$ having only first q (p) dimensions of $\mathcal{L}$, the KLR problem in $\mathcal{S}$ is

$$\min_{\beta_0, \beta^*} \frac{1}{n} \sum_{i=1}^n \ln \left(1 + e^{-y_i f(\mathbf{x}_i^*)} \right) + \lambda \|\beta\|_{\mathcal{H}_k}^2.$$

Above is equivalent to

$$\min_{\beta_0 < \beta^*} \frac{1}{n} \sum_{i=1}^n \ln \left(1 + e^{-y_i f(\mathbf{x}_i^*)} \right) \text{ subject to } \|\beta^*\|_2^2 \leq s, f(\mathbf{x}_i^*) = \beta_0 + \mathbf{h}(\mathbf{x}_i^*)^T \beta^*, i=1, \dots, n. \quad (4.3)$$

Let us denote its solution by $\hat{\beta}^*(s)$. Then for the submodel in $\mathcal{S}$ we have two propositions followed by Theorem 4.2.

Proposition 4.1. $Q(p, \mathbf{x}, \mathcal{L}) = Q(q, \mathbf{x}^*, \mathcal{S})$.

On the other hand, when the true feature space is $\mathcal{F}_\mathcal{S}$ and feature information is stored in the first q coordinates of $\mathbf{x}$ and its last $p - q$ coordinates are noises, we can partition

$\mathbf{x}_{n \times p} = (\mathbf{x}_{n \times q}^*, \mathbf{x}_{n \times p-q}^\dagger)$, where $\mathbf{x}^\dagger$ is a noise vector. We define a robust version of margin in the whole space $\mathcal{F}_\mathcal{L}$ to guard against all possible noises:

$$Q(\mathbf{x}^*, \mathcal{L}) = \max_{\mathbf{x}^\dagger} Q(p, (\mathbf{x}^*, \mathbf{x}^\dagger), \mathcal{L}),$$

that is, the shortest distance from the training data to the hyperplane separating $\mathbf{x} = (\mathbf{x}^*, \mathbf{x}^\dagger)$ for all noise $\mathbf{x}^\dagger$ cases. Denote the first q coordinates of solution β by β^* .

Proposition 4.2. $Q(\mathbf{x}^*, \mathcal{L}) = Q(q, \mathbf{x}^*, \mathcal{S})$, and thus $\hat{\beta}^* = \tilde{\beta}^*$.

Theorem 4.2. If training data is separable in $\mathcal{S}(\subseteq \mathcal{L})$ and the solution for the equivalent KLR problem in $\mathcal{S}$ and $\mathcal{L}$ are respectively $\hat{\beta}^*(s)$ and $\tilde{\beta}(s)$, then as $s \rightarrow \infty$

$$\hat{\beta}^*(s)/s - \tilde{\beta}^*(s)/s \rightarrow \hat{\beta}^* - \tilde{\beta} = 0.$$

where $\tilde{\beta}^*(s)$ are the first q coordinates of $\tilde{\beta}(s)$.

Theorem 4.2 shows optimality of the classifier under complete separability assumption in the subspace $\mathcal{F}_\mathcal{S}$.

4.1 Feature Import Vector Machine

Since the KLR lies in the heart of the FIVM, first we would like to describe the basic steps of KLR, where we want to minimize

$$H = \frac{1}{n} \sum_{i=1}^n \ln [1 + e^{-y_i f(\mathbf{x}_i)}] + \lambda \|f\|_{\mathcal{H}_k}^2. \quad (4.4)$$

Let us define,

$$p_i = \frac{1}{1 + e^{-y_i f(\mathbf{x}_i)}}, \text{ for } i=1, 2, \dots, n, \quad (4.5)$$

$$\mathbf{a} = (a_1, a_2, \dots, a_n)', \quad (4.6)$$

$$\mathbf{p} = (p_1, p_2, \dots, p_n)', \quad (4.7)$$

$$\mathbf{y} = (y_1, y_2, \dots, y_n)', \quad (4.8)$$

$$\mathbf{K} = (K(\mathbf{x}_i, \mathbf{x}_j))_{i,j=1}^n, \quad (4.9)$$

$$\mathbf{W} = \text{Diag}(p_1(1-p_1), p_2(1-p_2), \dots, p_n(1-p_n)). \quad (4.10)$$

Following Zhu and Hastie (2005) with little abuse of notations, equation (4.4) can be rewritten in matrix notations as

$$H = \frac{1}{n} \mathbf{1}' \ln(\mathbf{1} + e^{-\mathbf{y} \cdot \mathbf{K} \mathbf{a}}) + \lambda \mathbf{a}' \mathbf{K} \mathbf{a}, \quad (4.11)$$

where “ $\cdot$ ” represents element-wise multiplication. To find $\mathbf{a}$, we set the derivative of H with respect to $\mathbf{a}$ equals to zero and use the Newton-Raphson method to iteratively solve the score equation. With a little bit of algebra, it can be shown that Newton-Raphson step is a weighted least square method

$$\mathbf{a}^{(k)} = \left(\frac{1}{n} \mathbf{K}' \mathbf{W} \mathbf{K} + \lambda \mathbf{K} \right)^{-1} \mathbf{K}' \mathbf{W} \mathbf{z},$$

where $\mathbf{a}^{(k)}$ is the value of $\mathbf{a}$ obtained in the k th step and

$$\mathbf{z} = \frac{1}{n} \left(\mathbf{K} \mathbf{a}^{(k-1)} + \mathbf{W}^{-1} (\mathbf{y} \cdot \mathbf{p}) \right).$$

Alternatively other minimization method can be used directly to minimize (4.4) to obtain $\mathbf{a}$. Quasi-Newton method implemented in IMSL (Numerics, 1987) library is an interesting possibility. The advantage of this method is that no differentiation and matrix inversion is necessary. Details of other nonlinear function optimization methods can be found in Dennis and Schnabel (Dennis and Schnabel, 1983) and reference therein.

4.2 FIVM Algorithm

As mentioned earlier we want to find a subset $\mathcal{S}$ of $\mathcal{L}$ such that it is a good approximation of the full model. For simplicity, we assume that the observations on the $\mathbf{x}$ are normalized and centered, that is $\frac{1}{n} \sum_{i=1}^n x_{ij}^2 = 1$ and $\frac{1}{n} \sum_{i=1}^n x_{ij} = 0$. This also ascertains same scaling factor (θ) across all dimensions. We use a greedy search technique to bypass the huge combinatorial search problem involving too many features. We start with the null model or $\mathcal{S} = \emptyset$, then iteratively build up $\mathcal{S}$ by adding one feature at a time. Our main goal is to add a feature from $\mathcal{L} \setminus \mathcal{S}$ which will produce minimum value of the regularized negative log-likelihood, until it satisfies some stopping criteria.

Algorithm 1: FIVM Algorithm

1. Let $\mathcal{S} = \emptyset$, $\mathcal{L} = \{1, 2, \dots, p\}$ and $k = 1$.
2. For each $l \in \mathcal{L} \setminus \mathcal{S}$, let

$$f_l(\mathbf{x}) = \sum_{i=1}^n a_i K_{\mathcal{S} \cup l}(\mathbf{x}, \mathbf{x}_i), \text{ where } K_{\mathcal{S} \cup l}(\mathbf{x}_i, \mathbf{x}_j) = \exp \left\{ -\frac{\sum_{t \in \mathcal{S} \cup l} (x_{it} - x_{jt})^2}{2\theta^2} \right\}.$$

Define $\mathbf{K}_1 = (K_{\mathcal{S} \cup l}(\mathbf{x}_i, \mathbf{x}_j))_{i,j=1}^n$. Use Newton-Raphson or other function minimization method to find $\mathbf{a}$ which minimizes,

$$H(\mathbf{x}_l) = \frac{1}{n} \mathbf{1}' \log(\mathbf{1} + e^{-\mathbf{y} \cdot \mathbf{K}_1 \mathbf{a}}) + \lambda \mathbf{a}' \mathbf{K}_1 \mathbf{a}. \quad (4.12)$$

3. Find l^* such that

$$l^* = \operatorname{argmin}_{l \in \mathcal{L}} H(\mathbf{x}_l) \quad (4.13)$$

Let $\mathcal{S} = \mathcal{S} \cup \{l^*\}$, $\mathcal{L} = \mathcal{L} \setminus \{l^*\}$, $H_k = H(\mathbf{x}_{l^*})$ and $k = k + 1$.

4. Repeat Steps 2 and 3 until convergence criteria are satisfied.

The dimensions in $\mathcal{S}$ are called imported features.

4.3 Convergence Criteria

In their original IVM algorithm, Zhu and Hastie (2005) compared the quantity H_k in different iterations. At step k they compare H_k with $H_{k-\Delta k}$, where Δk is a pre chosen small integer, say $\Delta k = 1$. If the ratio $|\frac{H_k - H_{k-\Delta k}}{H_k}|$ is less than a pre chosen small number ε , say $\varepsilon = 0.001$ the algorithm stops adding new observations. This convergence criterion is fine in IVM context as their algorithm compares individual observations without altering dimensions. For FIVM in a specific iteration different $H(\mathbf{x}_j)$'s (Step 2 of the Algorithm 1) are comparable as dimension remains fixed. However individual H_k 's are not comparable across different iterations as each of them becomes trans-dimensional quantity. For this kind of trans-dimensional comparison we need to think of a quantity independent of the dimensional effect. To do so in step 3 of the Algorithm 1, at the k -th iteration after selecting I^* we compute $p_k - p_{I^*}$ defined as the proportion of correctly classified training observations with k imported features. If the ratio $|\frac{p_k - p_{k-\Delta k}}{p_k}|$ is less than a pre chosen small number ε , say $\varepsilon = 0.001$ the algorithm stops adding new features.

Though it was not discussed in their original paper we would like to mention that the convergence criterion described in Zhu and Hastie (2005) has a mild assumption of no repetition of observations to be successfully applicable. In case there are two (or more) identical data points and $\Delta k = 1$ is chosen and if it turns out that in one of those identical points is also an imported observation in any step of the iteration, then the algorithm will

stop there as it will produce $|\frac{H_k - H_{k-\Delta k}}{H_k}| = 0 < \varepsilon$. For FIVM we make two similar assumptions. First, there are no two features having exactly the same value for all observations in the training set. Second, there exists no feature having a constant value throughout. It is trivial to prove that for these cases $|\frac{p_k - p_{k-\Delta k}}{p_k}| = 0 < \varepsilon$, hence the algorithm will not be successful. Note that for very small sample size p_k could be an unstable quantity. While this is true, the way we define relative ratio of p_k and $p_{k-\Delta k}$ (to compare with ε) will somewhat normalize this numerical stability issue.

4.4 Choice of the Regularization parameter λ

Our goal is to construct a classifier which will not only reduce the misclassification error in training set, but also produce good overall generalization capability. To choose optimum value of λ we closely follow the computational consideration of Zhu and Hastie (2005). Following their suggestions we randomly split all the data into a training set and a tuning set. We choose optimal value of λ by decreasing it from a larger value to a smaller value until we hit the optimum misclassification value in the tuning set. Details of these steps could be found in Zhu and Hastie (2005) and hence not repeated here.

5 Experimental Result

In this section we test the FIVM algorithm on few experimental data sets. Our first data set is a synthetic one, where we deliberately introduce few extra features which essentially contain white noise. Our goal is to see whether the algorithm can filter the informative

features from the noisy ones. Second, we consider a real life genomic breast cancer data studied earlier in classification context through singular value decomposition (SVD) by West et al. (2001). Finally we consider Colon Cancer data set of Alon et al. (1999), studied and benchmarked extensively in classification and clustering context.

5.1 Exploration With Synthetic Data

The data is generated following mixture of normal distributions. Following the suggestion of Hastie et al. (Zhu and Hastie, 2005; Hastie et al., 2001), we first generate 10 means μ_k from a bivariate normal distribution $\mathcal{N}(1.5, 0)', 2\mathbf{I})$ and label this class as +1. Similarly we generate another 10 means from $\mathcal{N}(1.5, 0)', 2\mathbf{I})$ and label them as -1. Then from each class we generate 100 observations as follows: For each observation a mean μ_k is chosen with probability 1/10, and then generate a $\mathcal{N}(\mu_k, \mathbf{I}/5)$, which leads to a mixture of normal clusters for each class. Next for each observations we deliberately add another 8 dimensions and fill them with white noise. So our final data set has 200 observations and each observation has 10 dimensions/features. Scatter plots for these data points are represented in Figure 2(a), (b) and (c).

We expect FIVM algorithm to filter out informative features from the noisy ones. Figure 3(a) and Table 1 shows how the tuning parameter λ is chosen. The optimal λ is found to be 0.5 which corresponds to misclassification rate in training set as 0.105. Notably only first two dimensions are selected and we keep $\varepsilon = 0.05$ fixed. Figure 3(b) represents number of selected dimensions for different values of ε . If we choose very small value of $\varepsilon (< 0.0001)$ naturally all the dimensions will be selected on the other hand for very large value of ε only one dimension will be enough. Clearly from the graph optimal cutoff value will lie between 0.02 to 0.06, which justifies our choice for $\varepsilon = 0.05$. Figure 3(c) represents the training error rate if we forcefully add more dimension by lowering ε value. There is very minor improvement when more than two dimensions are selected as they contain white noise.

We next present testing data set in Figure 4(a). Testing data is generated in similar fashion as that of training data. However only those dimensions (namely first two) are used for classification which are selected in the training data by FIVM algorithm. It should be noted that the order of the informative dimensions is a non-issue as FIVM will search and select (see step 3 of the Algorithm 1) those dimensions sequentially which decreases regularized NLL the most. Figure 4(b) represents plot after classification. In Table 1 we summarize the performance of FIVM, SVM and L_1 penalized logistic regression. The scheme of representing testing/training error is following: number without bracket represents the mean while inside the bracket (if any) is standard error. Mean testing error for FIVM is 0.125. Notably for SVM with Gaussian kernel this is 0.18. For L_1 penalized logistic regression testing error is 0.145. Ordinary SVM will not do any dimension selection and performance of SVM does not improve if we include all dimensions in the present setup. We next increase number of testing sample size to see if the testing error rate gets effected for large number of data points. Figure 4(c) represents a plot for testing error rate for different number of testing data set. As seen in this plot testing error rate is a decreasing function of sample size. To consider the generalization capability of FIVM further, we also generated ten separate synthetic training and test data sets each with 200 observations with different

random seeds as described in Section 5.1. The average training error for FIVM is $0.106(\pm 0.016)$ while average testing error is $0.175(\pm 0.019)$. In each case FIVM is able to select two informative dimensions for appropriate choice of λ . We would like to mention in the synthetic data set we have considered $p = 10$, with only two informative dimensions. To consider $n < p$ scenario it is possible to add as many noisy dimensions desired. We would like to report addition of these noisy dimensions will not affect FIVM's feature selection accuracy, thus making it useful for both $n < p$ and $p < n$ scenario. In Table 1 we also report the exact runtime based on a single processor Windows-7-based machine. FIVM is slower than L_1 penalized logistic regression but its performance is comparable to SVM. This is not surprising since both SVM and FIVM are non-linear techniques, however in our experience this computation cost is not too prohibitive.

5.2 Exploration With Breast Cancer Data

We next test our algorithm with a real life data set. This data set is studied earlier by West et al. (2001) in classification context through SVD with stochastic regularization by using Bayesian analysis. Their basic idea was to explore whether gene expression information generated by DNA microarray analysis of human tumors can provide molecular phenotyping that identifies distinct tumor classifications. Tumors were either positive for both the estrogen and progesterone receptors or negative for both receptors. The final collection of tumors consisted of 13 estrogen receptor (ER)+lymph node (LN)+tumors, 12 ER–LN+tumors, 12 ER+LN–tumors, and 12 ER– LN–tumors. Other biological details can be found in the original paper of West et al. (2001). The initial 49 tumors were classified as ER+ or ER– via immunohistochemistry (IHC) at time of diagnosis and then later via protein immunoblotting assay for ER to check the IHC results. As reported in their paper in five cases these results are found conflicting. We have selected these five and another nine samples randomly to create our test set. So on the basis of ER status our training set consist of 35 samples while the test set consists of 14 samples. Each sample consists of 7129 gene probes.

A second data set consists of the same set of genetic probes but an outcome variable based on Lymph Node status is also studied. We consider both data sets for determining efficacy of FIVM algorithm. It is customary in many of the classification papers for the large p , small n domain to screen out only a few genes based on some preliminary analysis, thus reducing dimensionality. In their paper West et al. (2001). described a simple screen to identify the 100 genes maximally correlated with outcome. This choice often made for convenience rather than being confirmatory. We consider an alternative approach through FIVM algorithm by means of which we like to select maximally important subset of genes. The normalized data set is divided into training and testing data set as described earlier. Tuning parameters are chosen via ten fold cross-validation and grid search similar to the synthetic data set over the training set and found to be $\lambda = 0.5$ and $\theta = 1$. For each experiment we explore two different values of the convergence parameter ε namely, 0.05 and 0.001 with $k = 1$.

5.2.1 Classification on the Basis of Estrogen Receptor Status—For $\varepsilon = 0.05$ only four genes are selected while for $\varepsilon = 0.001$ six genes are selected. Selected genes and details

of the results are presented in Table 2. The number outside the bracket indicates mean bootstrap error (over 20 replications) over the test set and the number inside the bracket is standard deviation. Though testing accuracy is not high for FIVM, it consistently outperforms SVM and additionally performs feature selection. We next present four 3D scatter plots in Figure 5 based on selected set of genes. Figure 5(a) represents a scatter plot for training data set based on the first three selected genes which decreases regularized NLL most. Figure 5(b) represents a scatter plot based on second set of three selected genes. Similar plots for test set are drawn in Figure 5(c) and 5(d). Notably out of top six genes selected by FIVM, 4 are found to be common with the West et al. (2001).

5.2.2 Classification on the Basis of Lymph Node Status—We have done similar analysis for lymph node (LN) data set. For two different choices of $\varepsilon = 0.05$ and $\varepsilon = 0.001$ three and five genes are selected respectively. Selected genes and details of the results are presented in Table 3. Notably mean bootstrap testing accuracy for SVM is close to 0.475, implying it is similar to a naive random classifier. We next present four 3D scatter plots in Figure 6 based on selected set of genes. Figure 6(a) represents a scatter plot for training data set based on the first three selected genes which decreases regularized NLL most. Figure 6(b) represents a scatter plot based on second set of three selected genes. Similar plots for test set are drawn in Figure 6(c) and 6(d).

5.3 Exploration With Colon Cancer Data

Alon et al. (1999) described a Gene expression profile for 40 tumor and 22 normal colon tissue samples, analyzed with an Affymetrix oligonucleotide array which is complementary to more than 6,500 human genes and expressed sequence tags (ESTs). Colon adenocarcinoma specimens were collected from patients and mRNA was extracted and hybridized to the array. The Affymetrix Hum6000 array contains about 65,000 features. For details regarding array preparation we would like to refer to original paper of Alon et al. (1999). Final data set contains intensities of 2,000 genes in 22 normal and 40 tumor colon tissues. The genes chosen are the 2,000 genes with highest minimal intensity across the samples. Each gene was also normalized based on all observations. The data set can be downloaded with all complimentary information from the website, <http://microarray.princeton.edu/oncology/affydata/index.html>.

This data set is heavily studied and benchmarked (Getz et al., 2003; Pochet et al., 2004; Li et al., 2004; Weston et al., 2000) in different context by many researchers extensively. FIVM algorithm is used to find its efficacy in gene screening and classification accuracy. We first divide them into training set and testing set by allocating normal and tumor samples randomly in the test and training set respectively. The resulting training set contains 50 observations while the testing data set has 12 observations. Normalized training data set is then presented to FIVM algorithm. Tuning parameter λ is found to be 0.35, with $\theta = 2$ via cross-validation and we choose $k = 1$. For this experiment we choose $\varepsilon = 0.001$ as the value of the convergence parameter. A lower value of ε will of course select more number of genes. Only three genes are selected by the algorithm with mean bootstrap testing error 0.258. The selected genes and details of the result are presented in Table 4. Notably the three selected genes has appeared in many different citations as the top significant genes that can

differentiate between normal and cancer tissue samples (Yang et al., 2003; Yang, 2010). FIVM again beats SVM in testing accuracy. Interestingly this data set is analyzed using similar training (50 out of 62) and testing (12 out of 62) split-up with 15 selected genes in Weston et al. (2000) and they report 12.8% testings error (no variance is reported). While they do not identify which genes are selected however we report similar accuracy with only three genes. Selecting additional genes by lowering the values of ε will of course improve the accuracy of our method. Next we present two 3D scatter plots in Figure 7 based on selected set of genes for both training and testing sets.

6 Multiclass Extension of FIVM

FIVM essentially follows KLR to select a submodel, hence it is readily extendible for multiclass case. We briefly outline the algorithm for multiclass FIVM (MFIVM), which is again adopted for dimension filtering from multiclass IVM described in Zhu and Hastie (2005). Following Section 3, for multiclass case the class label y assumes value from a finite set $C = \{1, \dots, M\}$. Let us denote the conditional probability of a point being classified in category $c \in C$ as $p_c(\mathbf{x}) = P(Y = c | \mathbf{X} = \mathbf{x})$. Then Bayes classification rule is given by

$$c(\mathbf{x}) = \arg \max_{c \in C} p_c(\mathbf{x}) = \arg \max_{c \in C} \frac{e^{f_c(\mathbf{x})}}{\sum_{k=1}^M e^{f_k(\mathbf{x})}}, \quad (6.1)$$

where $f_c(\mathbf{x}) \in \mathcal{H}_k$; $\mathcal{H}_k$ is RKHS generated by the positive definite kernel $K(\cdot, \cdot)$. As noted in Zhu and Hastie (2005) above model is not identifiable for a location shift of $f_c(\mathbf{x})$ and we need to put an identifiability constraint

$$\sum_{c=1}^M f_c(\mathbf{x}) = 0.$$

Following the similar path of Section 4.1 we first introduce multiclass KLR, in which we minimize regularized NLL of the multinomial likelihood represented as

$$H = -\frac{1}{n} \sum_{i=1}^n \ln p_{\mathbf{y}_i}(\mathbf{x}_i) + \lambda \|\mathbf{f}\|_{\mathcal{H}_k}^2 = -\frac{1}{n} \sum_{i=1}^n \left[\mathbf{y}_i \mathbf{f}(\mathbf{x}_i) - \ln \left(e^{f_1(\mathbf{x}_i)} \dots + e^{f_M(\mathbf{x}_i)} \right) \right] + \lambda \|\mathbf{f}\|_{\mathcal{H}_k}^2, \quad (6.2)$$

where $\mathbf{y}_i$ is a M dimensional vector with all zero elements except a 1 in position c indicating which class the i -th observation belongs to and

$$\mathbf{f}(\mathbf{x}_i) = (f_1(\mathbf{x}_i), \dots, f_M(\mathbf{x}_i))', \quad (6.3)$$

$$\|\mathbf{f}\|_{\mathcal{H}_k}^2 = \sum_{c=1}^M \|f_c\|_{\mathcal{H}_k}^2. \quad (6.4)$$

Again using representer theorem of Kimeldorf and Wahba (1971), it can be shown that $f_c(\mathbf{x})$ minimizes (6.2) and is given by

$$f_c(\mathbf{x}) = \sum_{i=1}^n a_{ic} K(\mathbf{x}_i, \mathbf{x}). \quad (6.5)$$

Following Zhu and Hastie (2005) with little abuse of notations, equation (6.2) can be rewritten in matrix notations as

$$H = \frac{1}{n} \sum_{i=1}^n \left[-\mathbf{y}'_i (\mathbf{K}(i, \cdot) \mathbf{A})' + \log \left(\mathbf{1}' e^{(\mathbf{K}(i, \cdot) \mathbf{A})'} \right) \right] + \lambda \sum_{c=1}^M a_c' \mathbf{K} a_c, \quad (6.6)$$

where $\mathbf{A} = (a_1, \dots, a_M)$ and $\mathbf{K}$ is defined as earlier in the two class problem. With the above setup we would like to apply FIVM algorithm described in Section 4.2 for selecting out important features. Computational burden of MFIVM is similar to two class FIVM and is proportional with the number of class labels. In Figure 8, we presents a simulation result for multiclass FIVM with three classes and with two informative dimensions. The data in each class are generated from a mixture of Gaussians (Hastie et al., 2001).

7 Summary and Conclusions

Classification literature is pretty rich both from probabilistic and deterministic domain. In this paper we have developed a classification method based on reproducing kernel Hilbert space principle. We first select useful features which decrease the regularized negative log-likelihood most. The final classifier is then constructed on the transformed feature space but by using only the selected features. This lead to great data reduction as well as non-linear classification. We have tested efficacy of our algorithm for synthetic as well as real life data examples.

One future problem for classification in high dimension will be when dimensions are having non ignorable association among each other. Notably this cross-correlation information is not being utilized in the dimension selection algorithm developed in this paper. Another exciting future possibility is simultaneous selection of feature and observation especially in the *Big data* context. With appropriate regularization scheme this may leads to huge data compression with minimal loss in classification accuracy.

Acknowledgments

The research of first author is partly supported by NIH grant number P30 ES020957. Authors would like to thank Dr. Dipak Dey for his comments. Authors would also like to thank two anonymous referees and the editorial team, whose comments provided additional insights and have greatly improved the scope and presentation of the paper.

Appendix

A Proof of Theorem 4.1

Since n training observations are completely separable in $\mathcal{S}$ we have

$$y_i f_{\mathcal{S}}(\mathbf{x}_i) = y_i \sum_{j=1}^n \alpha_j K_{\mathcal{S}}(\mathbf{x}_i, \mathbf{x}_j) > 0, \text{ for } i=1, 2, \dots, n.$$

Proof: For RBF it is easy to show $K_{\mathcal{G}}(\mathbf{x}_i, \mathbf{x}_j) = K_{\mathcal{S}}(\mathbf{x}_i, \mathbf{x}_j)$, $K_{\mathcal{G} \setminus \mathcal{S}}(\mathbf{x}_i, \mathbf{x}_j)$, where $\mathcal{S} \subseteq \mathcal{L}$ and RBF is a bounded kernel having the property $0 \leq K_{\Lambda}(\mathbf{x}_i, \mathbf{x}_j) \leq 1$ for any set Λ and

$$K_{\Lambda}(\mathbf{x}_i, \mathbf{x}_j | \theta) = \exp \left\{ -\frac{\sum_{t \in \Lambda} (x_{it} - x_{jt})^2}{2\theta^2} \right\}. \text{ Similarly for the set } \mathcal{L}$$

$$y_i f_{\mathcal{L}}(\mathbf{x}_i) = y_i \sum_{j=1}^n \beta_j K_{\mathcal{L}}(\mathbf{x}_i, \mathbf{x}_j) = y_i \sum_{j=1}^n \beta_j K_{\mathcal{S}}(\mathbf{x}_i, \mathbf{x}_j) K_{\mathcal{L} \setminus \mathcal{S}}(\mathbf{x}_i, \mathbf{x}_j). \quad (\text{A.1})$$

Since $\mathcal{S} \subseteq \mathcal{L}$, if we take all dimensions in $\mathcal{L} \setminus \mathcal{S}$ as constant (i.e., providing no information)

for all observations, then $\sum_{t \in \mathcal{L} \setminus \mathcal{S}} (x_{it} - x_{jt})^2 = 0$. This implies $K_{\mathcal{L} \setminus \mathcal{S}}(\mathbf{x}_i, \mathbf{x}_j) = 1$. Now if we take $\beta_j = \alpha_j$ for $j = 1, 2, \dots, n$, then $y_i f_{\mathcal{L}}(\mathbf{x}_i) > 0$ for $i = 1, 2, \dots, n$ and hence separable in $\mathcal{L}$

too. Notably we are not saying in real life $\sum_{t \in \mathcal{L} \setminus \mathcal{S}} (x_{it} - x_{jt})^2 = 0$ need to hold, in fact it will not be true in general. However since we already have complete separability in $\mathcal{S}$,

suppressing all $\mathcal{L} \setminus \mathcal{S}$ dimensions by some constant will render $\sum_{t \in \mathcal{L} \setminus \mathcal{S}} (x_{it} - x_{jt})^2 = 0$, hence the lower dimensional classifier in $\mathcal{S}$ will still achieve a complete separation in $\mathcal{L}$.

B Proof of Theorem 4.2

We begin with an important theorem of Zhu and Hastie (2005) which relates to the solution of KLR and SVM in an unified framework of margin maximizing loss function. We omit the subscript for denoting dimensional set, as in their original theorem dimension is not assumed to be a varying quantity, rather solutions for equivalent SVM and KLR problem considered only.

Theorem B.1. If training data is separable, that is $\exists \beta_0, \beta$, s.t. $y_i(\beta_0 + \sum_{j=1}^n \beta_j K(\mathbf{x}_i, \mathbf{x}_j)) > 0$ $\forall_i$ and if the solution for the KLR problem stated earlier is denoted as $\tilde{\beta}(s)$, then

$$\frac{\tilde{\beta}(s)}{s} \rightarrow \beta^* \text{ as } s \rightarrow \infty,$$

where β^* is the solution of the margin maximizing SVM, provided β^* is unique. If β^* is not

unique, then $\frac{\tilde{\beta}(s)}{s}$ may have multiple convergence points, but they will all represent margin-maximizing separating hyperplanes.

Regarding the proof of the theorem we refer to Rosset et al. (2004) and Zhu and Hastie (2005). This theorem tells us that margin maximizing property is shared by KLR and SVM

and their asymptotic solution is consistent. Again as stated earlier in Section 4, for any $\mathcal{S}$ having only first $q(p)$ dimensions of $\mathcal{L}$, the KLR problem in $\mathcal{S}$ is given in equation 4.3. Denote its solution by $\hat{\beta}^*(s)$. Then by Theorem B.1, as $s \rightarrow \infty$,

$$\frac{\hat{\beta}^*(s)}{s} \rightarrow \hat{\beta}^*.$$

Again for the complete space $\mathcal{H}$ having p dimensions, by Theorem B.1, we have as $s \rightarrow \infty$,

$$\frac{\tilde{\beta}(s)}{s} \rightarrow \beta^*.$$

Let $\tilde{\beta}^*(s)$ be the first q coordinates of $\tilde{\beta}(s)$. Then as $s \rightarrow \infty$,

$$\frac{\tilde{\beta}^*(s)}{s} \rightarrow \tilde{\beta}^*.$$

We will give the proof of Theorem 4.2 through two propositions.

B.1 Proof of Proposition 4.1

Proof: It is a consequence of Theorem 4.1 and the definitions of $Q(p, \mathbf{x}, \mathcal{L})$ and $Q(q, \mathbf{x}^*, \mathcal{S})$

B.2 Proof of Proposition 4.2

Proof: Suppose training data $\mathbf{x} = (\mathbf{x}^*, \mathbf{x}^\dagger)$ are separable in $\mathcal{L}$. Since $\mathcal{S} \subset \mathcal{L}$. Denote

$$\mathcal{L} = \mathcal{S} + \mathcal{S}^\perp$$

We show that a hyperplane in $\mathcal{L}$ separating $\mathbf{x} = (\mathbf{x}^*, \mathbf{x}^\dagger)$ for all possible noise $\mathbf{x}^\dagger$ must perpendicularly pass through a hyperplane in $\mathcal{S}$ that separates $\mathbf{x}^*$.

Because the noise $\mathbf{x}^\dagger$ is random and may take all possible values, it is not possible to separate $\mathbf{x} = (\mathbf{x}^*, \mathbf{x}^\dagger)$ using only $\mathbf{x}^\dagger$, and any hyperplane ℓ in $\mathcal{L}$ separating $\mathbf{x} = (\mathbf{x}^*, \mathbf{x}^\dagger)$ for all possible $\mathbf{x}^\dagger$ must also separate $\mathbf{x}^*$ in $\mathcal{S}$. That is, any hyperplane in $\mathcal{L}$ separating $\mathbf{x} = (\mathbf{x}^*, \mathbf{x}^\dagger)$ for all possible $\mathbf{x}^\dagger$ must pass a separating hyperplane in $\mathcal{S}$.

We next show that if the hyperplane separates $\mathbf{x} = (\mathbf{x}^*, \mathbf{x}^\dagger)$ for all possible noise $\mathbf{x}^\dagger$, it has to be perpendicular to $\mathcal{S}$. Otherwise, suppose a hyperplane ℓ in $\mathcal{L}$ is not perpendicular to $\mathcal{H}$. Then ℓ is tilted to some coordinates of $\mathbf{x}^\dagger$, say these coordinates are x_r , $r > q$. We may select large positive values for these $|x_r|$, such that $\mathbf{x} = (\mathbf{x}^*, \mathbf{x}^\dagger)$ and $\mathbf{x} = (\mathbf{x}^*, -\mathbf{x}^\dagger)$ fall in two sides of the hyperplane ℓ . Thus, the hyperplane ℓ cannot separate $\mathbf{x} = (\mathbf{x}^*, \mathbf{x}^\dagger)$ for both $\mathbf{x}^\dagger$ and $-\mathbf{x}^\dagger$. Therefore, the hyperplane in $\mathcal{L}$ separating $\mathbf{x} = (\mathbf{x}^*, \mathbf{x}^\dagger)$ for all possible $\mathbf{x}^\dagger$ with the largest

distance to $\mathbf{x} = (\mathbf{x}^*, \mathbf{x}^\dagger)$ is the same hyperplane in $\mathcal{S}$ separating $\mathbf{x}^*$ with the largest distance to $\mathbf{x}^*$, that is $Q(\mathbf{x}^*, \mathcal{L}) = Q(q, \mathbf{x}^*, \mathcal{S})$. This implies $\hat{\beta}^* = \tilde{\beta}^*$ and hence the proposition is proved.

Now by Theorem B.1 and above proposition we obtain

$$\hat{\beta}^*(s)/s - \tilde{\beta}^*(s)/s \rightarrow \hat{\beta}^* - \tilde{\beta}^* = 0.$$

This completes the proof of Theorem 4.2.

References

- Alon, U.; Barkai, N.; Notterman, DA.; Gish, K.; Ybarra, S.; Mack, D.; Levine, AJ. Proceedings of the National Academy of Sciences. Vol. 96. National Academy of Sciences; 1999. Broad patterns of gene expression revealed by clustering analysis of tumor and normal colon tissues probed by oligonucleotide arrays.
- Aronszajn N. Theory of Reproducing Kernels. Transactions of the American Mathematical Society. 1950; 68:337–404.
- Canu S, Grandvalet Y. Adaptive Scaling for Feature Selection in SVMs. NIPS. 2002; 15
- Dennis, JE.; Schnabel, RB. Numerical Methods for Unconstrained Optimization and Nonlinear Equations. Englewood Cliffs, NJ: Prentice Hall; 1983.
- Evgeniou, T.; Pontil, M.; Poggio, T. AI Memo 1654. Massachusetts Institute of Technology; 1999. A unified framework for Regularization Networks and Support Vector Machines.
- Getz G, Gal H, Kela I, N DA, D E. Coupled Two-way Clustering Analysis of Breast Cancer and Colon Cancer Gene Expression Data. Bioinformatics. 2003; 19:1079–1089. [PubMed: 12801868]
- Hastie, T.; Tibshirani, R.; Friedman, J. The Elements of Statistical Learning. New York: Springer Series in Statistics; 2001.
- Kimeldorf G, Wahba G. Some Results on Tchebyshevan Spline Functions. Journal of Mathematical Analysis and Application. 1971; 33:82–95.
- Lee Y, Lin Y, Wahba G. Multicategory Support Vector Machines, Theory, and Application to the Classification and Microarray Data and Satellite Radiance Data. Journal of American Statistical Association. 2004; 99:67–81.
- Li X, Rao S, Wang Y, Gong B. Gene mining: a novel and powerful ensemble decision approach to hunting for disease genes using microarray expression profiling. Nucleic Acids Res. 2004; 32:2685–2694. [PubMed: 15148356]
- Lin Y. Support Vector Machines and the Bayes Rule in Classification. Data Mining and Knowledge Discovery. 2002; 6:259–275.
- Lin Y, Wahba G, Xiang D, Gao F, Klein B. Smoothing Spline ANOVA Models for Large Data Sets with Bernoulli Observations and the Randomized GACV. The Annals of Statistics. 2000; 28:1570–1600.
- Nguyen M, Torre F. Optimal feature selection for support vector machines. Pattern Recognition. 2010; 43:584–591.
- Numerics, V. IMSL C Library Users's Manual. Vol. 2. IMSL; Houston, Texas: 1987. version 5.5
- Park M, Hastie T. Penalized logistic regression for detecting gene interactions. Bio-statistics. 2008; 9:30–50.
- Parzen, E. Proceedings of the 12th Biennial Seminar. Montreal: Canadian Mathematical Congress; 1970. Statistical Inference on Time Series by the RKHS Methods; p. 1-37.
- Pochet N, Smet FD, Suykens JAK, De Moor BLR. Systematic benchmarking of microarray data classification: assessing the role of non-linearity and dimensionality reduction. Bioinformatics. 2004; 20:3185–3195. [PubMed: 15231531]
- Rosset S, Zhu J, Hastie T. Margin maximizing loss functions. Neural Information Processing Systems. 2004; 16

- Roth V. The generalized LASSO. IEEE Transactions on Neural Networks. 2004; 15:16–28. [PubMed: 15387244]
- Smola, A.; Scholkopf, B. Proceedings of the Seventeenth International Conference on Machine Learning. San Francisco, CA: Morgan Kaufmann; 2000. Sparse Greedy Matrix Approximation for Machine Learning; p. 911-918.
- Tibshirani R. Regression shrinkage and selection via the lasso. Journal of the Royal Statistical Society, B. 1996; 58:267–288.
- Vapnik, VN. The Nature of Statistical Learning Theory. Berlin: Springer Verlag; 1995.
- Vapnik, VN. Statistical Learning Theory. New York: Wiley; 1998.
- Wahba, G. Spline Models for Observational Data. Philadelphia: SIAM; 1990.
- West, M.; Blanchette, C.; Dressman, H.; Huang, E.; Ishida, S.; Spang, R.; Zuzan, H.; Olson, A.J.; Marks, J.R.; Nevins, J.R. Predicting the clinical status of human breast cancer by using gene expression profiles. In: Bickel, P.J., editor. Proceedings of the National Academy of Sciences. Vol. 98. National Academy of Sciences; 2001. p. 11462-11467.
- Weston J, Mukherjee A, Chapelle O, Pontil M, Poggio T, Vapnik V. Feature selection for SVMs. Advances in Neural Information Processing Systems. 2000; 13
- Weston, W.; Watkins, C. Tech Rep CSD-TR-98-04. Univ. of London; Royal Holloway: 1999. Multi-class Support Vector Machines.
- Yang, S. Comparison of RankGene Measures for Colon Cancer Data of Alon et al. 2010. available at <http://genomics10.bu.edu/yangsu/rankgene/compare-alon-colon-cancer-top100.html>
- Yang S, Murali T, Pavlovic V, S M, K S. RankGene: identification of diagnostic genes based on expression data. Bioinformatics. 2003; 19:1578–1579. [PubMed: 12912841]
- Zhu J, Hastie T. Kernel logistic regression and the import vector machine. Journal of Computational and Graphical Statistics. 2005; 14:185–205.

Figure 1. Hinge loss of SVM and NLL of binomial distribution for two class classification, $y \in \{-1, 1\}$.

Figure 2.

Scatter plots of the synthetic data over different dimensions. Plot (a) is between first two dimensions. Plot (b) and plot(c) include noisy dimensions. Clearly last two scatter plots do not represent any classification pattern, as expected. Plot (d) represents data points after classification. Only dimensions selected (first and second one) by the FIVM algorithm is used for final classification. Misclassification rate in training set is 0.105.

Figure 3. Radial basis kernel is used throughout, $n = 200$, $\theta = 1$, $k = 1$. The stopping criterion is satisfied when $|\mathcal{S}| = 2$. For plot (c), ε value is adjusted so that desired number of dimensions could be chosen.

Figure 4. Plot (a) shows scatter plots of the synthetic testing data over selected dimensions. Plot (b) represents data points after FIVM classification. Plot (c) shows misclassification rate as a function of sample size.

3D Plot based on First set of Three selected genes

(a)

3D Plot based on Second set of Three selected genes

(b)

3D Plot based on First set of Three selected genes

(c)

3D Plot based on Second set of Three selected genes

(d)

Figure 5.

Scatter plots of the Breast Cancer Data set (ER Status). Plot (a) and (c) are drawn among first three selected genes which decreases regularized NLL the most. Plot (b) and (d) are drawn among second set of three selected genes. Each gene is marked by the accession number in each axis. Top two plots (a and b) are for training set and bottom two plots (c and d) are for test set.

3D Plot based on First set of Three selected genes

(a)

3D Plot based on Second set of Three selected genes

(b)

3D Plot based on First set of Three selected genes

(c)

3D Plot based on Second set of Three selected genes

(d)

Figure 6.

Scatter plots of the Breast Cancer Data set (LN Status). Plot (a) and (c) are drawn among first three selected genes which decreases regularized NLL the most. Plot (b) and (d) are drawn among second set of three selected genes. Each gene is marked by the Accession number in each axis. Top two plots (a and b) are for training set and bottom two plots (c and d) are for test set.

Figure 7.

Scatter plots of the Colon Cancer Data set. Plots are drawn on the basis of three selected genes. Each gene is marked by the Gene number in each axis. Plot (a) is for training set and Plot (b) is for testing set.

Figure 8. Radial kernel is used. $C = 3$, $n = 300$, $\lambda = 1.5$, Training error=0.089, Testing Error=0.1.

Table 1

Testing performance of SVM, FIVM and L_1 Logistic Regression. SVM will not do any dimension selection, so dimensions are selected first by FIVM. L_1 logistic regression is also able to select two informative features for appropriate λ . For all cases tuning parameters are selected through extensive grid search and via ten fold cross-validation over the training set. We searched over the range $[2^{-6}, 2^6]$ for θ and $[2^{-10}, 2^{10}]$ for λ .

SVM Performance	FIVM Performance	L_1 Logistic Regression
$\theta = 2, \lambda = 4$, Train error=0.125	$\theta = 1, \lambda = 0.5$, Train error=0.105	$\lambda = 0.1$, Train error=0.12
Test error= 0.18, 36 Support points	Test errors 0.125	Test error= 0.145
Run Time=55 sec	Run Time=51 sec	Run Time=31 sec

Table 2
Estrogen Receptor (ER) Classification Result for Breast Cancer Data

$\varepsilon = 0.05$	$\varepsilon = 0.001$
Four Genes Selected by FIVM	Six Genes Selected by FIVM
Z84721_cds2, Contains alpha and zeta globin genes and ESTs	Z84721_cds2, Contains alpha and zeta globin genes and ESTs
X83425, Homo Sapiens LU gene for Lutheran blood group glycoprotein	X83425, Homo Sapiens LU gene for Lutheran blood group glycoprotein
X55037, Homo Sapiens GATA-3 mRNA	X55037, Homo Sapiens GATA-3 mRNA
U33147, Human mammaglobin mRNA, complete cds	U33147, Human mammaglobin mRNA, complete cds
	HG1205-HT1205, Collagen, Type IV, Alpha 2, N-Terminus
	X03635, Human mRNA for oestrogen receptor
Testing error= 0.344(± 0.032)	Testing error= 0.416(± 0.027)
SVM Testing Performance & Tuning Parameters	
$\theta = 0.5, \lambda = 1$ Error=0.406(± 0.052), 12(± 4) support points	$\theta = 1, \lambda = 1$ Error=0.41(± 0.054), 20(± 3) support points

Table 3
Lymph Node (LN) Classification Result for Breast Cancer Data

$\varepsilon = 0.05$	$\varepsilon = 0.001$
Three Genes Selected by FIVM	Five Genes Selected by FIVM
HG3638-HT3849_s_at Amyloid Beta (A4) Precursor Protein, Alt. Splice 2,	HG3638-HT3849_s_at Amyloid Beta (A4) Precursor Protein, Alt. Splice 2,
U02493, Human 54 kDa protein mRNA, complete cds	U02493, Human 54 kDa protein mRNA, complete cds
U51678, Human small acidic protein mRNA, complete cds	U51678, Human small acidic protein mRNA, complete cds
	HG3432-HT3620_s_at Fibroblast Growth Receptor K-Sam, Alt. Splice 3
	X80692, Homo Sapiens ERK3 mRNA
Testing error= 0.304(± 0.017)	Testing error= 0.411(± 0.022)
SVM Testing Performance & Tuning Parameters	
$\theta = 0.25, \lambda = 2$ Error=0.475(± 0.012), 25(± 4) support points	$\theta = 0.5, \lambda = 4$ Error=0.475(± 0.012), 25(± 4) support points

Table 4
Classification Result for Colon Cancer Data set

$\varepsilon = 0.001$
Three Genes Selected by FIVM
T71025: Human (HUMAN)
M76378: Human cysteine-rich protein (CRP) gene, exons 5 and 6.
M63391: Human desmin gene, complete cds.
Testing error= 0.153(± 0.033)
SVM Testing Performance & Tuning Parameters
$\theta = 0.25$ & $\lambda = 8$ Testing error=0.22(± 0.047), 10(± 3) support points