
Sparse Greedy Minimax Probability Machine Classification

Thomas R. Strohmann

Department of Computer Science
University of Colorado, Boulder
strohman@cs.colorado.edu

Andrei Belitski

Department of Computer Science
University of Colorado, Boulder
Andrei.Belitski@colorado.edu

Gregory Z. Grudic

Department of Computer Science
University of Colorado, Boulder
grudic@cs.colorado.edu

Dennis DeCoste

Machine Learning Systems Group
NASA Jet Propulsion Laboratory
decoste@aig.jpl.nasa.gov

Abstract

The Minimax Probability Machine Classification (MPMC) framework [Lanckriet *et al.*, 2002] builds classifiers by minimizing the maximum probability of misclassification, and gives direct estimates of the probabilistic accuracy bound Ω . The only assumptions that MPMC makes is that good estimates of means and covariance matrixes of the classes exist. However, as with Support Vector Machines, MPMC is computationally expensive and requires extensive cross validation experiments to choose kernels and kernel parameters that give good performance. In this paper we address the computational cost of MPMC by proposing an algorithm that constructs nonlinear *sparse* MPMC (SMPMC) models by incrementally adding basis functions (i.e. kernels) one at a time – greedily selecting the next one that maximizes the accuracy bound Ω . SMPMC automatically chooses both kernel parameters and feature weights *without* using computationally expensive cross validation. Therefore the SMPMC algorithm simultaneously addresses the problem of kernel selection and feature selection (i.e. feature weighting), based solely on maximizing the accuracy bound Ω . Experimental results indicate that we can obtain reliable bounds Ω , as well as test-set accuracies that are comparable to state of the art classification algorithms.

1 Introduction

The goal of a binary classifier is to maximize the probability that unseen test data will be classified correctly. Assuming that the test data is generated from the same probability distribution as the training data, it is possible to derive specific probability bounds for the case that the decision boundary is a hyperplane. The following result due to Marshall and Olkin [1] and extended by Bertsimas and Popescu [2] provides the theoretical basis for

assigning probability bounds to hyperplane classifiers:

$$\sup_{E[\mathbf{z}]=\bar{\mathbf{z}}, Cov[\mathbf{z}]=\Sigma_{\mathbf{z}}} Pr\{\mathbf{a}^T \mathbf{z} \geq b\} = \frac{1}{1 + \omega^2} \quad \omega^2 = \inf_{\mathbf{a}^T \mathbf{t} \geq b} (\mathbf{t} - \bar{\mathbf{z}})^T \Sigma_{\mathbf{z}}^{-1} (\mathbf{t} - \bar{\mathbf{z}}) \quad (1)$$

where $\mathbf{a} \in \mathbf{R}^d$, b are the hyperplane parameters, \mathbf{z} is a random vector, and \mathbf{t} is an ordinary vector. Lanckriet et al (see [3] and [4]) used the above results to build the Minimax Probability Machine for binary classification (MPMC). From (1) we note that the only required relevant information of the underlying probability distribution for each class is its mean and covariance matrix. No other estimates and/or assumptions are needed, which implies that the obtained bound (which we refer to as Ω) is essentially distribution-free, i.e. it holds for any distribution with a certain mean and covariance matrix.

As with other classification algorithms such as Support Vector Machines (SVM) (see [5]), the main disadvantage of current MPMC implementations is that they are computationally expensive (same complexity as SVM), and require extensive cross validation experiments to choose kernels and kernel parameter to give good performance on each dataset. The goal of this paper is to propose a kernel based MPMC algorithm that directly addresses these computational issues.

Towards this end, we propose a sparse greedy MPMC (SMPMC) algorithm that efficiently builds classifiers, while at the same time maintains the distribution free probability bound of MPM type algorithms. To achieve this goal, we propose to use an iterative algorithm which adds basis functions (i.e. kernels) one by one, to an initially "empty" model. We are considering basis functions that are induced by Mercer kernels, i.e. functions of the following form $\Phi_i(\mathbf{z}) = K_i(\mathbf{z}, \mathbf{z}_i)$ (where \mathbf{z}_i is an input vector of the training data). Bases are added in a greedy way: we select the particular \mathbf{z}_i that maximizes the MPMC objective Ω . Furthermore, SMPMC chooses optimal kernel parameters that maximize this metric (hence the subscript i in K_i), including automatically weighting input features by $\gamma_j \geq 0$ for each kernel added, such that $\mathbf{z}_i = (\gamma_1 z_{i1}, \gamma_2 z_{i2}, \dots, \gamma_d z_{id})$ for d dimensional data. The proposed SMPMC algorithm automatically selects kernels and re-weights features (i.e. does feature selection) for each new added basis function, by minimizing the error bound (i.e. maximizing Ω). Thus the large computational cost of cross validation (typically used by SVM and MPMC) is avoided.

The paper is organized as follows: Section 2.1 reviews the standard MPMC; Section 2.2 describes the proposed sparse greedy MPMC algorithm (SMPMC); and Sections 2.3-2.5 show how we can use sparse MPMC to determine optimal kernel parameters. In section 3 we compare our results to the ones described in the original MPMC paper (see [4]), showing the probability bounds and the test set accuracies for different binary classification problems. The conclusion is presented in Section 4. Matlab source code for the SMPMC algorithm is available online: <http://nago.cs.colorado.edu/~strohman/papers.html>

2 Classification model

In this section we develop a sparse version of the Minimax Probability Machine for binary classification. We show that besides a significant reduction in computational cost, the sparse MPMC algorithm allows us to do automated kernel and feature selection.

2.1 Minimax Probability Machine for binary classification

We will briefly describe the underlying concepts of the MPMC framework as developed in (see [4]). The goal of MPMC is to find a decision boundary $\mathcal{H}(\mathbf{a}, b) = \{\mathbf{z} | \mathbf{a}^T \mathbf{z} = b\}$ such that the minimum probability $\Omega_{\mathcal{H}}$ of classifying future data correctly is maximized. If we assume that the two classes are generated from random vectors \mathbf{x} and \mathbf{y} , we can express

this probability bound just in terms of the means and covariances of these random vectors:

$$\Omega_{\mathcal{H}} = \inf_{\mathbf{x} \sim (\bar{\mathbf{x}}, \Sigma_{\mathbf{x}}), \mathbf{y} \sim (\bar{\mathbf{y}}, \Sigma_{\mathbf{y}})} Pr\{\mathbf{a}^T \mathbf{x} \geq b \wedge \mathbf{a}^T \mathbf{y} \leq b\} \quad (2)$$

Note that we do not make any distributional assumptions other than that $\bar{\mathbf{x}}$, $\Sigma_{\mathbf{x}}$, $\bar{\mathbf{y}}$, and $\Sigma_{\mathbf{y}}$ are bounded. Exploiting a theorem from Marshall and Olkin [1], it is possible to rewrite (2) as a closed form expression:

$$\Omega_{\mathcal{H}} = \frac{1}{1 + m^2} \quad (3)$$

where

$$m = \min_{\mathbf{a}} \sqrt{\mathbf{a}^T \Sigma_{\mathbf{x}} \mathbf{a}} + \sqrt{\mathbf{a}^T \Sigma_{\mathbf{y}} \mathbf{a}} \quad s.t. \quad \mathbf{a}^T (\bar{\mathbf{x}} - \bar{\mathbf{y}}) = 1 \quad (4)$$

The optimal hyperplane parameter \mathbf{a}_* is the vector that minimizes (4). The hyperplane parameter b_* can then be computed as:

$$b_* = \mathbf{a}_*^T \bar{\mathbf{x}} - \frac{\sqrt{\mathbf{a}_*^T \Sigma_{\mathbf{x}} \mathbf{a}_*}}{m} \quad (5)$$

A new data point z_{new} is classified according to $sign(\mathbf{a}_*^T \mathbf{z}_{new} - b_*)$; if this yields $+1$, z_{new} is classified as belonging to class \mathbf{x} , otherwise it is classified as belonging to class \mathbf{y} .

2.2 Sparse MPM classification

One of the appealing properties of Support Vector Machines is that their models typically rely only on a small fraction of the training examples, the so called support vectors. The models obtained from the kernelized MPMC (see [4]), however, use *all* of the training examples, i.e. the decision hyperplane will look like:

$$\sum_{i=1}^{N_x} a_i^{(x)} K(\mathbf{x}_i, \mathbf{z}) + \sum_{i=1}^{N_y} a_i^{(y)} K(\mathbf{y}_i, \mathbf{z}) = b \quad (6)$$

where in general all $a_i^{(x)}, a_i^{(y)} \neq 0$.

This brings up the question whether one can also construct sparse models for the MPMC where most of the coefficients $a_i^{(x)}$ or $a_i^{(y)}$ are zero. In this paper we propose to do this by starting with an initially "empty" model and then adding basis vectors one by one. As we will see shortly, this approach is speeding up both learning and evaluation time while it is still maintaining the distribution free probability bounds of the MPMC.

Before we outline the algorithm we introduce some notation:

- N = $N_x + N_y$ the total number of training examples
- ℓ = $(\ell_1, \dots, \ell_N)^T \in \{-1, 1\}^N$ the labels of the training data
- $\hat{\ell}^{(k)}$ = $(\hat{\ell}_1^{(k)}, \dots, \hat{\ell}_N^{(k)})^T \in \mathbf{R}^N$ output of the model after adding the k th basis function
- $\mathbf{a}^{(k)}$ = the MPMC coefficients when adding the k th basis function
- $b^{(k)}$ = the MPMC offset when adding the k th basis function
- $\Phi_{\mathbf{b}}$ = $(K_{\mathbf{b}}(\mathbf{b}, \mathbf{x}_1), \dots, K_{\mathbf{b}}(\mathbf{b}, \mathbf{x}_{N_x}), K_{\mathbf{b}}(\mathbf{b}, \mathbf{y}_1), \dots, K_{\mathbf{b}}(\mathbf{b}, \mathbf{y}_{N_y}))^T$
basis function evaluated on all training examples
- $\Phi_{\mathbf{x}_b}$ = $(K_{\mathbf{b}}(\mathbf{b}, \mathbf{x}_1), \dots, K_{\mathbf{b}}(\mathbf{b}, \mathbf{x}_{N_x}))^T$ evaluated only on positive examples
- $\Phi_{\mathbf{y}_b}$ = $(K_{\mathbf{b}}(\mathbf{b}, \mathbf{y}_1), \dots, K_{\mathbf{b}}(\mathbf{b}, \mathbf{y}_{N_y}))^T$ evaluated only on negative examples

Note that $\hat{\ell}^{(k)}$ is a vector of real numbers (the distances of the training data to the hyperplane before applying the *sign* function). $\mathbf{b} \in \mathbf{R}^d$ is the training example generating the basis function $\Phi_{\mathbf{b}}$. We will simply write $\Phi^{(k)}, \Phi_{\mathbf{x}}^{(k)}, \Phi_{\mathbf{y}}^{(k)}$ for the k th basis.

For the first basis we are solving the one dimensional MPMC:

$$m = \min_a \sqrt{a\sigma_{\Phi_{\mathbf{x}}^{(1)}}^2} + \sqrt{a\sigma_{\Phi_{\mathbf{y}}^{(1)}}^2} \quad s.t. \quad a(\overline{\Phi_{\mathbf{x}}^{(1)}} - \overline{\Phi_{\mathbf{y}}^{(1)}}) = 1 \quad (7)$$

Because of the constraint the feasible region contains just one value for $a^{(1)}$:

$$\begin{aligned} a^{(1)} &= 1/(\overline{\Phi_x^{(1)}} - \overline{\Phi_y^{(1)}}) \\ b^{(1)} &= a^{(1)}\overline{\Phi_x^{(1)}} - \frac{\sqrt{a\sigma_{\Phi_x^{(1)}}^2 a}}{\sqrt{a\sigma_{\Phi_x^{(1)}}^2 a + \sqrt{a\sigma_{\Phi_y^{(1)}}^2 a}}} = a^{(1)}\overline{\Phi_x^{(1)}} - \frac{\sigma_{\Phi_x^{(1)}}}{\sigma_{\Phi_x^{(1)}} + \sigma_{\Phi_y^{(1)}}} \end{aligned} \quad (8)$$

The first model evaluated on the training data looks like:

$$\widehat{\ell}^{(1)} = a^{(1)}\overline{\Phi}^{(1)} - b^{(1)} \quad (9)$$

All of the subsequent models use the previous estimation $\widehat{\ell}^{(k)}$ as one input and the next basis $\overline{\Phi}^{(k+1)}$ as the other input. More formally, we set up the two dimensional classification problem:

$$\begin{aligned} \mathbf{x}^{(k+1)} &= [\widehat{\ell}_x^{(k)}, \overline{\Phi_x^{(k+1)}}] \in \mathbf{R}^{N_x \times 2} \\ \mathbf{y}^{(k+1)} &= [\widehat{\ell}_y^{(k)}, \overline{\Phi_y^{(k+1)}}] \in \mathbf{R}^{N_y \times 2} \end{aligned} \quad (10)$$

And solve the following optimization problem:

$$m = \min_{\mathbf{a}} \sqrt{\mathbf{a}^T \Sigma_{\mathbf{x}^{(k+1)}} \mathbf{a}} + \sqrt{\mathbf{a}^T \Sigma_{\mathbf{y}^{(k+1)}} \mathbf{a}} \quad s.t. \quad \mathbf{a}^T (\overline{\mathbf{x}^{(k+1)}} - \overline{\mathbf{y}^{(k+1)}}) = 1 \quad (11)$$

Let $\mathbf{a}^{(k+1)} = (a_1^{(k+1)}, a_2^{(k+1)})^T$ be the optimal solution of (11). We set:

$$b^{(k+1)} = \mathbf{a}^{(k+1)T} \overline{\mathbf{x}^{(k+1)}} - \frac{\sqrt{\mathbf{a}^{(k+1)T} \Sigma_{\mathbf{x}^{(k+1)}} \mathbf{a}^{(k+1)}}}{\sqrt{\mathbf{a}^{(k+1)T} \Sigma_{\mathbf{x}^{(k+1)}} \mathbf{a}^{(k+1)} + \sqrt{\mathbf{a}^{(k+1)T} \Sigma_{\mathbf{y}^{(k+1)}} \mathbf{a}^{(k+1)}}} \quad (12)$$

and obtain the next model as:

$$\widehat{\ell}^{(k+1)} = a_1^{(k+1)}\widehat{\ell}^{(k)} + a_2^{(k+1)}\overline{\Phi}^{(k+1)} - b^{(k+1)} \quad (13)$$

As stated above, one computational advantage of sparse MPMC is that we typically use only a small number of training examples to obtain our final model (i.e. $k \ll N$). Another benefit is that we have to solve only one and two dimensional MPMC problems. As seen in (8) the one dimensional solution is trivial to compute. An analysis of the two dimensional problem shows that it can be reduced to the problem of finding the roots of a fourth order polynomial. Polynomials of degree 4 still have closed form solutions (see e.g. [6]) which can be computed efficiently. In the standard MPMC algorithm (see [4]), however, the solution \mathbf{a} for equation (4) has N dimensions and can therefore only be found by expensive numerical methods.

It may seem that the values of $\Omega = 1/(1 + m^2)$ which we obtain from (11) are not true for the whole model since we are considering only two dimensional problems and not all of the $k + 1$ dimensions we have added so far through our basis functions. But it turns out that the "local" bound (from the 2D MPMC) is indeed equal to the "global" bound (when considering all $k + 1$ dimensions). We state this fact more formally in the following theorem:

Theorem 1: Let $\widehat{\ell}^{(k)} = c_0 + c_1\overline{\Phi}^{(1)} + \dots + c_k\overline{\Phi}^{(k)}$ be the sparse MPMC model at the k th iteration ($k \geq 1$) and let $a_1^{(k+1)}, a_2^{(k+1)}, b^{(k+1)}$ be the solution of the two dimensional MPMC: $\widehat{\ell}^{(k+1)} = a_1^{(k+1)}\widehat{\ell}^{(k)} + a_2^{(k+1)}\overline{\Phi}^{(k+1)} - b^{(k+1)}$. Then the values of Ω for the two dimensional MPMC and for the $k + 1$ dimensional MPMC are the same.

Proof: see Appendix

2.3 Selection of bases and Gaussian Kernel widths

In our experiments we are using the Gaussian kernel which looks like:

$$K_\sigma(\mathbf{u}, \mathbf{v}) = \exp\left(-\frac{\|\mathbf{u} - \mathbf{v}\|_2^2}{2\sigma^2}\right) \quad (14)$$

Table 1: Bound Ω , Test set accuracy (TSA), and number of bases (K) for sparse and standard MPMC

Dataset	Sparse MPMC			Standard MPMC (Lanckriet et al.)		
	Ω	TSA	K	Ω	TSA	K
Twonorm	$86.4 \pm 0.1\%$	$98.3 \pm 0.4\%$	25	$91.3 \pm 0.1\%$	$95.7 \pm 0.5\%$	270
Breast Cancer	$90.9 \pm 0.1\%$	$96.8 \pm 0.3\%$	50	$89.1 \pm 0.1\%$	$96.9 \pm 0.3\%$	614
Ionosphere	$77.7 \pm 0.2\%$	$91.6 \pm 0.5\%$	25	$89.3 \pm 0.2\%$	$91.5 \pm 0.7\%$	315
Pima Diabetes	$38.2 \pm 0.1\%$	$75.4 \pm 0.7\%$	50	$32.5 \pm 0.2\%$	$76.2 \pm 0.6\%$	691
Sonar	$78.5 \pm 0.2\%$	$86.4 \pm 1.0\%$	80	$99.9 \pm 0.1\%$	$87.5 \pm 0.9\%$	187

where σ is the so called kernel width. As mentioned before one typically has to choose σ manually or determine it by tenfold cross validation (see [4]). The sparse MPMC algorithm greedily selects a basis function – out of a randomly chosen *candidate set* – to maximize Ω which is equivalent to minimizing the value of m in (7) and (11). Before we state the optimization problem for the one and two dimensional MPMC we rewrite (14) so that we can get rid of the denominator:

$$K_\gamma(\mathbf{u}, \mathbf{v}) = \exp(-\gamma \|\mathbf{u} - \mathbf{v}\|_2^2) \quad \gamma \geq 0 \quad (15)$$

The optimization problem we solve for the first iteration is then:

$$\min_\gamma m(\gamma) = \min_a \sqrt{a\sigma_{\Phi_x}^2} a + \sqrt{a\sigma_{\Phi_y}^2} a \quad s.t. \quad a(\overline{\Phi_x^{(1)}} - \overline{\Phi_y^{(1)}}) = 1 \quad (16)$$

note that – even though we did not state it explicitly – the variables a , $\sigma_{\Phi_x}^2$, $\sigma_{\Phi_y}^2$, $\overline{\Phi_x^{(1)}}$, and $\overline{\Phi_y^{(1)}}$ all depend on the kernel parameter γ .

The two dimensional problem that has to be solved for all subsequent iterations $k \geq 2$ turns into the following optimization problem for γ :

$$\min_\gamma m(\gamma) = \min_{\mathbf{a}} \sqrt{\mathbf{a}^T \Sigma_{\mathbf{x}^{(k+1)}} \mathbf{a}} + \sqrt{\mathbf{a}^T \Sigma_{\mathbf{y}^{(k+1)}} \mathbf{a}} \quad s.t. \quad \mathbf{a}^T (\overline{\mathbf{x}^{(k+1)}} - \overline{\mathbf{y}^{(k+1)}}) = 1 \quad (17)$$

2.4 Feature selection

For doing feature selection with Gaussian kernel one has to replace the uniform kernel width γ with a d dimensional vector $\vec{\gamma}$ of kernel widths:

$$K_{\vec{\gamma}}(\mathbf{u}, \mathbf{v}) = \exp(-\sum_{l=1}^d \gamma_l (u_l - v_l)^2) \quad (\gamma_l \geq 0 \quad l = 1, \dots, d) \quad (18)$$

Note that now the optimization problems (16) and (17) for the one respectively two dimensional MPMC are d dimensional instead of just one dimensional.

3 Experiments

In this section we describe the results we obtained for sparse MPMC on various classification benchmarks. We used the same data sets as Lanckriet et al. in [4] for the standard MPMC. The data sets were randomly divided into 90% training data and 10% test data and the results were averaged over 50 runs for each of the five problems (see Table 3). In all the experiments listed in Table 3 we used the feature selection algorithm and had a candidate set of size 5, i.e. at each iteration the best basis out of 5 randomly chosen candidates was selected. The results as measured in test set accuracy are comparable to the ones reported by Lanckriet et al. However, sparse MPMC always uses in all cases significantly less basis functions while still obtaining good models. Note also, that for the Sonar data set the probability bound Ω of sparse MPMC holds while the one obtained from Standard MPMC is invalid. The two plots in figure 1 show what typical learning curves for sparse MPMC look like. As the number of basis function increases, both the bound Ω and the test set accuracy start to go up and after a while stabilize. The stabilization point usually occurs earlier when one does full feature selection (a γ weight for each input dimension) instead




Figure 1: Bound Ω and Test Set accuracy (TSA) for width selection (WS) and feature selection (FS)

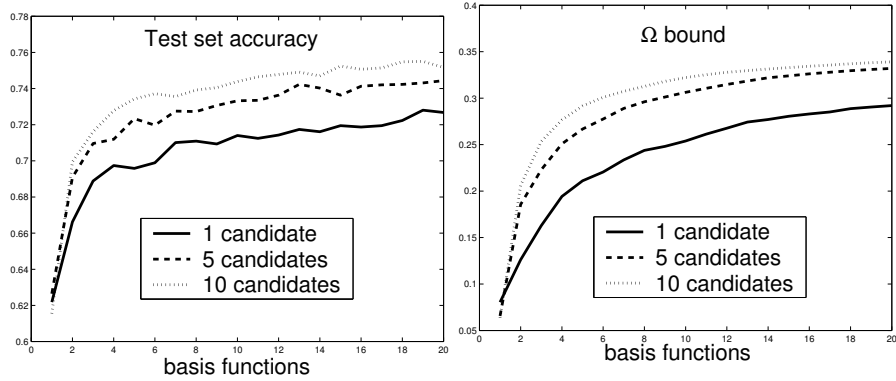


Figure 2: Accuracy and bound for the Diabetes data set using 1,5 or 10 basis candidates per iteration

of kernel width selection (one uniform γ for all dimensions). We also experimented with different sizes for the candidate set. The plots in figure 2 show what happens for 1,5, and 10 candidates. The overall behavior is that the test set accuracy as well as the Ω value converge earlier for larger candidate sets (but note that a larger candidate set also increases the computational cost per iteration).

As seen in figure 1 feature selection gives usually better results in terms of the bound Ω and the test set accuracy. Furthermore, a feature selection algorithm should indicate which features are relevant and which are not. We set up an experiment for the Twonorm data (which has 20 input features) where we added 20 additional noisy features that were not related to the output. The results are shown in figure 3 and demonstrate that the feature selection algorithm obtained from sparse MPMC is able to distinguish between relevant and irrelevant features.

4 Conclusion & future work

This paper introduces a new algorithm (Sparse Minimax Probability Machine Classification - SMPMC) for building sparse classification models that provide a lower bound on the probability of classifying future data correctly. We have shown that the method of iteratively adding basis function has significant computational advantages over the standard MPMC, while it still maintains the distribution free probability bound Ω . Experimental




Figure 3: Average feature weighting for the Twonorm data set. The first 20 features are the original inputs, the last 20 features are additional noisy inputs

results indicate that automated selection of kernel parameters, as well as automated feature selection (weighting), both key characteristics of the SMPMC algorithm, result in error rates that are competitive with those obtained by models where these parameters must be tuned by computationally expensive cross-validation.

Future research on sparse MPMC will focus on establishing a theoretical framework for a stopping criteria, when adding more basis functions (kernels) will not significantly reduce error rates, and may lead to overfitting. Also, experiments have so far focused on using Gaussian kernels as basis functions. From the experience with other kernel algorithms, it is known that other type of kernels (polynomial, tanh) can yield better results for certain applications. Furthermore, our framework is not limited to Mercer kernels, and other types of basis functions are also worth investigating. To experiment with these variations we provide the Matlab source code of the algorithm online: <http://nago.cs.colorado.edu/~strohman/papers.html>

References

- [1] A. W. Marshall and I. Olkin. Multivariate chebyshev inequalities. *Annals of Mathematical Statistics*, 31(4):1001–1014, 1960.
- [2] I. Popescu and D. Bertsimas. Optimal inequalities in probability theory: A convex optimization approach. Technical Report TM62, INSEAD, Dept. Math. O.R., Cambridge, Mass, 2001.
- [3] G. R. G. Lanckriet, L. E. Ghaoui, C. Bhattacharyya, and M. I. Jordan. Minimax probability machine. In T. G. Dietterich, S. Becker, and Z. Ghahramani, editors, *Advances in Neural Information Processing Systems 14*, Cambridge, MA, 2002. MIT Press.
- [4] G. R. G. Lanckriet, L. E. Ghaoui, C. Bhattacharyya, and M. I. Jordan. A robust minimax approach to classification. *Journal of Machine Learning Research*, 3:555–582, 2002.
- [5] B. Schölkopf and A. Smola. *Learning with Kernels*. MIT Press, Cambridge, MA, 2002.
- [6] William H. Beyer. *CRC Standard Mathematical Tables*, page 12. CRC Press Inc., Boca Raton, FL, 1987.

Appendix: Proof of Theorem 1

We have to show that the values of m are equal for the two dimensional MPMC and the $k + 1$ dimensional MPMC. We will just show the equivalence for the first term $\sqrt{\mathbf{a}^T \Sigma_{\mathbf{x}} \mathbf{a}}$,

an analogue argumentation will hold for the second term.

For the two dimensional MPMC we have the following for the term under the square root:

$$\begin{aligned} & \begin{pmatrix} a_1^{(k+1)} & a_2^{(k+1)} \end{pmatrix} \begin{pmatrix} \sigma_{\widehat{\ell}_x^{(k)}}^2 & \sigma_{\widehat{\ell}_x^{(k)} \Phi_x^{(k+1)}} \\ \sigma_{\Phi_x^{(k+1)} \widehat{\ell}_x^{(k)}} & \sigma_{\Phi_x^{(k+1)}}^2 \end{pmatrix} \begin{pmatrix} a_1^{(k+1)} \\ a_2^{(k+1)} \end{pmatrix} \\ &= [a_1^{(k+1)}]^2 \sigma_{\widehat{\ell}_x^{(k)}}^2 + 2a_1^{(k+1)} a_2^{(k+1)} \sigma_{\widehat{\ell}_x^{(k)} \Phi_x^{(k+1)}} + [a_2^{(k+1)}]^2 \sigma_{\Phi_x^{(k+1)}}^2 \end{aligned}$$

Note that we can rewrite

$$\begin{aligned} \sigma_{\widehat{\ell}_x^{(k)}}^2 &= Cov(c_0 + c_1 \Phi_x^{(1)} + \dots + c_k \Phi_x^{(k)}, c_0 + c_1 \Phi_x^{(1)} + \dots + c_k \Phi_x^{(k)}) \\ &= \sum_{i=1}^k \sum_{j=1}^k c_i c_j Cov(\Phi_x^{(i)}, \Phi_x^{(j)}) \\ \sigma_{\widehat{\ell}_x^{(k)} \Phi_x^{(k+1)}} &= Cov(c_0 + c_1 \Phi_x^{(1)} + \dots + c_k \Phi_x^{(k)}, \Phi_x^{(k+1)}) \\ &= \sum_{i=1}^k c_i Cov(\Phi_x^{(i)}, \Phi_x^{(k+1)}) \end{aligned}$$

by using properties of the sample covariance.

For the $k + 1$ dimensional MPMC let us first determine the $k + 1$ coefficients:

$$\begin{aligned} \widehat{\ell}^{(k+1)} &= a_1^{(k+1)}(c_0 + c_1 \Phi_x^{(1)} + \dots + c_k \Phi_x^{(k)}) + a_2^{(k+1)} \Phi_x^{(k+1)} - b^{(k+1)} \\ &= a_1^{(k+1)} c_1 \Phi_x^{(1)} + \dots + a_1^{(k+1)} c_k \Phi_x^{(k)} + a_2^{(k+1)} \Phi_x^{(k+1)} + a_1^{(k+1)} c_0 - b^{(k+1)} \end{aligned}$$

The term under the square root then looks like:

$$\begin{pmatrix} a_1^{(k+1)} c_1 \\ \dots \\ a_1^{(k+1)} c_k \\ a_2^{(k+1)} \end{pmatrix}^T \begin{pmatrix} \sigma_{\Phi_x^{(1)}}^2 & \dots & \sigma_{\Phi_x^{(1)} \Phi_x^{(k)}} & \sigma_{\Phi_x^{(1)} \Phi_x^{(k+1)}} \\ \dots & \dots & \dots & \dots \\ \sigma_{\Phi_x^{(k)} \Phi_x^{(1)}} & \dots & \sigma_{\Phi_x^{(k)}}^2 & \sigma_{\Phi_x^{(k)} \Phi_x^{(k+1)}} \\ \sigma_{\Phi_x^{(k+1)} \Phi_x^{(1)}} & \dots & \sigma_{\Phi_x^{(k+1)} \Phi_x^{(k)}} & \sigma_{\Phi_x^{(k+1)}}^2 \end{pmatrix} \begin{pmatrix} a_1^{(k+1)} c_1 \\ \dots \\ a_1^{(k+1)} c_k \\ a_2^{(k+1)} \end{pmatrix}$$

If we multiply out the all the elements that are in row $\leq k$ and in column $\leq k$ of the covariance matrix we obtain:

$$\begin{aligned} & \sum_{i=1}^k \sum_{j=1}^k a_1^{(k+1)} c_i Cov(\Phi_x^{(i)}, \Phi_x^{(j)}) a_1^{(k+1)} c_j \\ &= [a_1^{(k+1)}]^2 \sum_{i=1}^k \sum_{j=1}^k c_i c_j Cov(\Phi_x^{(i)}, \Phi_x^{(j)}) \\ &= [a_1^{(k+1)}]^2 \sigma_{\widehat{\ell}_x^{(k)}}^2 \end{aligned}$$

If we multiply out all the elements that are either in row $k + 1$ and in column $\leq k$ or in row $\leq k$ and in column $k + 1$ we obtain:

$$\begin{aligned} & \sum_{i=1}^k a_1^{(k+1)} c_i Cov(\Phi_x^{(i)}, \Phi_x^{(k+1)}) a_2^{(k+1)} + \sum_{j=1}^k a_2^{(k+1)} Cov(\Phi_x^{(k+1)}, \Phi_x^{(j)}) a_1^{(k+1)} c_j \\ &= 2a_1^{(k+1)} a_2^{(k+1)} \sum_{i=1}^k c_i Cov(\Phi_x^{(i)}, \Phi_x^{(k+1)}) \\ &= 2a_1^{(k+1)} a_2^{(k+1)} \sigma_{\widehat{\ell}_x^{(k)} \Phi_x^{(k+1)}} \end{aligned}$$

Finally, we multiply out the last element in row $k + 1$ and column $k + 1$

$$[a_2^{(k+1)}]^2 \sigma_{\Phi_x^{(k+1)}}^2$$

As we see, all of the three terms match up exactly with the two dimensional MPMC. Since this will also hold for the $\sqrt{\mathbf{a}^T \Sigma_y \mathbf{a}}$ term in m , we have shown that m (and therefore Ω) is equal for the two dimensional and the $k + 1$ dimensional MPMC. \square