Flexible High-dimensional Classification Machines and Their Asymptotic Properties

Xingye Qiao*

Department of Mathematical Sciences

State University of New York, Binghamton, NY 13902-6000.

E-mail: `qiao@math.binghamton.edu`

Lingsong Zhang

Department of Statistics

Purdue University, West Lafayette, IN 47907.

E-mail: `lingsong@purdue.edu`

---

*Corresponding author

I

Abstract

Classification is an important topic in statistics and machine learning with great potential in many real applications. In this paper, we investigate two popular large margin classification methods, Support Vector Machine (SVM) and Distance Weighted Discrimination (DWD), under two contexts: the high-dimensional, low-sample size data and the imbalanced data. A unified family of classification machines, the FLexible Assortment MachinE (FLAME) is proposed, within which DWD and SVM are special cases. The FLAME family helps to identify the similarities and differences between SVM and DWD. It is well known that many classifiers overfit the data in the high-dimensional setting; and others are sensitive to the imbalanced data, that is, the class with a larger sample size overly influences the classifier and pushes the decision boundary towards the minority class. SVM is resistant to the imbalanced data issue, but it overfits high-dimensional data sets by showing the undesired data-piling phenomena. The DWD method was proposed to improve SVM in the high-dimensional setting, but its decision boundary is sensitive to the imbalanced ratio of sample sizes. Our FLAME family helps to understand an intrinsic connection between SVM and DWD, and improves both methods by providing a better trade-off between sensitivity to the imbalanced data and overfitting the high-dimensional data. Several asymptotic properties of the FLAME classifiers are studied. Simulations and real data applications are investigated to illustrate the usefulness of the FLAME classifiers.

Key Words and Phrases: Classification; Discriminant analysis; Fisher consistency; High-dimensional, low-sample size asymptotics; Imbalanced data; Support Vector Machine.

1 Introduction

Classification refers to predicting the class label, $y \in \mathcal{C}$, of a data object based on its covariates, $\boldsymbol{x} \in \mathcal{X}$. Here $\mathcal{C}$ is the space of class labels, and $\mathcal{X}$ is the space of the covariates. Usually we consider $\mathcal{X} \equiv \mathbb{R}^d$, where d is the number of variables or the dimension. See Duda et al. (2001) and Hastie et al. (2009) for comprehensive introductions to many popular classification methods. When $\mathcal{C} = \{+1, -1\}$, this is an important class of classification problems, called binary classification. The classification rule for a binary classifier usually has the form $\phi(\boldsymbol{x}) = \text{sign}\{f(\boldsymbol{x})\}$, where $f(\boldsymbol{x})$ is called the discriminant function. Linear classifiers are the most important and the most commonly used classifiers, as they are often easy to interpret in addition to reasonable classification performance. We focus on linear classifier in this article. In the above formula, linear classifiers correspond to $f(\boldsymbol{x}; \boldsymbol{\omega}, \beta) = \boldsymbol{x}^T\boldsymbol{\omega} + \beta$. The sample space is divided into halves by the *separating hyperplane*, also known as the *classification boundary*, defined by $\left\{\boldsymbol{x} : f(\boldsymbol{x}) \equiv \boldsymbol{x}^T\boldsymbol{\omega} + \beta = 0\right\}$. Note that the coefficient vector $\boldsymbol{\omega} \in \mathbb{R}^d$ defines the normal vector, and hence the direction, of the classification boundary, and the intercept term $\beta \in \mathbb{R}$ defines the location of the classification boundary.

In this paper, two popular classification methods, Support Vector Machine (SVM; Cortes and Vapnik, 1995, Vapnik, 1998, Cristianini and Shawe-Taylor, 2000) and Distance Weighted Discrimination (DWD; Marron et al., 2007, Qiao et al., 2010) are investigated under two important contexts: the High-Dimensional, Low-Sample Size (HDLSS) data and the imbalanced data. Both methods are large margin classifiers (Smola et al., 2000), that seek separating hyperplanes which maximize certain notions of *gap* (*i.e.*, distances) between the two classes. The investigation of the performance of SVM and DWD motivates the invention of a novel family of classifiers, the FLexible Assortment MachinE (FLAME), which unifies the two classifiers, and helps to understand their connections and differences.

1.1 Motivation: Pros and Cons of SVM and DWD

SVM is a very popular classifier in statistics and machine learning. It has been shown to have Fisher consistency, *i.e.*, when sample size goes to infinity, its decision rule converges to the Bayes rule (Lin, 2004). SVM has several nice properties. 1) Its dual formulation is relatively easy to implement (by the Quadratic Programming). 2) SVM is robust to the model specification, which makes it very popular in various real applications. However, when being applied to HDLSS data, it has been observed that a large portion of the data (usually the support vectors, to be properly defined later) lie on two hyperplanes parallel to the SVM classification boundary. This is known as the *data-piling* phenomena (Marron et al., 2007, Ahn and Marron, 2010). Data-piling of SVM indicates a type of overfitting. Other overfitting phenomena of SVM under the HDLSS context include:

1. The angle between the SVM direction and the Bayes rule direction is usually large.

2. The variability of the sampling distribution of the SVM direction $\boldsymbol{\omega}$ is very large (Zhang and Lin, 2011). Moreover, because the separating hyperplane is decided only by the support vectors, the SVM direction tends to be unstable, in the sense that small turbulence or measurement error to the support vectors can lead to big change of the direction.

3. In some cases, the out-of-sample classification performance may not be optimal due to the suboptimal direction of the estimated SVM discrimination direction.

DWD is a recently developed classifier to improve SVM in the HDLSS setting. It uses a different notion of gap from SVM. While SVM is to maximize the smallest distance between classes, DWD is to maximize a special average distance (harmonic mean) between classes. It has been shown in many earlier simulations that DWD largely overcomes the overfitting (data-piling) issue and it usually gives a better discrimination direction.

On the other hand, the intercept term β of the DWD method is sensitive to the sample size ratio between the two classes, *i.e.*, to the imbalanced data (Qiao et al., 2010). Note that, even though a good discriminant direction $\boldsymbol{\omega}$ is more important in revealing the profiling

difference between the two populations, the classification/prediction performance heavily depends on the intercept β, more than on the direction $\boldsymbol{\omega}$. As shown in Qiao et al. (2010), usually the β of the SVM classifier is not sensitive to the sample size ratio, while the β of the DWD method will become too large (or too small) if the sample size of the positive class (or negative class) is very large.

In summary, both methods have pros and cons. SVM has larger stochastic variability and usually overfits the data by showing data-piling phenomena, but is less sensitive to the imbalanced data issue. DWD usually overcomes the overfitting/data-piling issue, and has smaller sampling variability, but is very sensitive to the imbalanced data. Driven by their similarity, we propose a unified class of classifiers, FLAME, in which the above two classifiers are special cases. FLAME provides a framework to study the connections and differences between SVM and DWD. Each FLAME classifier has a parameter θ which is used to control the performance balance between overfitting the HDLSS data and the sensitivity to the imbalanced data. It turns out that the DWD method is FLAME with $\theta = 0$; and that the SVM method corresponds to FLAME with $\theta = 1$. The optimal θ depends on the trade-off among several factors: stochastic variability, overfitting and resistance against the imbalanced data. In this paper, we also propose two approaches to select θ, where the resulting FLAME have a balanced performance between the SVM and DWD methods.

1.2   Outlines

The rest of the paper is organized as follows. Section 2 provides toy examples and highlights the strengths and drawbacks of SVM and DWD on classifying the HDLSS and imbalanced data. We develop the FLAME method in Section 3, which is motivated by the investigation of the loss functions of SVM and DWD. Section 4 provides suggestions for the parameters. Three types of asymptotic results for the FLAME classifier are studied in Section 5. Sections 6 discusses its properties using simulation experiments. A real application is discussed in Section 7. Some concluding remarks and discussions are made in Section 8.

2 Comparison of SVM and DWD

In this section, we use several toy examples to illustrate the strengths and drawbacks of SVM and DWD under two contexts: HDLSS data and imbalanced data.

2.1 Overfitting HDLSS data

We use simulations to compare SVM and DWD. The results show that the stochastic variability of the SVM direction is usually larger than that of the DWD method, and SVM directions are deviated farther away from Bayes rule directions. In addition, the new proposed FLAME machine (see details in Section 3) is also included in the comparison, and it turns out that FLAME is between the other two.

Figure 1 shows the comparison results between SVM, DWD and FLAME (with some chosen tuning parameters). We simulate 10 samples with the same underlying distribution. Each simulated data set contains 12 variables and two classes, with 120 observations in each class. The two classes have mean difference on only the first three dimensions and the within-class covariances are diagonal, that is, the variables are independent. For each simulated data set, we plot the first three components of the resulting discriminant directions from SVM, DWD and FLAME (after normalizing the 3D vectors to have unit norms), as shown in Figure 1. It clearly shows that the DWD directions (the blue down-pointing triangles) are the closest ones to the *true* Bayes rule direction (shown as the cyan diamond marker) among the three approaches. In addition, the DWD directions have a smaller variation (*i.e.*, more stable) over different samples. The SVM directions (the red up-pointing triangles) are farthest from the *true* Bayes rule direction and have a larger variation than the other two methods. To highlight the direction variabilities of the three methods, we introduce a novel measure for the variation (unstableness) of the discriminant directions: the trace of the sample covariance of the resulting direction vectors over the 10 replications, which we name as *dispersion*. The dispersion for the DWD method (0.0031) is much smaller than
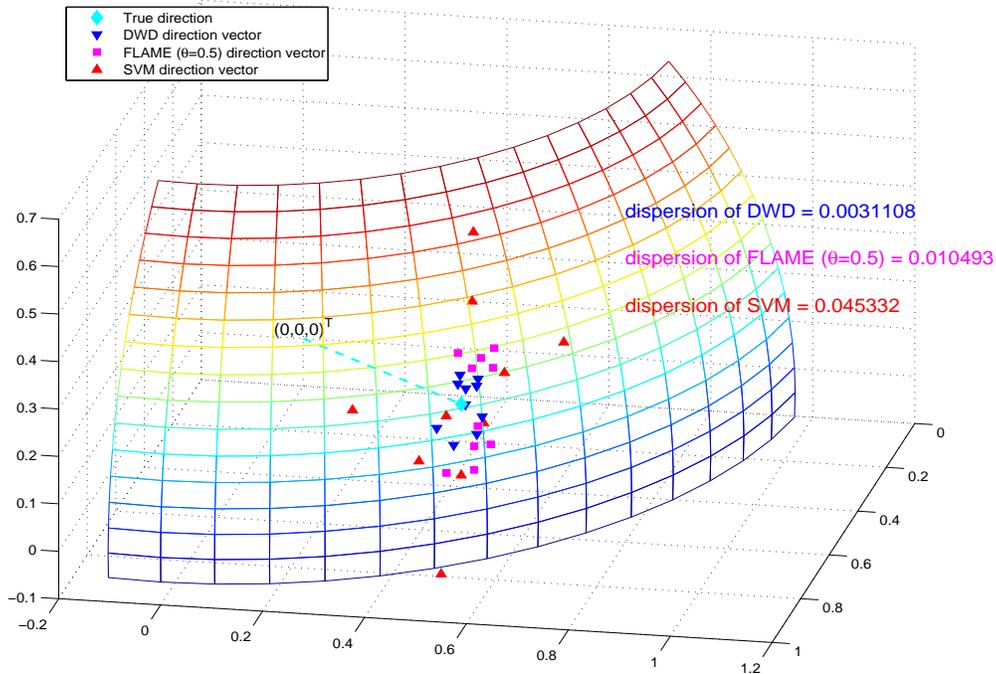
Figure 1: The true population mean difference direction vector (the cyan dashed line and diamond marker; equivalent to the Bayes rule direction), the DWD directions (blue down-pointing triangles), the FLAME directions with $\theta = 0.5$ (magenta squares), and the SVM directions (red up-pointing triangles) for 10 realizations of simulated data. Each direction vector has norm 1 and thus is depicted as a point on the 3D unit sphere. On average, all machines have their discriminant direction vectors scattering around the true direction. The DWD directions are the closest to the true direction and have a smaller variation. The SVM directions have the largest variation and are farthest from the true direction. The variation of the intermediate FLAME direction vectors is between the two machines above. The variation of a machine is also measured by the trace of the sample covariance calculated from the 10 resulting direction vectors for the 10 simulations.

that of the SVM method (0.0453), as highlighted in the figure as well. The new FLAME classifiers usually have the performance between DWD and SVM. Figure 1 shows the results of a specific FLAME ($\theta = 0.5$, the magenta squares), which is better than SVM but worse than DWD.

Besides the stochastic variability and the deviation from the true direction comparisons shown above, DWD outperforms SVM in terms of stability in the presence of small perturbations applied to some observations. In Figure 2, we use a two-dimensional example to illustrate this phenomenon. We simulate a perfectly separable 2-dimensional data set. The
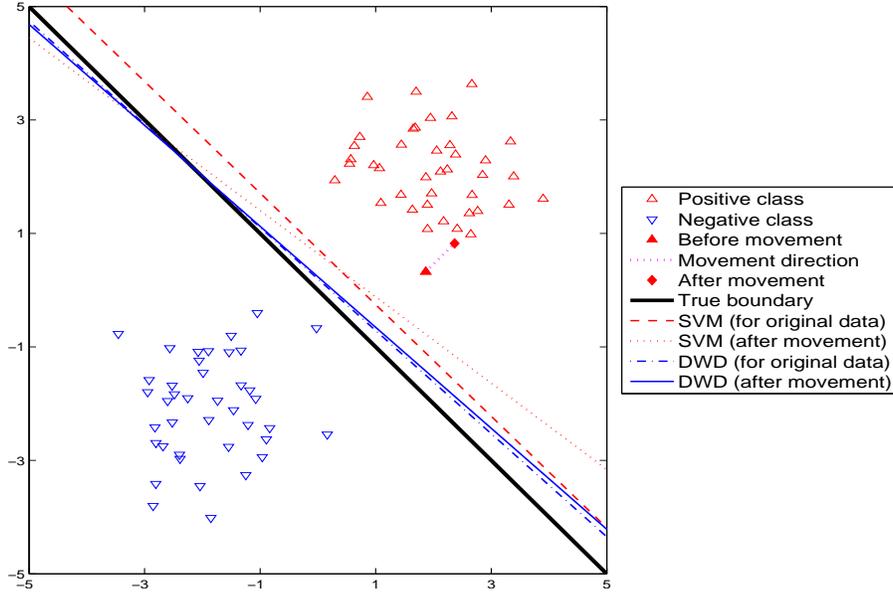
Figure 2: A 2D example shows that the unstable SVM boundary has changed due to a small turbulence of a support vector (the solid red triangle and diamond) while the DWD boundary remains almost still.

theoretical Bayes rule decision boundary is shown as the thick black line. The dashed red line and the dashed dotted blue line are the SVM and the DWD classification boundaries before the perturbation. We then move one observation in the positive group a little (from the solid triangle to the solid diamond as shown in the figure). This perturbation leads to a large change of direction in SVM (shown as the dotted red line), but a small change for DWD (shown as the solid blue line). Note that all four hyperplanes are capable of classifying this training data set perfectly. But it may not be true for an out-of-sample test set. This example shows small perturbation may lead to unstableness in SVM.

2.2   Sensitivity to imbalanced data

In the last subsection, we have shown that DWD outperforms SVM in estimating the discrimination direction, that is, DWD directions are closer to the Bayes rule discrimination directions and have smaller variability. However, it was found that the location of DWD classification boundary, which is characterized by the intercept β, is sensitive to the sample
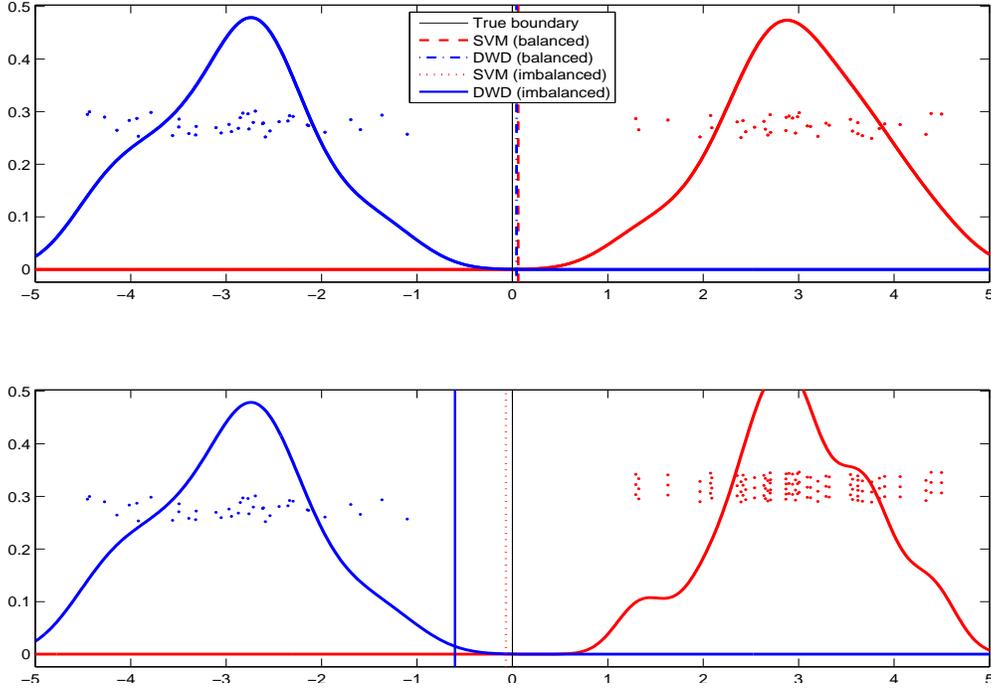
Figure 3: A 1D example shows that the DWD boundary is pushed towards the minority class (blue) when the majority class (red) has tripled its sample size.

size ratio between the two classes (Qiao et al., 2010).

Usually, a good discriminant direction $\boldsymbol{\omega}$ helps to reveal the profiling difference between two classes of populations. But the classification/prediction performance heavily depends on the location coefficient β. We define the *imbalance factor $m \geq 1$* as the sample size ratio between the majority class and the minority class. It turns out that β in the SVM classifier is not sensitive to m. However, the β for the DWD method is very sensitive to m. We also notice that, as a consequence, the DWD separating hyperplane will be pushed toward the minority class, when the ratio m is close to infinity, *i.e.*, DWD classifiers intend to ignore the minority class. Again, we use a toy example in order to better illustrate the impact of the imbalanced data on β and on the classification performance.

Figure 3 uses a one-dimensional example, so that estimating $\boldsymbol{\omega}$ is not needed. This also corresponds to a multivariate data set, where $\boldsymbol{\omega}$ is estimated correctly first, after which the data set is projected to $\boldsymbol{\omega}$ to form the one-dimensional data. In this plot, the x-coordinates of

the red dots and the blue dots are the values of the data while the y-coordinates are random jitters for better visualization. The red and blue curves are the kernel density estimations for both classes. In the top subplot of Figure 3, where $m = 1$ (*i.e.*, the balanced data), both the DWD (blue lines) and SVM (red lines) boundaries are close to the Bayes rule boundary (black solid line), which sits at 0. In the bottom subplot, the sample size of the red class is tripled, which corresponds to $m = 3$. Note that the SVM boundary moves a little towards the minority (blue) class, but still fairly close to the true boundary. The DWD boundary, however, is pushed towards the minority. Although this does not impose immediate problems for the training data set, the DWD classifier will suffer from a great loss of classification performance when it is applied to an out-of-sample data set. It can be shown that when m goes to infinity, the DWD classification boundary will tends to negative infinity, which totally ignores the minority group (see our Theorem 3). However, SVM will not suffer from severe imbalanced data problems. One reason is that SVM only needs a small fraction of data (called support vectors) for estimating both $\boldsymbol{\omega}$ and β, which mitigate the imbalanced data issue naturally.

Imbalanced data issues have been investigated in both statistics and machine learning. See an extensive survey in Chawla et al. (2004). Recently, Owen (2007) studied the asymptotic behavior of infinitely imbalanced binary logistic regression. In addition, Qiao and Liu (2009) and Qiao et al. (2010) proposed to use adaptive weighting approaches to overcome the imbalanced data issue.

In summary, the performance of DWD and SVM is different in the following ways: 1) The SVM direction usually has a larger variation and deviates farther from the Bayes rule direction than the DWD direction does, which are indicators of overfitting HDLSS data. 2) The SVM intercept is not sensitive to the imbalanced data, but the DWD intercept is. This motivates us to investigate their similarity and differences. In the next section, a new family of classifier will be proposed, which unifies the above two classifiers.

3 FLAME Family

In this section, we introduce FLAME, a family of classifiers which is motivated by a thorough investigation of the loss functions of SVM and DWD in Section 3.1. The formulation and implementation of the FLAME classifiers are given in Section 3.2.

3.1 SVM and DWD loss functions

The key factors that drive the very distinct performances of the SVM and the DWD methods are their associated loss functions (see Figure 4.)
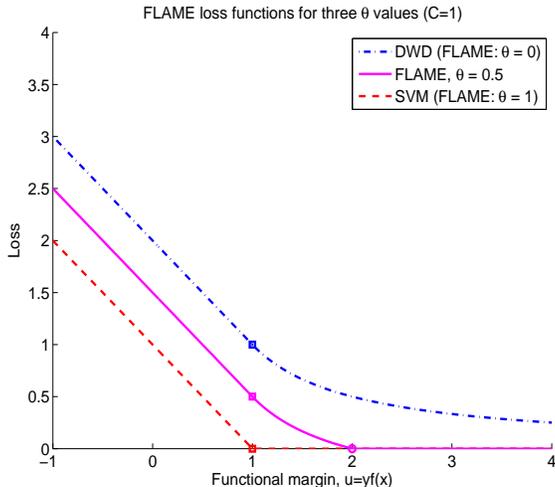
Figure 4: FLAME loss functions for three θ values: $\theta = 0$ (equivalent to SVM/Hinge loss), $\theta = 0.5$, $\theta = 1$ (equivalent to DWD). The parameter C is set to be 1.

Figure 4 displays the loss functions of SVM, DWD and FLAME with some specific tuning parameters. SVM uses the Hinge loss function, $H(u) = (1 - u)_+$ (the red dashed curve in Figure 4), where u corresponds to the functional margin $u \equiv yf(\boldsymbol{x})$. Note that the functional margin u can be viewed as the distance of vector $\boldsymbol{x}$ from the separating hyperplane (defined by $\{\boldsymbol{x} : f(\boldsymbol{x}) = 0\}$). When $u > 0$ and is large, the data vector is correctly classified and is far from the separating hyperplane; and $u < 0$, the data vector is wrongly classified. Note that when $u > 1$, the corresponding Hinge loss equals zero. Thus, only those observations with

$u \leq 1$ will contribute to estimating $\boldsymbol{\omega}$ and β. These observations are called *support vectors*. This is why SVM is insensitive to observations that are far away from the decision boundary, and why it is less sensitive to the imbalanced data issue. However, the only influence by the support vectors makes the SVM solution subject to the overfitting (data-piling) issue. This can be explained by that the optimization of SVM would try to push vectors towards small loss, *i.e.*, large functional margin u. But once a vector is pushed to the point where $u = 1$, the optimization lacks further incentive to continue pushing it towards a larger function margin as the Hinge loss cannot be reduced for this vector. Therefore many data vectors are piling along the hyperplanes corresponding to $u = 1$. Data-piling is bad for generalization because small turbulence to the support vectors could lead to big difference of the discriminant direction vector (recall the examples in Section 2.1).

The DWD method corresponds to a different DWD loss function,

$$V(u) = \begin{cases} 2\sqrt{C} - Cu & \text{if } u \leq \frac{1}{\sqrt{C}}, \\ 1/u & \text{otherwise.} \end{cases} \tag{1}$$

Here C is a pre-defined constant. Figure 4 shows the DWD loss function with $C = 1$. It is clear that the DWD loss function is very similar to the SVM loss function when u is small (both are linearly decreasing with respect to u). The major difference is that the DWD loss is always positive. This property will make the DWD method behave in a very different way than SVM. As there is always incentive to make the function margin to be larger (and the loss to be smaller), the DWD loss function kills data-piling, and mitigates the overfitting issue for HDLSS data.

On the other hand, the DWD loss function makes the DWD method very sensitive to the imbalanced data issue, since each observation will have some influence, and thus the larger class will have larger influence. The decision boundary of the DWD method will intend to ignore the smaller class, because sacrificing the smaller class (boundary being closer to the smaller class and farther from the larger class) can lead to a dramatic reduction of the loss, which ultimately lead to a minimized overall loss.

3.2 FLAME

We propose to borrow strengths from both methods to simultaneously deal with both the imbalanced data and the overfitting (data-piling) issues. We first highlight the connections between the DWD loss and an modified version of the Hinge loss (of SVM). Then we modify the DWD loss so that samples far from the classification boundary will have zero loss.

Let $f(\boldsymbol{x}) = \boldsymbol{x}^T \boldsymbol{\omega} + \beta$. The formulation of SVM can be rewritten (see details in the appendix) in the form of $\underset{\boldsymbol{\omega}, \beta}{\mathrm{argmin}} \sum_i H^*(y_i f(\boldsymbol{x}_i))$, s.t. $\|\boldsymbol{\omega}\|^2 \leq 1$ where the modified Hinge loss function H^* is defined as

$$H^*(u) = \begin{cases} \sqrt{C} - Cu & \text{if } u \leq \frac{1}{\sqrt{C}}, \\ 0 & \text{otherwise.} \end{cases} \tag{2}$$

Comparing the DWD loss (1) and this modified Hinge loss (2), one can easily see their connections: for $u \leq \frac{1}{\sqrt{C}}$, the DWD loss is greater than the Hinge loss of SVM by an exact constant $\sqrt{C}$, and for $u > \frac{1}{\sqrt{C}}$, the DWD loss is $1/u$ while the SVM Hinge loss equals 0. Clearly the modified Hinge loss (2) is the result of soft-thresholding the DWD loss at $\sqrt{C}$. In other words, SVM can be seen as a special case of DWD where the losses of those vectors with $u = y_i f(\boldsymbol{x}_i) > 1/\sqrt{C}$ are shrunken to zero. To allow different levels of soft-thresholding, we propose to use a new loss function which (soft-)thresholds the DWD loss function by constant $\theta\sqrt{C}$ where $0 \leq \theta \leq 1$, that is, a fraction of $\sqrt{C}$. The new loss function is

$$L(u) = \left[V(u) - \theta\sqrt{C} \right]_+ = \begin{cases} (2 - \theta)\sqrt{C} - Cu & \text{if } u \leq \frac{1}{\sqrt{C}}, \\ 1/u - \theta\sqrt{C} & \text{if } \frac{1}{\sqrt{C}} \leq u < \frac{1}{\theta\sqrt{C}}, \\ 0 & \text{if } u \geq \frac{1}{\theta\sqrt{C}}, \end{cases} \tag{3}$$

that is, to reduce the DWD loss by a constant, and truncate it at 0. The magenta solid curve in Figure 4 is the FLAME loss when $C = 1$ and $\theta = 0.5$. This simple but useful modification unifies the DWD and SVM methods. When $\theta = 1$, the new loss function (when $C = 1$) reduces to the SVM Hinge loss function; while when $\theta = 0$, it remains as the DWD loss.

Note that $L(u) = 0$ for $u > 1/(\theta\sqrt{C})$. Thus, those data vectors with large functional margins will have zero loss. For DWD loss, because it corresponds to $\theta = 0 \Rightarrow 1/(\theta\sqrt{C}) = \infty$,

no data vector can have zero loss. For SVM loss, all the data vector with $u > 1/(\theta\sqrt{C}) = 1/\sqrt{C}$ will have zero loss. Training a FLAME classifier with $0 < \theta < 1$ can be interpreted as sampling a portion of data which are farther from the boundary than $1/\theta\sqrt{C}$ and assign zero loss to them. Alternatively, it can be viewed as sampling data that are closer to the boundary than $1/\theta\sqrt{C}$ and assign positive loss to them. Note that the larger θ is, the fewer data are sampled to have positive loss. As one can flexibly choose θ, the new classification method with this new loss function is called the FLexible Assortment MachinE (FLAME).

FLAME can be implemented by a Second-Order Cone Programming algorithm (Toh et al., 1999, Tütüncü et al., 2003). Let $\theta \in [0, 1]$ be the FLAME parameter. The proposed method minimizes $\min_{\boldsymbol{\omega}, b, \boldsymbol{\xi}} \sum_{i=1}^{n} \left(\dfrac{1}{r_i} + C\xi_i - \theta\sqrt{C} \right)_+$. A slack variable $\varphi_i \geq 0$ can be introduced to absorb the $(\cdot)_+$ function. The optimization of the FLAME can be written as

$$\min_{\boldsymbol{\omega}, b, \boldsymbol{\xi}} \sum_i \varphi_i,$$

$$\text{s.t.} \quad \left(\frac{1}{r_i} + C\xi_i - \theta\sqrt{C} \right) - \varphi_i \leq 0, \ \varphi_i \geq 0,$$

$$r_i = y_i(\boldsymbol{x}_i^T \boldsymbol{\omega} + \beta) + \xi_i, \ r_i \geq 0 \text{ and } \xi_i \geq 0,$$

$$\|\boldsymbol{\omega}\|^2 \leq 1.$$

A `Matlab` routine has been implemented and is available at the authors' personal websites. See the online supplementary materials for more details on the implementation.

4   Choice of Parameters

There are two tuning parameters in the FLAME model: one is the C, inherited from the DWD loss, which controls the amount of allowance for misclassification; the other is the FLAME parameter θ, which controls the level of soft-thresholding. Similar to the discussion in DWD (Marron et al., 2007), the classification performance of FLAME is insensitive to different values of C. In addition, it can be shown for any C, FLAME is Fisher consistent, by applying the general results in Lin (2004). Thus, the default value for C as proposed

in Marron et al. (2007) will be used in FLAME. In this section, we introduce two ways of choosing the second parameter θ. As the property and the performance of FLAME depends on the choice of this parameter, it is important to select the right amount of thresholding. In the following two subsections, we discuss two options of choosing parameter θ. The first option is based on empirical plots resulting from the training data and is of practical useful, and it is the θ value that we suggest. The second option is motivated by a theoretical consideration and is heuristically meaningful as well.

4.1   Parameter from performance trade-off

Note that, an optimal θ depends on the nature of the data and problems that users have. The optimal θ also depends on two performance measures: insensitive to the imbalanced data, and resistance to overfitting. However, without prior knowledge, we may want to have a "good" trade-off between them. In this subsection, we suggest the following method to choose θ if no further information is provided.

As will become clear from the simulation examples shown in Section 6.3, we have observed that several performance measures for the FLAME classifiers, for example, the within-group error (see the definition in Section 6 and also in Qiao and Liu (2009)), are monotonically decreasing with respect to θ. On the other hand, performance measures such as the *RankComp* (see also in Section 6), are monotonically increasing functions of θ. The RankComp measure is more related to the overfitting phenomena, and the within-group error is designed for measuring the performance against the imbalanced data. The lesson is that with θ increases, FLAME becomes less sensitive to the imbalanced data issue, but is subject to more overfitting. This motivates us to use the following parameter: the two curves of the two measures are normalized to be between 0 and 1. When $\theta = 0$, the FLAME classifier (equivalent to DWD) has the smallest RankComp measure 0, but the largest within-group error 1. When $\theta = 1$, the FLAME classifier (equivalent to SVM) has the smallest within-group error 0, but the largest RankComp 1. The suggested θ is chosen as the value where the two normalized

13

curves intersect, that is, the normalized within-group error is the same as the normalized RankComp for this θ. This suggested parameter represents a natural trade-off between the two measures: neither measure is absolutely optimal, but each measure compromises by the same relative amount. This suggested parameter is called the *equal-trade-off parameter*.

4.2   Parameter from adaptive modeling

Having observed that the DWD discrimination direction is usually closer to the Bayes rule direction, but its location term β is sensitive to the imbalanced data issue, we propose the following alternative data-driven approach to select an appropriate θ. Without loss of generality, we assume that the negative class is the majority class with sample size n_- and the positive class is the minority class with sample size n_+. We point out that the main reason that DWD is sensitive to the imbalanced data issue is that it uses all vectors in the *majority* class to build up a classifier. A heuristic strategy to correct this would be to force the optimization to use the same number of vectors from both classes to build up a classifier: we first apply DWD to the data set, and calculate the distances of all data in the majority class to the current DWD classification boundary; we then train FLAME with a carefully chosen parameter θ which assigns positive loss to the closet n_+ data vectors in the majority class to the classification boundary. As a consequence, each class will have exactly n_+ vectors which have positive loss. In other words, while keeping the least imbalance (because we have the same numbers of vectors from both classes that have influence over the optimization), we obtain a model with the least possible overfitting (because $2n_+$ vectors have influence, instead of only the limited support vectors as in SVM.)

In practice, since the new FLAME classification boundary using the θ chosen above may be different from the initial DWD classification boundary, the n_+ closest points to the FLAME classification boundary may not be the same n_+ closest points to the DWD boundary. This means that it is not guaranteed that exactly n_+ points from the majority class will have positive loss. However, one can expect that reasonable approximation can

be achieved. Moreover, an iterative scheme for finding θ is introduced as follows in order to minimize such discrepancy.

For simplicity, we let $(\boldsymbol{x}_i, y_i)$ with index i be an observation from the positive/minority class and $(\boldsymbol{x}_j, y_j)$ with index j be an observation from the negative/majority class.

ALGORITHM 1. (**Adaptive parameter**)

1. Initiate $\theta_0 = 0$.

2. For $k = 0, 1, \cdots,$

    (a) Solve FLAME solutions $\boldsymbol{\omega}(\theta_k)$ and $\beta(\theta_k)$ given parameter θ_k.

    (b) Let $\theta_{k+1} = \max\left(\theta_k, \ \left\{g_{(n_+)}(\theta_k)\sqrt{C}\right\}^{-1}\right)$, where $g_j(\theta_k)$ is the functional margin $u_j \equiv y_j(\boldsymbol{x}_j^T \boldsymbol{\omega}(\theta_k) + \beta(\theta_k))$ of the jth vector in the negative/majority class and $g_{(l)}(\theta_k)$ is the lth order statistic of these functional margins.

3. When $\theta_k = \theta_{k-1}$, the iteration stops.

The goal of this algorithm is to make $g_{(n_+)}(\theta_k)$ to be the greatest functional margin among all the data vectors that have positive loss in the negative/majority class. To achieve this, we calibrate θ by aligning $g_{(n_+)}(\theta_k)$ to the turning point $u = 1/(\theta\sqrt{C})$ in the definition of the FLAME loss (3), that is $g_{(n_+)}(\theta_k) = 1/(\theta\sqrt{C}) \Rightarrow \theta = \left(g_{(n_+)}(\theta_k)\sqrt{C}\right)^{-1}$.

We define the equivalent sample objective function of FLAME for the iterative algorithm above, $s(\boldsymbol{\omega}, \beta, \theta) = \dfrac{1}{n_+ + n_-}\left[\displaystyle\sum_{i=1}^{n_+} L((\boldsymbol{x}_i^T \boldsymbol{\omega} + \beta), \theta) + \sum_{j=1}^{n_-} L(-(\boldsymbol{x}_j' \boldsymbol{\omega} + \beta), \theta)\right] + \dfrac{\lambda}{2}\|\boldsymbol{\omega}\|^2$. Then the convergence of this algorithm is shown in Theorem 1.

THEOREM 1. *In Algorithm 1, $s(\boldsymbol{\omega}_k, \beta_k, \theta_k)$ is non-increasing in k. As a consequence, Algorithm 1 converges to a stationary point $s(\boldsymbol{\omega}_\infty, \beta_\infty, \theta_\infty)$ where $s(\boldsymbol{\omega}_k, \beta_k, \theta_k) \geq s(\boldsymbol{\omega}_\infty, \beta_\infty, \theta_\infty)$. Moreover, Algorithm 1 terminates finitely.*

Ideally, one would hope to get an optimal parameter θ^* which satisfies $\theta^* = \left(g_{(n_+)}(\theta^*)\sqrt{C}\right)^{-1}$. In practice, θ_∞ will approximate θ^* very well. In addition, we notice that one-step iteration usually gives decent results for simulation examples and some real examples.

5 Theoretical Properties

In this section, several important theoretical properties of the FLAME classifiers are investigated. We first prove the Fisher consistency (Lin, 2004) of the FLAME in Section 5.1. As one focus of this paper is imbalanced data classification, the asymptotic properties for FLAME under extremely imbalanced data setting is studied in Section 5.2. Lastly, a novel HDLSS asymptotics where n is fixed and $d \to \infty$, the other focus of this article, is studied in Section 5.3.

5.1 Fisher consistency

Fisher consistency is a very basic property for a classifier. A classifier is Fisher consistent means that the minimizer of the conditional risk of the classifier given observation $\boldsymbol{x}$ has the same sign as the Bayes rule, $\underset{k \in \{+1, -1\}}{\operatorname{argmax}} P(Y = k | \boldsymbol{X} = \boldsymbol{x})$. It has been shown that both SVM and DWD are Fisher consistent (Lin, 2004, Qiao et al., 2010). The following proposition states that the FLAME classifiers are Fisher consistent too.

PROPOSITION 2. *Let f^* be the global minimizer of $E[L(Yf(\boldsymbol{X}), \theta)]$, where $L(\cdot, \theta)$ is the loss function for FLAME given parameter θ. Then $\operatorname{sign}(f^*(\boldsymbol{x})) = \operatorname{sign}(P(Y = +1 | \boldsymbol{X} = \boldsymbol{x}) - 1/2)$.*

5.2 Asymptotics under imbalanced setting

In this subsection, we investigate the asymptotic performance of SVM, DWD and FLAME. The asymptotic setting we focus on is when the minority sample size n_+ is fixed and the majority sample size $n_- \to \infty$, which is similar to the setting in Owen (2007). We will show that DWD is sensitive to the imbalanced data, while FLAME with proper choices of parameter θ and SVM are not.

Let $\overline{\boldsymbol{x}}_+$ be the sample mean of the positive/minority class. Theorem 3 shows that in the imbalanced data setting, when the size of the negative/majority class grows while that of the positive/minority class is fixed, the intercept term for DWD tends to negative infinity, in

the order of $\sqrt{m}$. Therefore, DWD will classify all the observations to the negative/majority class, that is, the minority class will be 100% misclassified.

THEOREM 3. *Let n_+ be fixed. Assume that the conditional distribution of the negative majority class $F_-(\boldsymbol{x})$ surrounds $\overline{\boldsymbol{x}}_+$ by the definition given in* Owen (2007), *and that γ is a constant satisfying* $\inf\limits_{\|\boldsymbol{\omega}\|=1} \int_{(\boldsymbol{x}-\overline{\boldsymbol{x}}_+)'\boldsymbol{\omega}>0} dF_-(\boldsymbol{x}) > \gamma \geq 0$, *then the DWD intercept $\widehat{\beta}$ satisfies*

$$\widehat{\beta} < -\sqrt{\frac{\gamma}{C}m} - \overline{\boldsymbol{x}}_+^T\boldsymbol{\omega} = -\sqrt{\frac{n_-\gamma}{n_+C}} - \overline{\boldsymbol{x}}_+^T\boldsymbol{\omega}.$$

In Section 4.2, we have introduced an iterative approach to select the parameter θ. Theorem 4 shows that with the optimal parameter θ^* found by Algorithm 1, the discriminant direction of FLAME is in the same direction of the vector that joins the sample mean of the positive class and the *tilted* population mean of the negative class. Moreover, in contrast to DWD, the intercept term of FLAME in this case is finite.

THEOREM 4. *Suppose that $n_- \gg n_+$ and $\boldsymbol{\omega}^*$ and β^* are the FLAME solutions trained with the parameter θ^* that satisfies $\theta^* = \left(g_{(n_+)}(\theta^*)\sqrt{C}\right)^{-1}$. Then $\boldsymbol{\omega}^*$ and β^* satisfy that*

$$\boldsymbol{\omega}^* = \frac{C}{(1+m)\lambda}\left[\overline{\boldsymbol{x}}_+ - \frac{\int(\boldsymbol{x}^T\boldsymbol{\omega}^* + \beta^*)^{-2}\boldsymbol{x}dF_-(\boldsymbol{x} \mid E)}{\int(\boldsymbol{x}^T\boldsymbol{\omega}^* + \beta^*)^{-2}dF_-(\boldsymbol{x} \mid E)}\right],$$

where E is the event that $[Y(\boldsymbol{X}^T\boldsymbol{\omega}^* + \beta^*)]^{-1} \geq \theta^*\sqrt{C}$ where $(\boldsymbol{X}, Y)$ is a random sample from the negative/majority class, and that

$$\int(\boldsymbol{x}^T\boldsymbol{\omega}^* + \beta^*)^{-2}dF_-(\boldsymbol{x} \mid E) = \frac{n_+}{n^o}C, \text{where } 0 < n^o \leq n_+.$$

As a consequence of Theorem 4, when $m = n_-/n_+ \to \infty$, we have $\|\boldsymbol{\omega}^*\| \to 0$. Since the right-hand-side of the last equation above is positive finite, β^* does not diverge. In addition, since $\mathrm{P}(\overline{E}) \to 1$ with probability converging to 1, $\beta^* < -1/(\theta\sqrt{C})$.

The following theorem shows the performance of SVM under the imbalanced data context, which completes our comparisons between SVM, DWD and FLAME.

THEOREM 5. *Suppose that $n_- \gg n_+$. The solutions $\widehat{\boldsymbol{\omega}}$ and $\widehat{\beta}$ to SVM satisfy that*

$$\widehat{\boldsymbol{\omega}} = \frac{1}{(1+m)\lambda}\left\{\overline{\boldsymbol{x}}_+ - \int \boldsymbol{x}dF_-(\boldsymbol{x} \mid G)\right\},$$

where G is the event that $1 - Y(\boldsymbol{X}^T\widehat{\boldsymbol{\omega}} + \widehat{\beta}) > 0$ where $(\boldsymbol{X}, Y)$ is a random sample from the negative/majority class, and that

$$\mathrm{P}(\overline{G}) = \mathrm{P}(1 + \boldsymbol{X}^T\widehat{\boldsymbol{\omega}} + \widehat{\beta} \leq 0) = 1 - 1/m.$$

The last statement in Theorem 5 means that with probability converging to 1, $\widehat{\beta} \leq -1$. However, note this is the only restriction that SVM solution has for the intercept term (recall that the counterpart in DWD is $\widehat{\boldsymbol{\beta}} < -\sqrt{\frac{\gamma}{C}m} - \overline{\boldsymbol{x}}_+^T\boldsymbol{\omega}$).

5.3 High-dimensional, low-sample size asymptotics

HDLSS data are emerging in many areas of scientific research. The HDLSS asymptotics is a recently developed theoretical framework. Hall et al. (2005) gave a geometric representation for the HDLSS data, which can be used to study these new 'n fixed, $d \to \infty$' asymptotic properties of binary classifiers such as SVM and DWD. Ahn et al. (2007) weakened the conditions under which the representation holds. Qiao et al. (2010) improved the conditions and applied this representation to investigate the performance of the weighted DWD classifier. The same geometric representation can be used to analyze FLAME. See summary of some previous HDLSS results in the online supplementary materials. We develop the HDLSS asymptotic properties of the FLAME family by providing conditions in Theorem 6 under which the FLAME classifiers always correctly classify HDLSS data.

We first introduce the notations and give some regularity assumptions, then state the main theorem. Let $k \in \{+1, -1\}$ be the class index. For the kth class and given a fixed n_k, consider a sequence of random data matrices $\boldsymbol{X}_1^k, \boldsymbol{X}_2^k, \cdots \boldsymbol{X}_d^k, \cdots$, indexed by the number of rows d, where each column of $\boldsymbol{X}_d^k$ is a random observation vector from $\mathbb{R}^d$ and each row represents a variable. Assume that each column of $\boldsymbol{X}_d^k$ comes from a multivariate distribution with dimension d and with covariance matrix $\boldsymbol{\Sigma}_d^k$ independently. Let $\lambda_{1,d}^k \geq \cdots \geq \lambda_{d,d}^k$ be the eigenvalues of the covariance, and $\left(\sigma_d^k\right)^2 = d^{-1}\sum_{i=1}^d \lambda_{i,d}^k$ the average eigenvalue. The eigenvalue decomposition of $\boldsymbol{\Sigma}_d^k$ is $\boldsymbol{\Sigma}_d^k = \boldsymbol{V}_d^k\boldsymbol{\Lambda}_d^k\left(\boldsymbol{V}_d^k\right)^T$. We may define the square root of $\boldsymbol{\Sigma}_d^k$ as $\left(\boldsymbol{\Sigma}_d^k\right)^{1/2} = \boldsymbol{V}_d^k\left(\boldsymbol{\Lambda}_d^k\right)^{1/2}$, and the inverse square root $\left(\boldsymbol{\Sigma}_d^k\right)^{-1/2} =$

$\left(\mathbf{\Lambda}_d^k\right)^{-1/2} \left(\mathbf{V}_d^k\right)^T$. With minimal abuse of notation, let $\mathbb{E}(\mathbf{X}_d^k)$ denote the expectation of columns of $\mathbf{X}_d^k$. Lastly, the $n^k \times n^k$ *dual* sample covariance matrix is denoted by $\mathbf{S}_{D,d}^k = d^{-1}\left\{\mathbf{X}_d^k - \mathbb{E}(\mathbf{X}_d^k)\right\}^T \left\{\mathbf{X}_d^k - \mathbb{E}(\mathbf{X}_d^k)\right\}$.

ASSUMPTION 1. There are five components:

(i) Each column of $\mathbf{X}_d^k$ has mean $\mathbb{E}(\mathbf{X}_d^k)$ and the covariance matrix $\mathbf{\Sigma}_d^k$ of its distribution is positive definite.

(ii) The entries of $\mathbf{Z}_d^k \equiv \left(\mathbf{\Sigma}_d^k\right)^{-\frac{1}{2}} \left\{\mathbf{X}_d^k - \mathbb{E}(\mathbf{X}_d^k)\right\} = \left(\mathbf{\Lambda}_d^k\right)^{-\frac{1}{2}} \left(\mathbf{V}_d^k\right)^T \left\{\mathbf{X}_d^k - \mathbb{E}(\mathbf{X}_d^k)\right\}$ are independent.

(iii) The fourth moment of each entry of each column is uniformly bounded by $M > 0$ and the Wishart representation holds for each dual sample covariance matrix $\mathbf{S}_{D,d}^k$ associated with $\mathbf{X}_d^k$, that is,

$$dS_{D,d}^k = \left\{\left(\mathbf{Z}_d^k\right)^T \left(\mathbf{\Lambda}_d^k\right)^{1/2} \left(\mathbf{V}_d^k\right)^T\right\} \left\{\mathbf{V}_d^k \left(\mathbf{\Lambda}_d^k\right)^{1/2} \mathbf{Z}_d^k\right\} = \sum_{i=1}^d \lambda_{i,d}^k \mathbf{W}_{i,d}^k,$$

where $\mathbf{W}_{i,d}^k \equiv \left(Z_{i,d}^k\right)^T Z_{i,d}^k$ and $Z_{i,d}$ is the ith row of $\mathbf{Z}_d^k$ defined above. It is called Wishart representation because if $\mathbf{X}_d^k$ is Gaussian, then each $\mathbf{W}_{i,d}^k$ follows the Wishart distribution $\mathcal{W}_{n^k}(1, \mathbf{I}_{n^k})$ independently.

(iv) The eigenvalues of $\mathbf{\Sigma}_d^k$ are sufficiently diffused, in the sense that

$$\epsilon_d^k = \frac{\sum_{i=1}^d (\lambda_{i,d}^k)^2}{\left(\sum_{i=1}^d \lambda_{i,d}^k\right)^2} \to 0 \text{ as } d \to \infty. \tag{4}$$

(v) The sum of the eigenvalues of $\mathbf{\Sigma}_d^k$ is the same order as d, in the sense that $\left(\sigma_d^k\right)^2 = O(1)$ and $1/\left(\sigma_d^k\right)^2 = O(1)$.

ASSUMPTION 2. The distance between the two population expectations satisfies,

$$d^{-1}\left\|\mathbb{E}(\mathbf{X}_d^{(+1)}) - \mathbb{E}(\mathbf{X}_d^{(-1)})\right\|^2 \to \mu^2, \text{ as } d \to \infty.$$

Moreover, there exist constants σ^2 and τ^2, such that

$$\left(\sigma_d^{(+1)}\right)^2 \to \sigma^2, \text{ and } \left(\sigma_d^{(-1)}\right)^2 \to \tau^2.$$

Let $\nu^2 \equiv \mu^2 + \sigma^2/n_+ + \tau^2/n_-$. The following theorem gives the sure classification condition for FLAME, which includes SVM and DWD as special cases.

19

THEOREM 6. *Without loss of generality, assume that $n_+ \leq n_-$. The situation of $n+ > n_-$ is similar and omitted.*

- *If either one of the following three conditions is satisfied,*

    1. *for $\theta \in \left[0, (1 + \sqrt{m^{-1}})/(\nu\sqrt{dC})\right)$, $\mu^2 > (n_-/n_+)^{\frac{1}{2}}\sigma^2/n_+ - \tau^2/n_- > 0$;*

    2. *for $\theta \in \left[(1 + \sqrt{m^{-1}})/(\nu\sqrt{dC}), 2/(\nu\sqrt{dC})\right)$, $\mu^2 > T - \tau^2/n_- > 0$ where $T := \left(1/(2\theta\sqrt{dC}) + \sqrt{1/(4\theta^2 dC) + \sigma^2/n_+}\right)^2 - \sigma^2/n_+$;*

    3. *for $\theta \in \left[2/(\nu\sqrt{dC}), 1\right]$, $\mu^2 > \sigma^2/n_+ - \tau^2/n_- > 0$,*

    then for a new data point $\boldsymbol{x}_0^+$ from the positive class (+1),

    $\mathrm{P}(\boldsymbol{x}_0^+$ *is correctly classified by FLAME*$) \to 1$, *as $d \to \infty$.*

    Otherwise, the probability above $\to 0$.

- *If either one of the following three conditions is satisfied,*

    1. *for $\theta \in \left[0, (1 + \sqrt{m^{-1}})/(\nu\sqrt{dC})\right)$, $(n_-/n_+)^{\frac{1}{2}}\sigma^2/n_+ - \tau^2/n_- > 0$;*

    2. *for $\theta \in \left[(1 + \sqrt{m^{-1}})/(\nu\sqrt{dC}), 2/(\nu\sqrt{dC})\right)$, $T - \tau^2/n_- > 0$;*

    3. *for $\theta \in \left[2/(\nu\sqrt{dC}), 1\right]$, $\sigma^2/n_+ - \tau^2/n_- > 0$,*

    then for any $\mu > 0$, for a new data point $\boldsymbol{x}_0^-$ from the negative class (−1),

    $\mathrm{P}(\boldsymbol{x}_0^-$ *is correctly classified by FLAME*$) \to 1$, *as $d \to \infty$.*

Theorem 6 has two parts. The first part gives the conditions under which FLAME correctly classifies a new data point from the positive class, and the second part is for the negative class. Each part lists three conditions based on three disjoint intervals of parameter θ. Note the first and third intervals of each part generalize results which were shown to hold only for DWD and SVM before (*c.f.* Theorem 1 and Theorem 2 in Hall et al., 2005). In particular, it shows that all the FLAME classifiers with θ falling into the first interval behave like DWD asymptotically. Similarly, all the FLAME classifiers with θ falling into the third interval behave like SVM asymptotically. This partially explains the shape of the within-group error curve that we will show in Figures 6, S.2, and S.3, which we will discuss in the next section.

In the first part, the condition for other FLAMEs (with θ in the second interval) is weaker

than the DWD-like FLAMEs (in the first interval), but stronger than the SVM-like FLAMEs (in the third interval). This means that it is easier to classify a new data point from the positive/minority class by SVM, than by an intermediate FLAME, which is easier than by DWD. Note that when $n_+ \leq n_-$, the hyperplane for FLAME is in general closer to the positive class.

In terms of classifying data points from the negative class, the order of the difficulties among DWD, FLAME and SVM reverses.

6 Simulations

FLAME is not only a unified representation of DWD and SVM, but also introduces a new family of classifiers which are capable of avoiding the overfitting HDLSS data issue and the sensitivity to imbalanced data issue. In this section, we use simulations to show the performance of FLAME at various parameter levels. We will show that with a range of carefully chosen parameters, FLAME can outperform both the DWD and the SVM methods in various simulation settings.

6.1 Measures of performance

Before we introduce our simulation examples, we first introduce the performance measures in this paper. Note that the Bayes rule classifier can be viewed as the "gold standard" classifier. In our simulation settings, we assume that data are generated from two Gaussian populations $MVN(\boldsymbol{\mu}_{\pm}, \boldsymbol{\Sigma})$ with different mean vectors $\boldsymbol{\mu}_+$ and $\boldsymbol{\mu}_-$ and same covariance matrices $\boldsymbol{\Sigma}$. This setting leads to the following Bayes rule.

$$\text{sign}(\boldsymbol{x}^T\boldsymbol{\omega}_B + \beta_B) \text{ where } \boldsymbol{\omega}_B = \boldsymbol{\Sigma}^{-1}(\boldsymbol{\mu}_+ - \boldsymbol{\mu}_-) \text{ and } \beta_B = -\frac{1}{2}(\boldsymbol{\mu}_+ + \boldsymbol{\mu}_-)'\boldsymbol{\omega}_B. \quad (5)$$

Five performance measures are evaluated in this paper:

1. The mean within-class error (MWE) for out-of-sample test set, which is defined as

$$MWE = \frac{1}{2n_+} \sum_{i=1}^{n_+} \mathbb{1}(\widehat{Y}_i^+ \neq Y_i^+) + \frac{1}{2n_-} \sum_{j=1}^{n_-} \mathbb{1}(\widehat{Y}_j^- \neq Y_j^-)$$

2. The deviation of the estimated intercept β from the Bayes rule intercept β_B: $|\beta - \beta_B|$.

3. Dispersion: a measure of the stochastic variability of the estimated discrimination direction vector $\boldsymbol{\omega}$. The dispersion measure was introduced in Section 1, as the trace of the sample covariance of the resulting discriminant direction vectors: disperson $=$ $\text{Var}([\boldsymbol{\omega}_r]_{r=1:R})$ where R is the number of repeated runs.

4. Angle between the estimated discrimination direction $\boldsymbol{\omega}$ and the Bayes rule direction $\boldsymbol{\omega}_B$: $\angle(\boldsymbol{\omega}, \boldsymbol{\omega}_B)$.

5. RankComp$(\boldsymbol{\omega}, \boldsymbol{\omega}_B)$: In general, for two direction vectors $\boldsymbol{\omega}$ and $\boldsymbol{\omega}^*$, RankComp is defined as the proportion of the pairs of variables, among all $d(d-1)/2$ pairs, whose relative importances (in terms of their absolute values) given by the two directions are different, $i.e.$,

$$\text{RankComp}(\boldsymbol{\omega}, \boldsymbol{\omega}^*) \equiv \frac{1}{d(d-1)/2} \sum_{1 \leq i < j \leq d} \mathbb{1}\left\{ (|\omega_i| - |\omega_j|) \times (|\omega_i^*| - |\omega_j^*|) < 0 \right\},$$

where ω_i and ω_i^* are the ith components of the vectors $\boldsymbol{\omega}$ and $\boldsymbol{\omega}^*$ respectively.

We report the RankComp between the estimated direction $\boldsymbol{\omega}$ and the Bayes rule direction $\boldsymbol{\omega}_B$ to measure their closeness.

We will investigate these measures based on different dimensions d and different imbalance factors m.

6.2 Effects of dimensions and imbalanced data

In Section 1, a specific FLAME ($\theta = .5$) has been compared with SVM ($\theta = 1$) and DWD ($\theta = 0$) in Figure 1, and on average, its discriminant directions are closer to the Bayes rule direction $\boldsymbol{\omega}_B$ compared to the SVM directions, but is less close than the DWD directions. In this subsection, we will further investigate the performance of FLAME with several different values of θ, and compare them with DWD and SVM under various simulation settings.
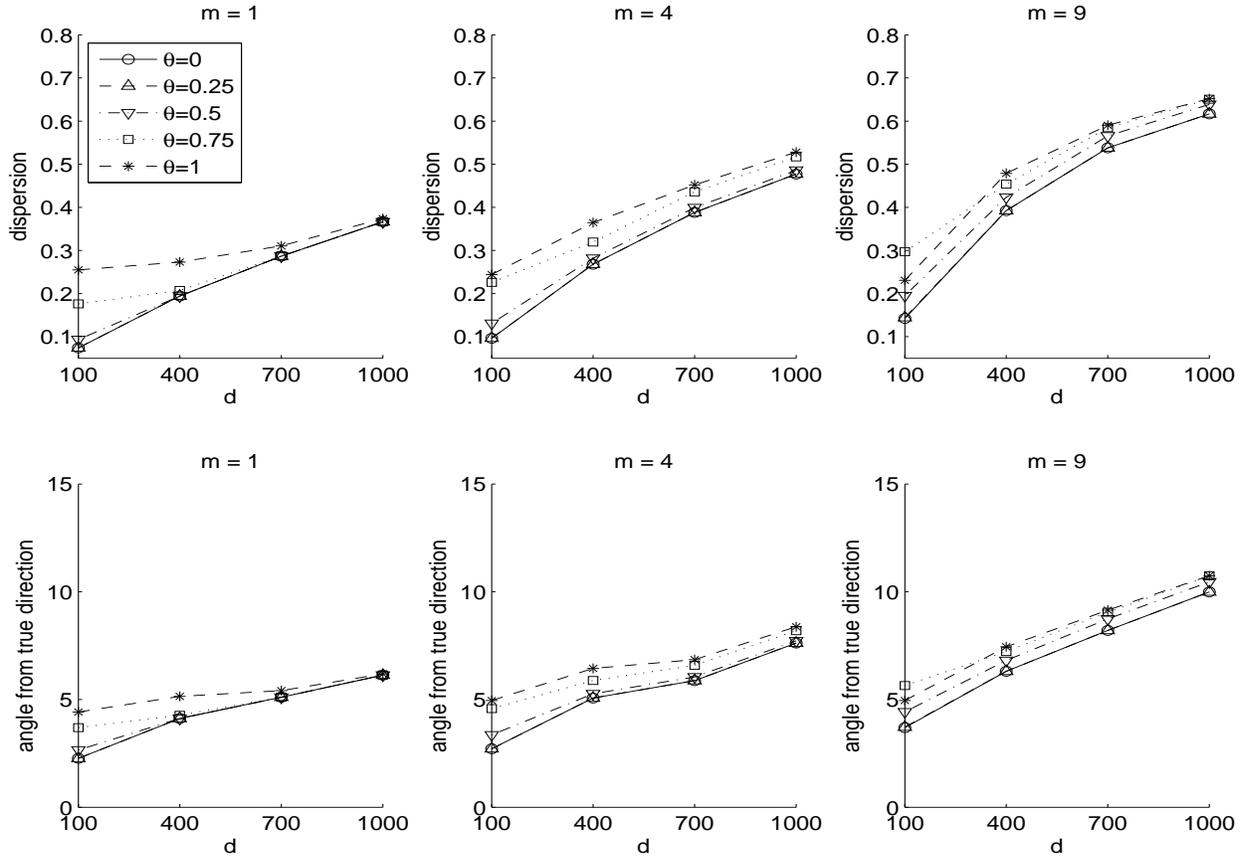
Figure 5: The dispersions (top row) and the angles between the FLAME direction and the Bayes direction (bottom row) for 50 runs of simulations, where the imbalance factors m are 1, 4 and 9 (the left, center and right panels), in the increasing dimension setting ($d = 100, 400, 700, 1000$; shown on the x-axes). The FLAME machines have $\theta = 0, 0.25, 0.5, 0.75, 1$ which are depicted using different curve styles (the first and the last cases correspond to DWD and SVM, respectively.) Note that with θ and the dimension d increase, both the dispersion and the deviation from the Bayes direction increase. The emergence of the imbalanced data (the increase of m) does not much deteriorate the FLAME directions except for large d.

Figure 5 shows the comparison results under the same simulation setting with various combinations of (d, m)'s. In this simulation setting, data are from multivariate normal distributions with identity covariance matrices $MVN_d(\boldsymbol{\mu}_\pm, \mathbf{I}_d)$, where $d = 100, 400, 700$ and 1000. We let $\boldsymbol{\mu}_0 = c(d, d-1, d-2, \cdots, 1)^T$ where $c > 0$ is a constant which scales $\boldsymbol{\mu}_0$ to have norm 2.7. Then we let $\boldsymbol{\mu}_+ = \boldsymbol{\mu}_0$ and $\boldsymbol{\mu}_- = -\boldsymbol{\mu}_0$. The imbalance factor varies among 1, 4 and 9 while the total sample size is 240. For each experiment, we repeat the simulation 50 times,

and plot the average performance measure in Figure 5. The Bayes rule is calculated according to (5). It is obvious that when the dimension increases, both the dispersion and the angle increase. They are indicators of overfitting HDLSS data. When the imbalance factor m increases, the two measures increases as well, although not as much as when the dimension increases. More importantly, it shows that when θ decreases (from 1 to 0, or equivalently FLAME changes from SVM to DWD), the dispersion and the angle both decrease, which is promising because it shows that FLAME improves SVM in terms of the overfitting issue.

6.3 Effects of tuning parameters with covariance

We also investigate the effect of different covariance structures, since independence structure among variables as in the last subsection is not realistic in real applications. We investigate three covariance structures: independent, interchangeable and block-interchangeable covariance. Data are generated from two multivariate normal distributions $MVN_{300}(\boldsymbol{\mu}_{\pm}, \boldsymbol{\Sigma})$ with $d = 300$. We fist let $\boldsymbol{\mu}_1 = (75, 74, 73, \cdots, 1, 0, 0, \cdots, 0)'$, then scale it by multiply a constant c such that the Mahalanobis distance between $\boldsymbol{\mu}_+ = c\boldsymbol{\mu}_1$ and $\boldsymbol{\mu}_- = -c\boldsymbol{\mu}_1$ equals 5.4, i.e., $(\boldsymbol{\mu}_+ - \boldsymbol{\mu}_-)'\boldsymbol{\Sigma}^{-1}(\boldsymbol{\mu}_+ - \boldsymbol{\mu}_-) = 5.4$. Note that this represents a reasonable signal-to-noise ratio.

We consider the FLAME machines with different parameter θ from a grid of 11 values $(0, 0.1, 0.2, \cdots, 1)$, and apply them to nine simulated examples (three different imbalance factors ($m = 2, 3, 4$) $\times$ three covariance structures). For the independent structure example, $\boldsymbol{\Sigma} = \mathbf{I}_{300}$; For the interchangeable structure example, $\boldsymbol{\Sigma}_{ii} = 1$ and $\boldsymbol{\Sigma}_{ij} = 0.8$ for $i \neq j$; For the block-interchangeable structure example, we let $\boldsymbol{\Sigma}$ be a block diagonal matrix with five diagonal blocks, the sizes of which are 150, 100, 25, 15, 10, and each block is an interchangeable covariance matrix with diagonal entries 1 and off-diagonal entries 0.8.

Figure 6 provides the summary results of the interchangeable structure example. Since the results are similar under different covariance structures, results from the other two covariance structures are included in the online supplementary materials to save space (Figure

for the independent structure, and Figure for the block-interchangeable covariance).
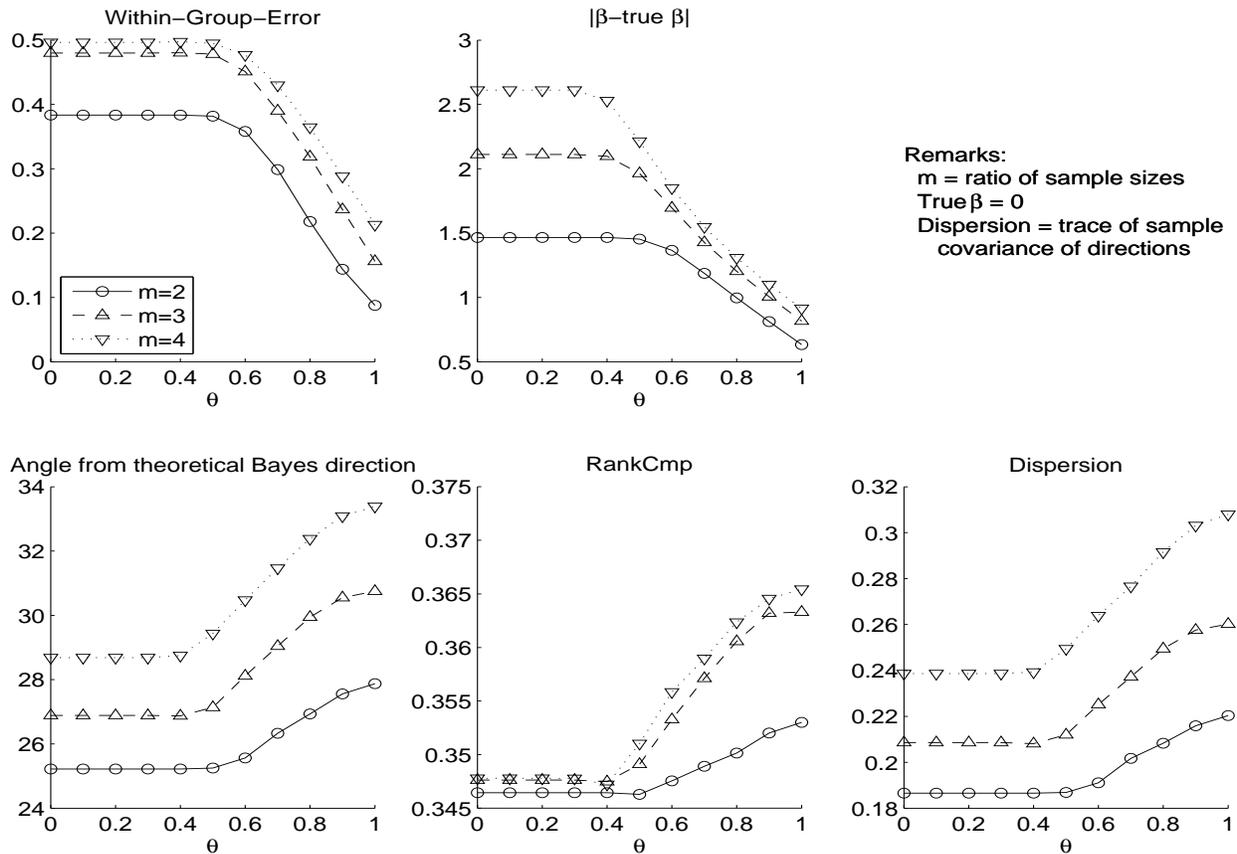
Figure 6: Interchangeable example. It can be seen that with FLAME turns from DWD to SVM (θ from 0 to 1), the within-class error decreases (top-left), thanks to the more accurate estimate of the intercept term (top-middle). On the other hand, this comes at the cost of larger deviation from the Bayes direction (bottom-left), incorrect rank of the importance of the variables (bottom-middle) and larger stochastic variability of the estimation directions (bottom-right).

In each plot, we include the within-group error (top-left), the absolute value of the difference between the estimated intercept and the Bayes intercept $|\beta - \beta_B|$ (top-middle), the angle between the estimated direction and the Bayes direction $\angle(\boldsymbol{\omega}, \boldsymbol{\omega}_B)$ (bottom-left), the RankComp between the estimated direction and the Bayes direction (bottom-middle) and the dispersion of the estimated directions (bottom-right).

We can see that in Figure 6 (and Figures S.2 and S.3 in the online supplementary materials), when we increase θ from 0 to 1, *i.e.*, when the FLAME moves from the DWD end to

the SVM end, the within-group error decreases. This is mostly due to the fact that the intercept term β comes closer to the Bayes rule intercept β_B. On the other hand, the estimated direction is deviating from the true direction (larger angle), is giving the wrong rank of the variables (larger RankComp) and is more unstable (larger dispersion). Similar observations hold for the other two covariance structures, with one exception in the block interchangeable setting (Figure S.3) where the RankComp first decreases then increases.

In the entire FLAME family, DWD represents one extreme which provides better estimation of the direction, is closer to the Bayes direction, provides the right order for all variables, and is more stable. But it suffers from the inaccurate estimation of β in the presence of imbalanced data; SVM represents the other extreme, which is not sensible to imbalanced data and usually provides a good estimation of β, but is in general outperformed by DWD in terms of closeness to the Bayes optimal direction. In most situations, within the FLAME family, there is no single machine that is better than the both ends from the two aspects at the same time.

The observations above motivate the use of the equal-trade-off parameter introduced in Section 4.1. In the next subsection, we will compare this parameter choice with other options.

6.4   Comparison of parameter options

We have suggested an equal-trade-off parameter based on the plots of within-group error and the RankComp (see Section 4.1) and justified the use of an adaptive θ based on an iterative procedure (see details in Section 4.2). Figure 7 compares FLAME with these two choices, and with $\theta = 0$ (DWD) and 1 (SVM). Various covariance structures (independent, interchangeable and block-interchangeable) are investigated. To save the space, we only show the results for the block-interchangeable dependence structure as this is more realistic in many real applications in genomic science and other applications. Here the full dimensions ($d = 80, 100, 180, 300, 600$ or 900) are divided to three blocks (50%, 25% and 25% of d). The
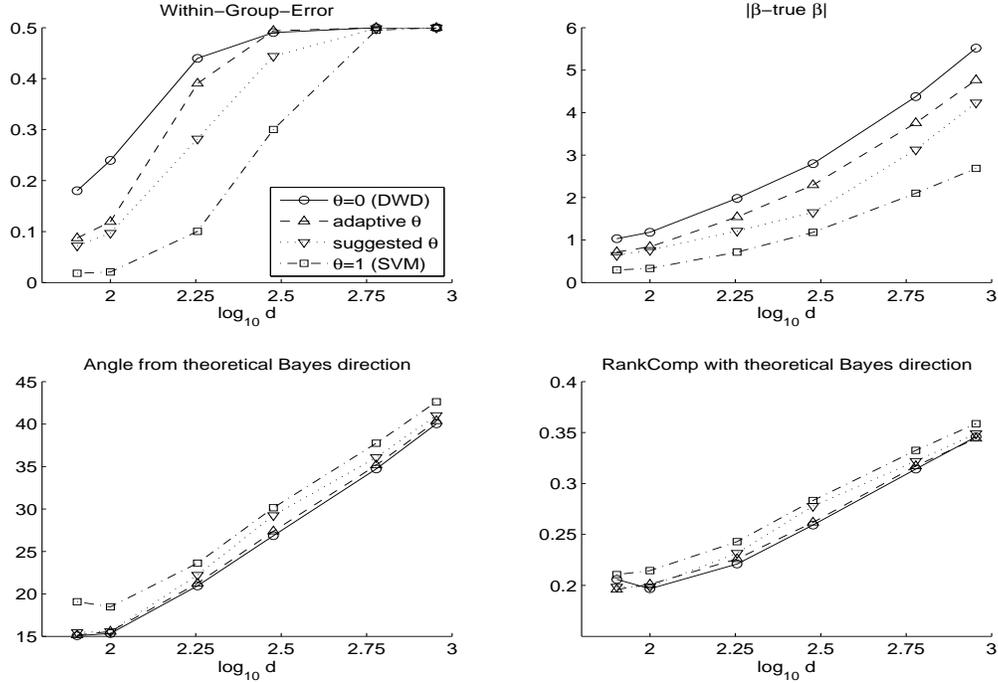
Figure 7: Comparison of four FLAMEs with $\theta = 0,\ 1$, the suggested θ introduced in Section 4.1 and the adaptive θ (after one step) introduced in Section 4.2 for a simulated example with block-interchangeable dependence structure, in terms of the within-group error, deviation from the true intercept term β_B, deviation from the true direction $\boldsymbol{\omega}_B$, and the RankComp from $\boldsymbol{\omega}_B$. Intermediate FLAMEs provide improvements over DWD for the first two measures and over SVM for the last two measures.

total sample size is 240 and the imbalance factor m is 3 (moderate imbalanced).

In Figure 7, we compare the within-group error, $|\beta - \beta_B|$, $\angle(\boldsymbol{\omega}, \boldsymbol{\omega}_B)$, and $\text{RankComp}(\boldsymbol{\omega}, \boldsymbol{\omega}_B)$. We see that these intermediate FLAMEs provide improvements over DWD for the first two measures and over SVM for the last two measures. For relatively small d, the equal-trade-off θ value is very similar to the adaptive θ. For large d, the adaptive θ is closer to DWD than the equal-trade-off θ. For very large d, all four machines encounter difficulty in classification.

7　Real Data Application

In this section we demonstrate the performance of FLAME on a real example: the Human Lung Carcinomas Microarray Dataset, which has been analyzed earlier in Bhattacharjee

et al. ([2001]).

The Human Lung Carcinomas Dataset contains six classes: adenocarcinoma, squamous, pulmonary carcinoid, colon, normal and small cell carcinoma, with sample sizes of 128, 21, 20, 13, 17 and 6 respectively. Liu et al. ([2008]) used this data as a test set to demonstrate their proposed significance analysis of clustering approach. We combine the first two subclasses and the last four subclasses to form the positive and negative classes respectively. The sample sizes are 149 and 56 with imbalance factor $m = 2.66$. The original data contain 12,625 genes. We first filter genes using the ratio of the sample standard deviation and sample mean of each gene and keep 2,530 of them with large ratios (Dudoit et al., [2002], Liu et al., [2008]).

We conduct a five-fold cross-validation (CV) to evaluate the within-group error for the two classes over 100 random splits. The RankComp measure is calculated based on the full data set instead of on the samples in a single fold. For each replication, we find a suggested value for θ that leads to the same normalized value for RankComp and within-group error over a grid of θ values. The adaptive value for θ (after one step) is calculated based on the DWD direction using all the samples in the data set.

The average within-group errors and the RankComp measures for DWD, SVM, and the two intermediate FLAMEs (using the adaptive θ and the suggested *equal-trade-off* θ) are shown in Figure [8].

Note that as we do not know the true rank of importance of genes for this real data application (no Bayes rule direction or the rank of importance that it implies), we use the DWD rank as a surrogate to the truth, since simulation examples show that on average its estimated direction is the closest to the Bayes rule direction. Therefore, the RankComp measure for DWD is 0.

This experiment does show that FLAME opens a new dimension of improving both the classification performance and the interpretative ability of the classifier. The compromise of the FLAME classifier, with the suggested equal-trade-off θ, in terms of the within-group error is very small compared to the improvement obtained in terms of the direction. We would
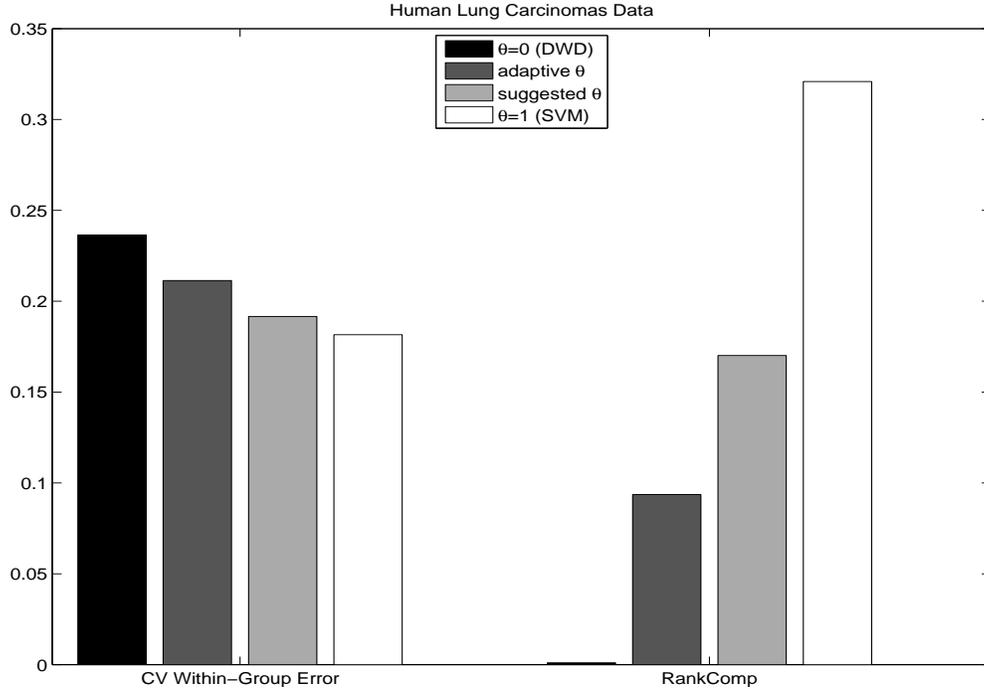
Figure 8: The average cross-validation within-group errors and the RankComp measures for DWD, FLAME with the adaptive and equal-trade-off θ parameters and SVM for the Human Lung Carcinomas Dataset.

not recommend using either DWD or SVM for conducting prediction and interpretation simultaneously due to their bad performance for at least one criterion. A FLAME classifier with an appropriate parameter may be more suitable for practical use.

8   Conclusion and Discussion

In this paper, we thoroughly investigate SVM and DWD on their performance when applied to the HDLSS and imbalanced data. A novel family of binary classifiers called FLAME is proposed, where SVM and DWD are the two ends of the spectrum. On the DWD end, the estimation of the intercept term is deteriorated while it provides better estimation of the direction vector, and thus better handle the HDLSS data. On the hand, SVM is good at estimating the intercept term but not the direction and is subject to overfitting, and thus is more suitable for imbalanced data but not HDLSS data.

We conduct extensive study of the asymptotic properties of the FLAME family in three different flavors, the 'd fixed, $n \to \infty$' asymptotics (Fisher consistency), the 'd and n_+ fixed, $n_- \to \infty$' asymptotics (extremely imbalanced data), and the 'n fixed, $d \to \infty$' asymptotics (the HDLSS asymptotics). These results explain the performance we have seen in the simulations and suggest that with a smart choice of θ, FLAME can properly handle both the HDLSS data and the imbalanced data, by improving the estimations of the direction and the intercept term.

The FLAME family can be immediately extended to multi-class classification, as was done for SVM and DWD such as in Weston and Watkins (1999), Crammer and Singer (2000), Lee et al. (2004) or Huang et al. (2012). Another natural extension is variable selection for FLAME.

The FLAME machines generalize the concepts of support vectors. In SVM, support vectors are referred to vectors that sit on or fall into the two hyperplanes corresponding to $u \leq 1$ (or $u \leq 1/\sqrt{C}$ for the modified version of Hinge loss (2)). In SVM, only support vectors have impacts on the final solution. DWD is the other extreme case where all the data vectors have some impacts. In the presence of imbalanced sample size, the fact that all the data vectors influence the solution cause the optimization to ignore the minority class. The FLAME with $0 < \theta < 1$ is somewhere in the middle. For FLAME, part of the data vectors, more than the support vectors, but fewer than all the vectors, have impacts. Smart choice of θ means that one needs to include as many vectors, and as balanced influential samples, as possible. More vectors usually lead to mitigated overfitting, and balanced sample size of the influential vectors from two classes means that the sensitivity issue of the intercept term can be alleviated.

The authors are aware that it is possible to implement a two step procedure to conduct binary linear classification. In the first step, a good direction is found, probably in the fashion of DWD; in the second step, a fine intercept is chosen by borrowing idea of SVM. However, the theoretical properties of such a procedure are known and will be left as a future

research direction.

The choice of θ usually depends on the nature of the data and the scientific context. If the users prefer better classification performance over reasonable discrimination direction for interpretation of the data, θ may be chosen to be closer to 1. If the right direction is the first priority, then θ should be chosen to be closer to 0. Note that, under some circumstances, the primary goal is to obtain a direction vector which can provide a score $\boldsymbol{x}^T\boldsymbol{\omega}$ for each observation for further use, and the intercept parameter β is of no use at all. For example, some users may use a receiver operating characteristic (ROC) curve as a graphical tool to evaluate classification performance over different β value instead of using a single β value given by the classifier. In this case, a FLAME machine close to the DWD method may be ideal.

Appendix

Derivation of the modified Hinge loss

Note that the original SVM formulation is $\underset{\tilde{\boldsymbol{\omega}},\tilde{\beta}}{\arg\min} \sum \left(1 - y_i\tilde{f}(\boldsymbol{x}_i)\right)_+$, s.t. $\|\tilde{\boldsymbol{\omega}}\|^2 \leq C$, where $\tilde{f}(\boldsymbol{x}) = \boldsymbol{x}^T\tilde{\boldsymbol{\omega}} + \tilde{\beta}$. Here the coefficient vector $\tilde{\boldsymbol{\omega}}$ does not have unit norm. We let $\boldsymbol{\omega} = \tilde{\boldsymbol{\omega}}/\sqrt{C}$, $\beta = \tilde{\beta}/\sqrt{C}$ and $f = \tilde{f}/\sqrt{C}$. Thus SVM solution is given by $\underset{\boldsymbol{\omega},\beta}{\arg\min} \sum \left(1 - \sqrt{C}y_if(\boldsymbol{x}_i)\right)_+$, s.t. $\|\boldsymbol{\omega}\|^2 \leq 1$, or equivalently, $\underset{\boldsymbol{\omega},\beta}{\arg\min} \sum \left(\sqrt{C} - Cy_if(\boldsymbol{x}_i)\right)_+$, s.t. $\|\boldsymbol{\omega}\|^2 \leq 1$.

Acknowledgements

References

Ahn, J. and Marron, J. (2010), "The maximal data piling direction for discrimination," *Biometrika*, 97, 254–259.

Ahn, J., Marron, J., Muller, K., and Chi, Y. (2007), "The high-dimension, low-sample-size geometric representation holds under mild conditions," *Biometrika*, 94, 760–766.

Bhattacharjee, A., Richards, W. G., Staunton, J., Li, C., Monti, S., Vasa, P., Ladd, C., Beheshti, J., Bueno, R., Gillette, M., et al. (2001), "Classification of human lung carcinomas by mRNA expression profiling reveals distinct adenocarcinoma subclasses," *Proceedings of the National Academy of Sciences*, 98, 13790–13795.

Chawla, N., Japkowicz, N., and Kotcz, A. (2004), "Editorial: special issue on learning from imbalanced data sets," *ACM SIGKDD Explorations Newsletter*, 6, 1–6.

Cortes, C. and Vapnik, V. (1995), "Support-vector networks," *Machine learning*, 20, 273–297.

Crammer, K. and Singer, Y. (2000), "On the Learnability and Design of Output Codes for Multiclass Problems," in *In Proceedings of the Thirteenth Annual Conference on Computational Learning Theory*.

Cristianini, N. and Shawe-Taylor, J. (2000), *An introduction to Support Vector Machines: and other kernel-based learning methods*, Cambridge University Press.

Duda, R., Hart, P., and Stork, D. (2001), *Pattern classification*, Wiley.

Dudoit, S., Fridlyand, J., and Speed, T. (2002), "Comparison of discrimination methods for the classification of tumors using gene expression data," *Journal of the American statistical association*, 97, 77–87.

Hall, P., Marron, J. S., and Neeman, A. (2005), "Geometric representation of high dimension, low sample size data," *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 67, 427–444.

Hastie, T., Tibshirani, R., and Friedman, J. (2009), *The elements of statistical learning: Data mining, inference, and prediction (second edition)*, Springer.

Huang, H., Liu, Y., Du, Y., Perou, C. M., Hayes, D. N., Todd, M. J., and Marron, J. (2012), "Multiclass Distance Weighted Discrimination," *Journal of Computational and Graphical Statistics*.

Lee, Y., Lin, Y., and Wahba, G. (2004), "Multicategory support vector machines," *Journal of the American Statistical Association*, 99, 67–81.

Lin, Y. (2004), "A note on margin-based loss functions in classification," *Statistics & probability letters*, 68, 73–82.

Liu, Y., Hayes, D., Nobel, A., and Marron, J. (2008), "Statistical significance of clustering for high-dimension, low-sample size data," *Journal of the American Statistical Association*, 103, 1281–1293.

Marron, J., Todd, M., and Ahn, J. (2007), "Distance-weighted discrimination," *Journal of the American Statistical Association*, 102, 1267–1271.

Owen, A. (2007), "Infinitely imbalanced logistic regression," *The Journal of Machine Learning Research*, 8, 761–773.

Qiao, X. and Liu, Y. (2009), "Adaptive weighted learning for unbalanced multicategory classification," *Biometrics*, 65, 159–168.

Qiao, X., Zhang, H., Liu, Y., Todd, M., and Marron, J. (2010), "Weighted distance weighted discrimination and its asymptotic properties," *Journal of the American Statistical Association*, 105, 401–414.

Smola, A. J., Bartlett, P. L., Schölkopf, B., and Schuurmans, D. (2000), *Advances in large margin classifiers*, vol. 1, MIT press Cambridge, MA.

Toh, K., Todd, M., and Tütüncü, R. (1999), "SDPT 3- a MATLAB software package for semidefinite programming, version 1.3," *Optimization Methods and Software*, 11, 545–581.

Tütüncü, R., Toh, K., and Todd, M. (2003), "Solving semidefinite-quadratic-linear programs using SDPT3," *Mathematical programming*, 95, 189–217.

Vapnik, V. (1998), *Statistical learning theory*, Wiley.

Weston, J. and Watkins, C. (1999), "Support vector machines for multi-class pattern recognition," in *European Symposium on Artificial Neural Networks*, pp. 219–224.

Zhang, L. and Lin, X. (2011), "Some considerations of classification for high dimension low-sample size data," *Statistical Methods in Medical Research*.

Online supplementary materials

This document provides some additional details for the main paper, *Flexible High-dimensional Classification Machines and Their Asymptotic Properties*. This document includes how we implement the FLAME machine with pre-defined θ, detailed proofs for several theorems and propositions, and figures for additional simulations.

Implementation

In order to implement the FLAME algorithm, we introduce several new notations. Let S_{d+1} be a second order cone in the $d+1$ dimensional space, $S_{d+1} = \left\{ (t_0, t_1, \cdots, t_d)' : t_0 \geq \sqrt{\sum_{i=1}^d t_i^2} \right\}$. Note that r_i and $1/r_i$ can be substituted by three axillary variables ρ_i, σ_i and τ_i which satisfy $\rho_i + \sigma_i = r_i$, $\rho_i - \sigma_i = 1/r_i$, and $\tau_i = 1$. Then $\rho_i^2 = \sigma_i^2 + \tau_i^2$, and thus $(\rho_i, \sigma_i, \tau_i) \in S_3$. Let $w = 1$, then $(w; \boldsymbol{\omega}) \in S_{d+1}$ since $\|\boldsymbol{\omega}\| \leq 1$. Let $\eta_i \geq 0$, $\varphi_i \geq 0$, where φ_i and η_i can be viewed as the positive and negative parts of $\left(\frac{1}{r_i} + C\xi_i - \theta\sqrt{C} \right)$, i.e., $\varphi_i - \eta_i = \left(\frac{1}{r_i} + C\xi_i - \theta\sqrt{C} \right)$. With the reparameterization above, FLAME can be viewed as the following optimization problem:

$$\min_{\beta, w, \boldsymbol{\omega}, \rho_i, \sigma_i, \tau_i, \xi_i, \eta_i, \varphi_i} \sum_{i=1}^n \varphi_i$$

$$\text{s.t.} \quad y_i(\boldsymbol{x}_i^T \boldsymbol{\omega} + \beta) + \xi_i - \rho_i - \sigma_i = 0$$

$$\rho_i - \sigma_i + C\xi_i - \theta\sqrt{C} + \eta_i - \varphi_i = 0$$

$$w = 1$$

$$\tau_i = 1$$

and $(w; \boldsymbol{\omega}) \in S_{d+1}, (\rho_i, \sigma_i, \tau_i)' \in S_3, \xi_i \geq 0, \eta_i \geq 0, \varphi_i \geq 0.$

Therefore, all the constraints can be converted to linear forms, all the variables are either nonnegative, free, or in second order cones, and the objective function is linear. Such problem is called Second Order Cone Programming (SOCP), and can be efficiently solved by softwares such as SDPT3 (Toh et al., 1999, Tütüncü et al., 2003).

Proof to Theorem 1

It suffices to show that $s(\boldsymbol{\omega}_k, \beta_k, \theta_k) \geq s(\boldsymbol{\omega}_{k+1}, \beta_{k+1}, \theta_{k+1})$. First, $s(\boldsymbol{\omega}_k, \beta_k, \theta_k) \geq s(\boldsymbol{\omega}_k, \beta_k, \theta_{k+1})$ due to the definition of θ_{k+1} and that $\theta_k \leq \theta_{k+1}$. Then $s(\boldsymbol{\omega}_k, \beta_k, \theta_{k+1}) \geq s(\boldsymbol{\omega}_{k+1}, \beta_{k+1}, \theta_{k+1})$ since $\boldsymbol{\omega}_{k+1}$ and β_{k+1} minimize $s(\boldsymbol{\omega}, \beta, \theta_{k+1})$. $\square$

Proof to Proposition 2

For any $\boldsymbol{x}$, denote $p(\boldsymbol{x}) = \mathrm{P}(Y = +1 \mid \boldsymbol{X} = \boldsymbol{x})$. The conditional risk is $R(f) \equiv E[L(Yf(\boldsymbol{X}), \theta) \mid \boldsymbol{X} = \boldsymbol{x}] = L(f(\boldsymbol{x}), \theta)p(\boldsymbol{x}) + L(-f(\boldsymbol{x}), \theta)(1 - p(\boldsymbol{x}))$. For simplicity, we write $L(f(\boldsymbol{x}), \theta)$ as $L(f)$. Thereby, $R(f) = L(f)p(\boldsymbol{x}) + L(-f)(1 - p(\boldsymbol{x}))$.

We can see that for fixed $p(\boldsymbol{x}) \in (0, 1)$, $R(f)$ is continuous and differentiable everywhere and convex. Thus we find $f^*(\boldsymbol{x})$ by solving $R'(f) = 0$, where $R'(f) = L'(f)p(\boldsymbol{x}) + [L(-f)]'(1 - p(\boldsymbol{x}))$. The FLAME loss is

$$
L(f) = \begin{cases} (2 - \theta)\sqrt{C} - Cf & \text{if } f \leq \frac{1}{\sqrt{C}} \\ \frac{1}{f} - \theta\sqrt{C} & \text{if } \frac{1}{\sqrt{C}} < f \leq \frac{1}{\theta\sqrt{C}} \\ 0 & \text{otherwise.} \end{cases}
$$

So direct calculation gives us that

$$
L'(f) = \begin{cases} -C & \text{if } f \leq \frac{1}{\sqrt{C}} \\ -\frac{1}{f^2} & \text{if } \frac{1}{\sqrt{C}} < f \leq \frac{1}{\theta\sqrt{C}} \\ 0 & \text{otherwise,} \end{cases} \quad \text{and} \quad [L(-f)]' = \begin{cases} C & \text{if } f \geq -\frac{1}{\sqrt{C}} \\ \frac{1}{f^2} & \text{if } -\frac{1}{\theta\sqrt{C}} < f \leq -\frac{1}{\sqrt{C}} \\ 0 & \text{otherwise.} \end{cases}
$$

Finally, if $\frac{1-p(\boldsymbol{x})}{p(\boldsymbol{x})} \neq 1$, the solution to $R'(f) = 0$ has to be either $-\frac{1}{\theta\sqrt{C}} < f \leq -\frac{1}{\sqrt{C}}$ or $\frac{1}{\sqrt{C}} < f \leq \frac{1}{\theta\sqrt{C}}$. In the former case, $f^* = -\frac{1}{\sqrt{C}}\sqrt{\frac{1-p(\boldsymbol{x})}{p(\boldsymbol{x})}}$ if $\frac{1-p(\boldsymbol{x})}{p(\boldsymbol{x})} > 1$. In the latter case, $f^* = +\frac{1}{\sqrt{C}}\sqrt{\frac{1-p(\boldsymbol{x})}{p(\boldsymbol{x})}}$ if $\frac{1-p(\boldsymbol{x})}{p(\boldsymbol{x})} < 1$. If $\frac{1-p(\boldsymbol{x})}{p(\boldsymbol{x})} = 1$, $R'(f^*) = 0$ for any $f^* \in [-\frac{1}{\sqrt{C}}, \frac{1}{\sqrt{C}}]$.

Therefore, f^* satisfies $\mathrm{sign}(f^*) = \mathrm{sign}(p(\boldsymbol{x}) - 0.5) = \mathrm{sign}(2p(\boldsymbol{x}) - 1) = \mathrm{sign}(p(\boldsymbol{x}) - (1 - p(\boldsymbol{x}))) = \mathrm{sign}(\frac{p(\boldsymbol{x})}{1-p(\boldsymbol{x})} - 1)$. $\square$

Proof to Theorem 3

Since there are infinitely many negative class samples, it is reasonable to assume that the classification boundary is pushed closer to the minority positive class, and therefore, the functional margin $u_i = y_i f(\boldsymbol{x}_i) = f(\boldsymbol{x}_i)$ for the ith vector from the minority positive class is small and its DWD loss is $2\sqrt{C} - Cu_i = 2\sqrt{C} - Cf(\boldsymbol{x}_i)$. Similarly, the DWD loss for the jth vector from the majority negative class is $1/[y_j f(\boldsymbol{x}_j)] = -1/f(\boldsymbol{x}_j)$. The objective function for DWD is therefore equivalent to $\dfrac{1}{n_+ + n_-}\left\{\displaystyle\sum_{i=1}^{n_+}[2\sqrt{C} - C(\boldsymbol{x}_i^T\boldsymbol{\omega} + \beta)] - \sum_{j=1}^{n_-}\dfrac{1}{\boldsymbol{x}_j^T\boldsymbol{\omega} + \beta}\right\} + \dfrac{\lambda}{2}\|\boldsymbol{\omega}\|^2$.

The second term inside the curly bracket above can be approximated by $n_- \displaystyle\int \dfrac{1}{\boldsymbol{x}^T\boldsymbol{\omega} + \beta} dF_-(\boldsymbol{x})$ where $F_-(\cdot)$ is the conditional cumulative distribution function for the negative class. The objective function is therefore

$$l_D = \frac{1}{n_+ + n_-}\left\{\sum_{i=1}^{n_+}[2\sqrt{C} - C(\boldsymbol{x}_i^T\boldsymbol{\omega} + \beta)] - n_-\int \frac{1}{\boldsymbol{x}^T\boldsymbol{\omega} + \beta}dF_-(\boldsymbol{x})\right\} + \frac{\lambda}{2}\|\boldsymbol{\omega}\|^2$$

Before we continue, we need the definition that a distribution has a point surrounded (Owen, 2007).

DEFINITION 1. The distribution F on $\mathbb{R}^d$ has the point $\boldsymbol{x}_*$ surrounded if

$$\int_{(\boldsymbol{x}-\boldsymbol{x}_*)'\boldsymbol{\omega}>\epsilon} dF(\boldsymbol{x}) > \delta,$$

for some $\delta > 0$, some $\epsilon > 0$ and all $\boldsymbol{\omega} \in \mathbb{R}^d$ with $\|\boldsymbol{\omega}\| = 1$. Consequentially, if F has $\boldsymbol{x}_*$ surrounded, then there exist γ satisfying

$$\inf_{\|\boldsymbol{\omega}\|=1}\int_{(\boldsymbol{x}-\boldsymbol{x}_*)'\boldsymbol{\omega}>0} dF(\boldsymbol{x}) > \gamma \geq 0. \tag{S.1}$$

We observe that

$$\frac{\partial l_D}{\partial \beta} = \frac{1}{n_+ + n_-}[-n_+C + n_-\int(\boldsymbol{x}^T\boldsymbol{\omega} + \beta)^{-2}dF_-(\boldsymbol{x})]$$

$$\geq \frac{1}{n_+ + n_-}[-n_+C + n_-\int_{(\boldsymbol{x}-\overline{\boldsymbol{x}}_+)'\boldsymbol{\omega}\geq 0}((\boldsymbol{x}-\overline{\boldsymbol{x}}_+)^T\boldsymbol{\omega} + \overline{\boldsymbol{x}}_+^T\boldsymbol{\omega} + \beta)^{-2}dF_-(\boldsymbol{x})]$$

$$\geq \frac{1}{n_+ + n_-}[-n_+C + n_-\int_{\boldsymbol{x}^T\boldsymbol{\omega}\geq 0}(\overline{\boldsymbol{x}}_+^T\boldsymbol{\omega} + \beta)^{-2}dF_-(\boldsymbol{x})]$$

$$\geq \frac{1}{n_+ + n_-}[-n_+C + n_-\frac{\gamma}{(\overline{\boldsymbol{x}}_+^T\boldsymbol{\omega} + \beta)^2}]$$

Now suppose that $-\sqrt{\dfrac{n_-\gamma}{n_+C}} < \overline{\boldsymbol{x}}_+^T\boldsymbol{\omega} + \beta < 0$, then $\dfrac{n_-\gamma}{(\overline{\boldsymbol{x}}_+^T\boldsymbol{\omega}+\beta)^2} > n_+C$ and $\partial l_D/\partial\beta > 0$. Given

the fact that l_D is a strictly convex function, the minimizer $\widehat{\beta} < -\sqrt{\frac{n_- \gamma}{n_+ C}} - \overline{\boldsymbol{x}}_+^T \boldsymbol{\omega}$. $\square$

Proof to Theorem 4

Again, with the imbalance assumption we assume that the functional margins for the minority positive class are always greater than 0. Note that the penalized empirical loss for the FLAME machine is approximated by

$$
l_F = \frac{1}{n_+ + n_-} \left\{ \sum_i^{n_+} \left[2\sqrt{C} - C(\boldsymbol{x}_i^T \boldsymbol{\omega} + \beta) - \theta\sqrt{C} \right] + \right.
$$
$$
\left. n_- \int \left(-\frac{1}{\boldsymbol{x}^T \boldsymbol{\omega} + \beta} - \theta\sqrt{C} \right)_+ dF_-(\boldsymbol{x}) \right\} + \frac{\lambda}{2} \|\boldsymbol{\omega}\|^2
$$

Let $g_j^* = -(\boldsymbol{x}_j^T \boldsymbol{\omega}^* + \beta^*)$, $j = 1, 2, \cdots n_-$ be the functional margins for the negative class. Because $\frac{1}{g_{(n_+)}^* \sqrt{C}} = \theta^*$, that is, the reduced loss for the jth sample is greater than or equal to 0, $\frac{1}{g_{(n_+)}^*} - \theta^* \sqrt{C} = 0$, we observe that $1/g_{(n_+)}^*$ is the n_+-th greatest among all the function margins of the negative class $1/g_j^* = -1/(\boldsymbol{x}_j^T \boldsymbol{\omega}^* + \beta^*)$. Thus there are at most n_+ samples whose reduced losses that are ≥ 0. Assume that there are $n^o \leq n_+$ such samples.

For a random sample $(\boldsymbol{X}, Y)$ from the negative class, let E be the event that $(Y(\boldsymbol{X}^T \boldsymbol{\omega}^* + \beta^*))^{-1} \geq \theta^* \sqrt{C}$. From the argument above, $\mathrm{P}(E)$ is approximately n^o/n_-.

Then the integration $\int \left(-\frac{1}{\boldsymbol{x}^T \boldsymbol{\omega} + \beta} - \theta\sqrt{C} \right)_+ dF_-(\boldsymbol{x})$ equals

$$
\mathbb{E}\left[(-\frac{1}{\boldsymbol{x}^T \boldsymbol{\omega} + \beta} - \theta\sqrt{C})_+ \mid \overline{E} \right] \mathrm{P}(\overline{E}) + \mathbb{E}\left[(-\frac{1}{\boldsymbol{x}^T \boldsymbol{\omega} + \beta} - \theta\sqrt{C})_+ \mid E \right] \mathrm{P}(E)
$$
$$
\approx \mathbb{E}\left[0 \mid \overline{E} \right] (1 - \frac{n^o}{n_-}) + \mathbb{E}\left[(-\frac{1}{\boldsymbol{x}^T \boldsymbol{\omega} + \beta} - \theta\sqrt{C}) \mid E \right] \frac{n^o}{n_-}
$$
$$
= \mathbb{E}\left[(-\frac{1}{\boldsymbol{x}^T \boldsymbol{\omega} + \beta} - \theta\sqrt{C}) \mid E \right] \frac{n^o}{n_-}
$$

We then have

$$
l_F = \frac{1}{n_+ + n_-} \left\{ \sum_i^{n_+} \left[2\sqrt{C} - C(\boldsymbol{x}_i^T \boldsymbol{\omega} + \beta) - \theta\sqrt{C} \right] + \right.
$$
$$
\left. n^o \int \left(-\frac{1}{\boldsymbol{x}^T \boldsymbol{\omega} + \beta} - \theta\sqrt{C} \right) dF_-(\boldsymbol{x} \mid E) \right\} + \frac{\lambda}{2} \|\boldsymbol{\omega}\|^2
$$

Here, $dF_-(\boldsymbol{x} \mid E)$ is the conditional distribution function of $\boldsymbol{X}$ for the negative class given event E.

Setting $\partial l_F / \partial \beta = 0 = \dfrac{1}{n_+ + n_-} \left\{ -Cn_+ + n^o \displaystyle\int (\boldsymbol{x}^T\boldsymbol{\omega}^* + \beta^*)^{-2} dF_-(\boldsymbol{x} \mid E) \right\}$, we have

$$\int (\boldsymbol{x}^T\boldsymbol{\omega}^* + \beta^*)^{-2} dF_-(\boldsymbol{x} \mid E) = C\dfrac{n_+}{n^o}.$$

Setting $\partial l_F / \partial \boldsymbol{\omega} = \boldsymbol{0} = \dfrac{1}{n_+ + n_-} \left\{ -Cn_+\overline{\boldsymbol{x}}_+ + n^o \displaystyle\int (\boldsymbol{x}^T\boldsymbol{\omega}^* + \beta^*)^{-2}\boldsymbol{x} dF_-(\boldsymbol{x} \mid E) \right\} + \lambda\boldsymbol{\omega}^*,$

we have $\displaystyle\int (\boldsymbol{x}^T\boldsymbol{\omega}^* + \beta^*)^{-2}\boldsymbol{x} dF_-(\boldsymbol{x} \mid E) = -\dfrac{n_+ + n_-}{n^o}\lambda\boldsymbol{\omega}^* + C\dfrac{n_+}{n^o}\overline{\boldsymbol{x}}_+.$ And furthermore,

$$\dfrac{\int (\boldsymbol{x}^T\boldsymbol{\omega}^* + \beta^*)^{-2}\boldsymbol{x} dF_-(\boldsymbol{x} \mid E)}{\int (\boldsymbol{x}^T\boldsymbol{\omega}^* + \beta^*)^{-2} dF_-(\boldsymbol{x} \mid E)} = -\dfrac{n_+ + n_-}{n_+}\dfrac{\lambda}{C}\boldsymbol{\omega}^* + \overline{\boldsymbol{x}}_+ = -(1 + m)\dfrac{\lambda}{C}\boldsymbol{\omega}^* + \overline{\boldsymbol{x}}_+.$$

That is,

$$\boldsymbol{\omega}^* = \dfrac{C}{(1 + m)\lambda} \left[\overline{\boldsymbol{x}}_+ - \dfrac{\int (\boldsymbol{x}^T\boldsymbol{\omega}^* + \beta^*)^{-2}\boldsymbol{x} dF_-(\boldsymbol{x} \mid E)}{\int (\boldsymbol{x}^T\boldsymbol{\omega}^* + \beta^*)^{-2} dF_-(\boldsymbol{x} \mid E)} \right]$$

$\square$

Proof to Theorem 5

For simplicity we use the original SVM formulation with the Hinge loss function instead of the FLAME formulation. The objective function for SVM is equivalent to

$$l_S = \dfrac{1}{n_+ + n_-} \left\{ \sum_{i=1}^{n_+}[1 - (\boldsymbol{x}_i^T\boldsymbol{\omega} + \beta)] + n_- \int \left[1 + (\boldsymbol{x}^T\boldsymbol{\omega} + \beta) \right]_+ dF_-(\boldsymbol{x}) \right\} + \dfrac{\lambda}{2}\|\boldsymbol{\omega}\|^2$$

$$= \dfrac{1}{n_+ + n_-} \left\{ \sum_{i=1}^{n_+}[1 - (\boldsymbol{x}_i^T\boldsymbol{\omega} + \beta)] + \right.$$

$$\left. n_- \int \left[1 + (\boldsymbol{x}^T\boldsymbol{\omega} + \beta) \right] \mathbb{1}_{\{1 + \boldsymbol{x}^T\boldsymbol{\omega} + \beta > 0\}} dF_-(\boldsymbol{x}) \right\} + \dfrac{\lambda}{2}\|\boldsymbol{\omega}\|^2$$

Setting $\partial l_S / \partial \beta = 0$, we have

$$\dfrac{\partial l_S}{\partial \beta} = \dfrac{1}{n_+ + n_-} \left\{ -n_+ + n_-\mathrm{P}(G; \boldsymbol{\omega}, \beta) + n_- \int (1 + \boldsymbol{x}^T\boldsymbol{\omega} + \beta)\delta\left(1 + \boldsymbol{x}^T\boldsymbol{\omega} + \beta\right) dF_-(\boldsymbol{x}) \right\}$$

$$= \dfrac{1}{n_+ + n_-} \left\{ -n_+ + n_-\mathrm{P}(G; \boldsymbol{\omega}, \beta) \right\} = 0,$$

where $\delta(\cdot)$ is the Dirac delta function.

$$\Rightarrow \mathrm{P}(G; \widehat{\boldsymbol{\omega}}, \widehat{\beta}) = \dfrac{n_+}{n_-} = \dfrac{1}{m}.$$

Moreover,

$$
\begin{aligned}
\frac{\partial l_S}{\partial \boldsymbol{\omega}} =& \frac{1}{n_+ + n_-} \left\{ -\sum_{i=1}^{n_+} \boldsymbol{x}_i + n_- \int \boldsymbol{x} \mathbb{1}_{\{1+\boldsymbol{x}^T\boldsymbol{\omega}+\beta>0\}} dF_-(\boldsymbol{x}) \right. \\
& \left. + n_- \int \left[1 + (\boldsymbol{x}^T\boldsymbol{\omega}+\beta) \right] \delta\left(1 + \boldsymbol{x}^T\boldsymbol{\omega}+\beta \right) \boldsymbol{x} dF_-(\boldsymbol{x}) \right\} + \lambda\boldsymbol{\omega} \\
=& \frac{1}{n_+ + n_-} \left\{ -n_+\overline{\boldsymbol{x}}_+ + n_- \int \boldsymbol{x} \mathbb{1}_{\{1+\boldsymbol{x}^T\boldsymbol{\omega}+\beta>0\}} dF_-(\boldsymbol{x}) \right\} + \lambda\boldsymbol{\omega} \\
\Rightarrow \widehat{\boldsymbol{\omega}} =& \frac{1}{(n_+ + n_-)\lambda} \left\{ n_+\overline{\boldsymbol{x}}_+ - n_- \int \boldsymbol{x} \mathbb{1}_{\{1+\boldsymbol{x}^T\widehat{\boldsymbol{\omega}}+\widehat{\beta}>0\}} dF_-(\boldsymbol{x}) \right\} \\
=& \frac{1}{(n_+ + n_-)\lambda} \left\{ n_+\overline{\boldsymbol{x}}_+ - n_- \int \boldsymbol{x} dF_-(\boldsymbol{x}|G) \mathrm{P}(G; \widehat{\boldsymbol{\omega}}, \widehat{\beta}) \right\} \\
\approx& \frac{n_+}{(n_+ + n_-)\lambda} \left\{ \overline{\boldsymbol{x}}_+ - \int \boldsymbol{x} dF_-(\boldsymbol{x}|G) \right\} \\
=& \frac{1}{(1+m)\lambda} \left\{ \overline{\boldsymbol{x}}_+ - \int \boldsymbol{x} dF_-(\boldsymbol{x}|G) \right\}.
\end{aligned}
$$

$\square$

Classification boundaries for HDLSS data

The geometric representation in Hall et al. (2005) leads to some theoretical properties of several binary classifiers. In particular, as $d \to \infty$, the positive class and negative class converge to two $(n_+ - 1)$ and $(n_- - 1)$ simplices with random rotation. Note that the (normalized) pairwise distances between observations within each class are the same, and the (normalized) distances between any two observations from two different classes are the same as well. The geometric representation for SVM and DWD in Hall et al. (2005) is summarized as the follows.

1. **SVM**: It was shown that the linear SVM hyperplane projected to the $(N - 1)$-dimensional subspace that is generated by the N data vectors is given asymptotically by the unique $(N - 2)$-dimensional hyperplane that bisects each of the edges of length l in the N-polyhedron formed by the N data vectors. There are $n_+ \times n_-$ such edges. Let O^+ be the centroid of the $(n_+ - 1)$-simplex $\mathcal{X}^+(d)$ and O^- the centroid of the $(n_- - 1)$-simplex $\mathcal{X}^-(d)$. It can be further shown that the SVM hyperplane bisects the

line segment between O^+ and O^-.

2. **DWD**: The case of DWD is a little different, especially in the case where $n_+ \ll n_-$ (or $m \gg 1$), which is our main focus here. We assume that the DWD hyperplane intersects O^+O^- at point P. It can be shown that the two simplices, the DWD hyperplane, and the SVM hyperplane, are all orthogonal to O^+O^-. Thus all the vertices in the simplex $\mathcal{X}^+$ are equally distanced from the DWD hyperplane. Such distance is denoted by a. Similarly, all the vertices in the simplex $\mathcal{X}^-$ are equally distanced from the DWD hyperplane by β. The general version DWD hyperplane minimizes the sum of the reciprocals of the distances of data vectors to the hyperplane, $(n_+/a + n_-/b)$, with the constraint that $a + b$ equals to a constant (determined by μ, σ, τ, n_+, n_-, and d). A simple calculus practice reveals that $a/b = (n_+/n_-)^{1/2}$.

For the General FLAME case, we need to learn how the hyperplane moves from the point determined by $a/b = (n_+/n_-)^{1/2}$ on O^+O^- to the midpoint of O^+O^- as θ grows from 0 (DWD) to 1 (SVM). First, we consider the general version of FLAME which seeks to minimize the sum of losses for all data points, $\sum \left(1/u - \theta\sqrt{C}\right)_+$, where the functional margin u is either a or β for samples from the positive or the negative classes respectively. When $\theta = 0$, the FLAME hyperplane is determined by $a/b = (n_+/n_-)^{1/2} = m^{-1/2} < 1$. Thus $b > a$, that is, the hyperplane is closer to the minority class. We renamed them as a^0 and b^0 where the superscript "0" represents the value of θ.

When $\theta > 0$ but smaller than $1/(b^0\sqrt{C})$, then the hyperplane does not move. This is because that the new loss for each data vector becomes $1/a^0 - \theta\sqrt{C}$ or $1/b^0 - \theta\sqrt{C}$ since both are greater than 0. The additional term "$-\theta\sqrt{C}$" does not change the minimizer and thus $a^{\theta_1}/b^{\theta_1} = (n_+/n_-)^{1/2}$ does remain unchanged.

If we keep increasing θ so that it becomes greater than $1/(b^0\sqrt{C})$, then if the hyperplane does not move, then the loss for the majority class becomes 0. In this case, there is space for improvement: the hyperplane would move gradually towards the majority class, because this can make the loss on the minority class smaller while keeping the loss on the majority

zero. The FLAME hyperplane is determined by $b = 1/(\theta\sqrt{C})$.

Finally, as θ increases, the distance a increases and the distance β decreases, until at a point where $a = b$, and both $1/a - \theta\sqrt{C} = 1/b - \theta\sqrt{C} < 0$. After this point, further increase of θ will not change the position of the FLAME hyperplane which will remain at the midpoint of O^+O^-.
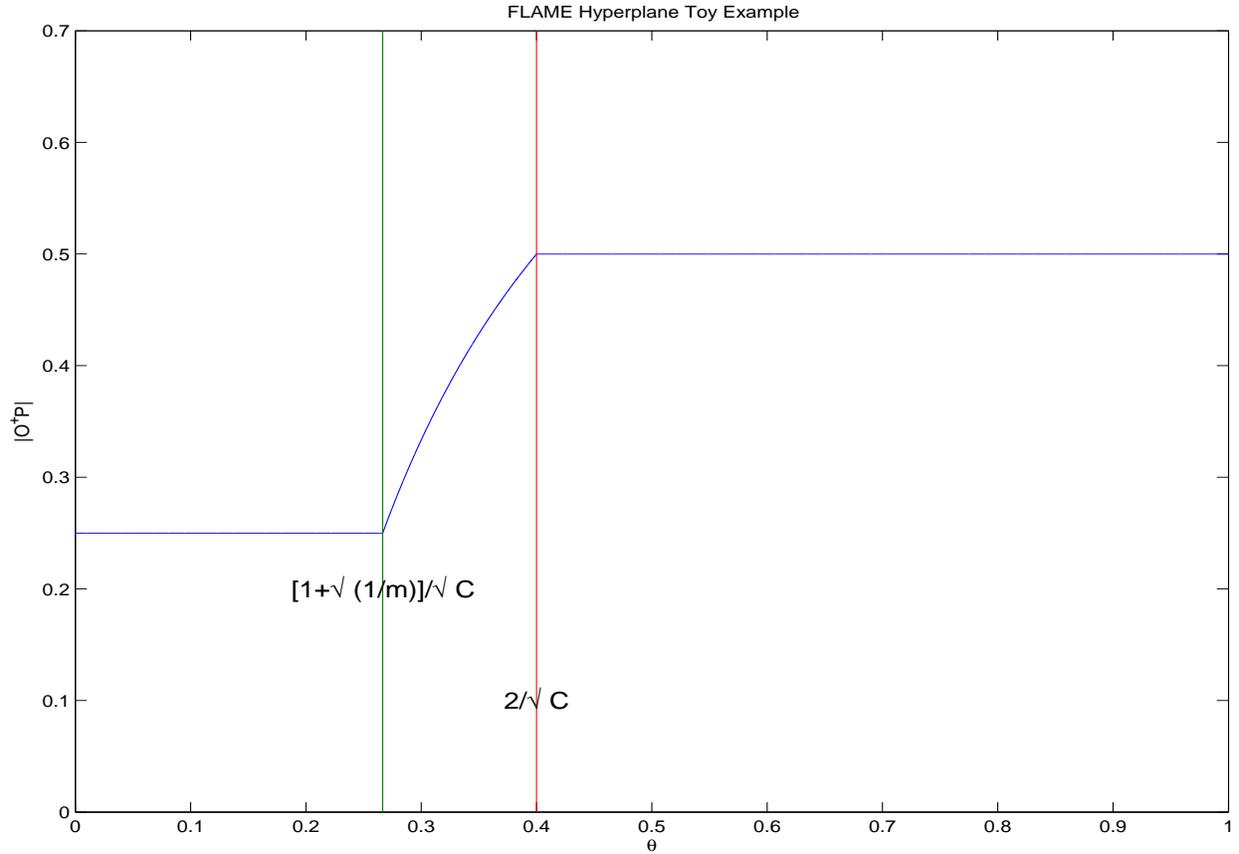
Figure S.1: A 1D toy example with $n_+ < n_-$ and $m = 9$ is used to mimic the d-asymptotic situation. The length of the line segment O^+O^- equals 1. As θ increases, the FLAME hyperplane stands still ($|O^+P|$ unchanged); when $\theta > \sqrt{1 + (1/m)}/\sqrt{C}$, $|O^+P|$ increases, which means the hyperplane moves towards the negative class, until $\theta = 2/\sqrt{C}$, after which the hyperplane remains at the midpoint of O^+O^-.

The derivation above assumes the distance between the two simplices are reasonable large, at least greater than $2/\sqrt{C}$. This is not difficult to achieve because we choose C to be a large number.

In summary, the intersection P of the FLAME hyperplane and O^+O^- stays closer to the minority class, and remains still as θ is small. When θ increases, the boundary moves towards the majority class, until reaching the midpoint of O^+O^-. This explains the simulation performance we observed in Figures S.2, 6 and S.3. We use a toy example and show the position of the FLAME hyperplane moving as θ increases in Figure S.1 in the same fashion we discussed above.

It is worth noting that the value of DWD/FLAME in terms of reducing overfitting is maximal when the dimension is greater than, but close to, the sample size. This is when data-piling starts to appear in SVM but not yet in DWD. Marron et al. (2007) showed some videos about such phenomenon. As a matter of fact, according to the geometric representation above, in the d-asymptotics, the discriminant directions for most classifiers are the same. Moreover, the projections of data points in the same class to O^+O^- are the same, which is the normal vector for the DWD, SVM and FLAME hyperplanes. Therefore, they all have data-piling in the d-asymptotics.

Derivation of the FLAME hyperplane in d asymptotics

The FLAME seeks to minimize

$$n_+ \left(1/a - \theta\sqrt{C}\right)_+ + n_- \left(1/b - \theta\sqrt{C}\right)_+ \tag{S.2}$$

$$\text{s.t. } a + b = \sqrt{d(\mu^2 + \sigma^2/n_+ + \tau^2/n_-)} = \nu\sqrt{d}. \tag{S.3}$$

When $\theta \in \left[0, (1 + \sqrt{m^{-1}})/(\nu\sqrt{dC})\right)$, it is easy to verify that both $(1/a - \theta\sqrt{C})_+$ and $(1/b - \theta\sqrt{C})_+$ are positive and equal to $1/a - \theta\sqrt{C}$ and $1/b - \theta\sqrt{C}$. In this case, the optimal solutions to problem S.2, a and β satisfy $a/b = (n_+/n_-)^{1/2} = \sqrt{m^{-1}}$. In particular, $a = \sqrt{m^{-1}}/(1 + \sqrt{m^{-1}})\nu\sqrt{d}$ and $b = 1/(1 + \sqrt{m^{-1}})\nu\sqrt{d}$.

When $\theta \in \left[(1 + \sqrt{m^{-1}})/(\nu\sqrt{dC}), 2/(\nu\sqrt{dC})\right)$, $1/b - \theta\sqrt{C} < 0$ and the optimized solutions to (S.2) are $b = 1/(\theta\sqrt{C})$ and $a = \sqrt{d}\nu - b$. Note that $a > b$.

When $\theta \in \left[2/(\nu\sqrt{dC}), 1\right]$, $a = b = 0.5\sqrt{d}\nu$

Proof to Theorem 6

We only need to prove the sure classification for the second interval, *i.e.*,

$$\theta \in \left[(1 + \sqrt{m^{-1}})/(\nu\sqrt{dC}), 2/(\nu\sqrt{dC}) \right).$$

The proofs for the other two intervals are similar to those in Qiao et al. (2010). It was shown in Hall et al. (2005) and Qiao et al. (2010) that the length of the line segment O^+O^- is $\sqrt{d}\nu$ and that the distance between the projection (denoted as P') of a new data point from the $\mathcal{X}^+$-population onto O^+O^- and the centroid of the positive class O^+ is $(\sigma^2/n_+)/(\mu^2 + \tau^2/n_-)$ times of its distance to the centroid of the negative class O^-, *i.e.*, $\frac{|O^+P'|}{|O^-P'|} = (\sigma^2/n_+)/(\mu^2 + \tau^2/n_-)$, where $|AB|$ is the length of the line segment connecting points A and B. Denote $|O^+P'|$ as a' and $|O^-P'|$ as b'. Because $a' + b' = \sqrt{d}\nu$, we must have $b' = \sqrt{d}(\mu^2 + \tau^2/n_-)/\nu$. In order for this new data point to be correctly classified to the positive class, P' has to be the same side as O^+ with respect to the intersection of the FLAME hyperplane and O^+O^-, that is,

$$b' > b$$

$$\Leftrightarrow \qquad \sqrt{d}\left(\mu^2 + \frac{\tau^2}{n_-} \right)/\nu > \frac{1}{\sqrt{C}\theta}$$

$$\Leftrightarrow \qquad \mu^2 + \frac{\tau^2}{n_-} > \frac{1}{\sqrt{dC}\theta}\nu$$

$$\Leftrightarrow \qquad \nu^2 - \frac{\sigma^2}{n_+} > \frac{1}{\sqrt{dC}\theta}\nu$$

$$\Leftrightarrow \qquad \nu^2 - \frac{1}{\sqrt{dC}\theta}\nu - \frac{\sigma^2}{n_+} > 0$$

$$\Leftrightarrow \qquad (\nu - \frac{1}{2\sqrt{dC}\theta})^2 - \frac{1}{4dC\theta^2} - \frac{\sigma^2}{n_+} > 0$$

$$\Leftarrow \qquad \nu > \sqrt{\frac{1}{4dC\theta^2} + \frac{\sigma^2}{n_+}} + \frac{1}{2\sqrt{dC}\theta}$$

$$\Leftrightarrow \qquad \mu^2 > \left[\sqrt{\frac{1}{4dC\theta^2} + \frac{\sigma^2}{n_+}} + \frac{1}{2\sqrt{dC}\theta} \right]^2 - \frac{\sigma^2}{n_+} - \frac{\tau^2}{n_-}$$

$$\Leftrightarrow \qquad \mu^2 > T - \frac{\tau^2}{n_-}.$$

We now assume that P' is the projection of a new data point from the $\mathcal{X}^-$-population. In

this situation, it can be shown that $a'/b' = (\mu^2 + \sigma^2/n_+)/(\tau^2/n_-)$ and thus $b' = \sqrt{d}\dfrac{\tau^2/n_-}{\nu}$.

To correctly classify this new data point, we only need to have $b' < b$. That is,

$$\sqrt{d}\frac{\tau^2/n_-}{\nu} < b = \frac{1}{\theta\sqrt{C}}$$

$$\Leftrightarrow \qquad \frac{\tau^2}{n_-} < \frac{1}{\theta\sqrt{dC}}\sqrt{\mu^2 + \tau^2/n_- + \sigma^2/n_+}$$

We only need to show that $\dfrac{\tau^2}{n_-} < \dfrac{1}{\theta\sqrt{dC}}\sqrt{\tau^2/n_- + \sigma^2/n_+}$. Let $q^2 = \tau^2/n_- + \sigma^2/n_+$. We need to show that

$$q^2 - \frac{\sigma^2}{n_+} < \frac{1}{\theta\sqrt{dC}}q$$

$$\Leftrightarrow \qquad \left(q - \frac{1}{2\theta\sqrt{dC}}\right)^2 - \frac{1}{4\theta^2 dC} - \frac{\sigma^2}{n_+} < 0$$

$$\Leftarrow \qquad q < \sqrt{\frac{1}{4\theta^2 dC} + \frac{\sigma^2}{n_+}} + \frac{1}{2\theta\sqrt{dC}}$$

$$\Leftrightarrow \qquad \frac{\tau^2}{n_-} < \left[\sqrt{\frac{1}{4\theta^2 dC} + \frac{\sigma^2}{n_+}} + \frac{1}{2\theta\sqrt{dC}}\right]^2 - \frac{\sigma^2}{n_+}$$

$$\Leftrightarrow \qquad \frac{\tau^2}{n_-} < T.$$

The last inequality is the condition stipulated in the theorem. $\qquad\square$

Additional figures
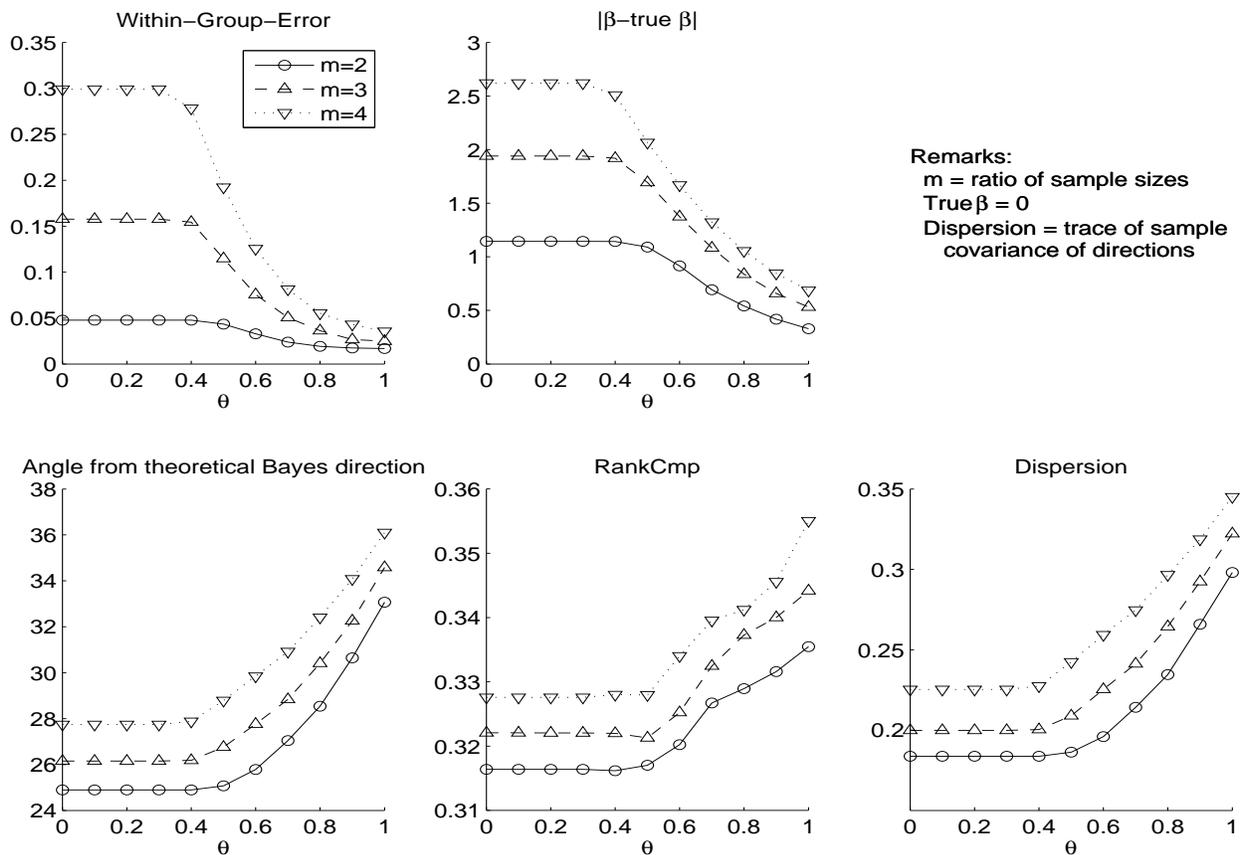
Figure S.2: Independent example. It can be seen that with FLAME turns from DWD to SVM (θ from 0 to 1), the within-class error decreases (top-left), thanks to the more accurate estimate of the intercept term (top-middle). On the other hand, this comes at the cost of larger deviation from the Bayes direction (bottom-left), incorrect rank of the importance of the variables (bottom-middle) and larger stochastic variability of the estimation directions (bottom-right).
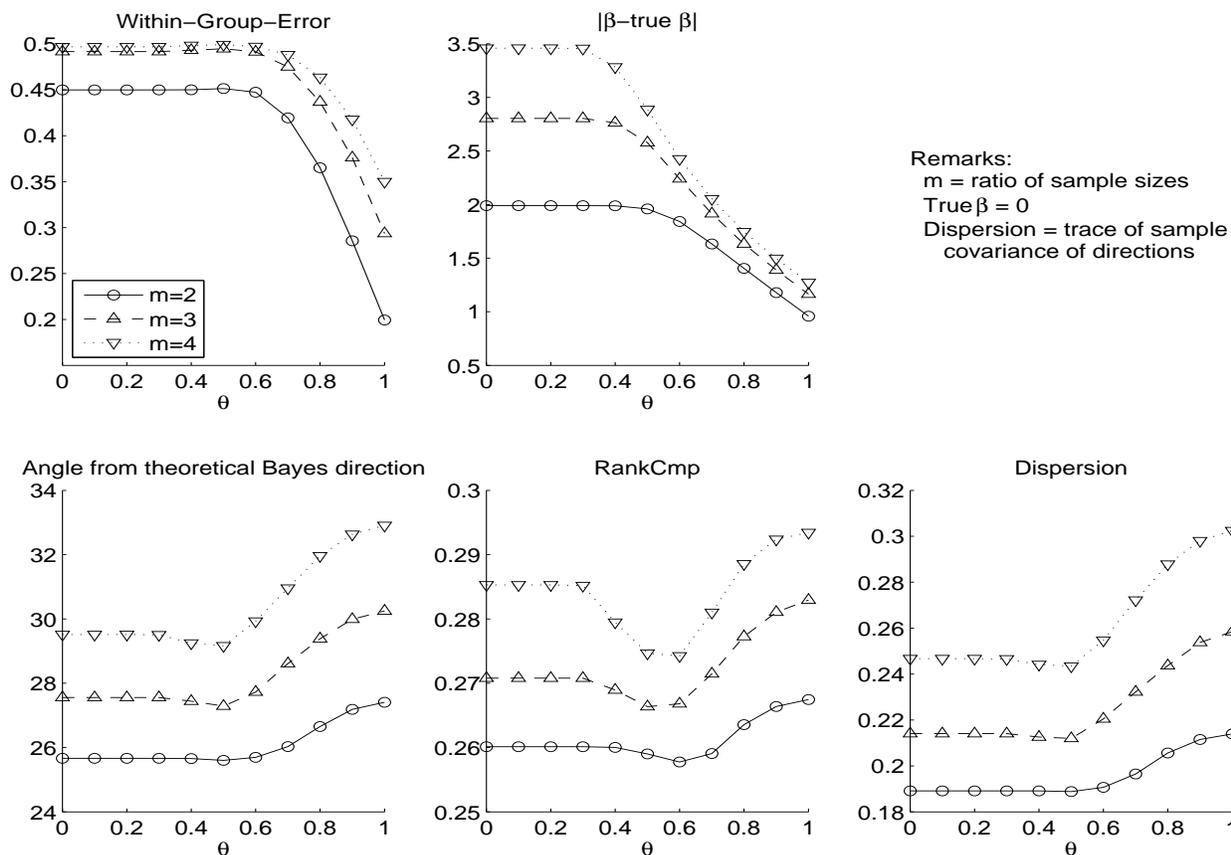
Figure S.3: Block interchangeable example. It can be seen that with FLAME turns from DWD to SVM (θ from 0 to 1), the within-class error decreases (top-left), thanks to the more accurate estimate of the intercept term (top-middle). On the other hand, this comes at the cost of larger deviation from the Bayes direction (bottom-left), incorrect rank of the importance of the variables (bottom-middle) and larger stochastic variability of the estimation directions (bottom-right). The RankComp measure in this example is an exception, in the sense that it decreases first then increases instead of monotonically increases.