

# Sparse Multiple Kernel Learning: Alternating Best Response and Semidefinite Relaxations

Anonymous authors

Paper under double-blind review

## Abstract

We study the Sparse Multiple Kernel Learning (SMKL), which is the problem of selecting a *sparse* convex combination of prespecified kernels for support vector binary classification. Unlike prevailing  $\ell_1$ -regularized approaches that approximate a sparsifying penalty, we formulate the problem by imposing an explicit cardinality constraint on the kernel weights and add an  $\ell_2$  penalty for robustness. We solve the resulting non-convex minimax problem via an *alternating best response* algorithm with two subproblems: the  $\alpha$ -subproblem is a standard kernel SVM dual solved via LIBSVM, while the  $\beta$ -subproblem admits an efficient solution via the Greedy Selector and Simplex Projector algorithm. We reformulate SMKL as a mixed integer semidefinite optimization problem and derive a hierarchy of semidefinite convex relaxations which can be used to certify near-optimality of the solutions returned by our best response algorithm and also to warm start it. On ten UCI benchmarks, our method with random initialization outperforms state-of-the-art MKL approaches in out of sample prediction accuracy on average by 3.34 percentage points (relative to the best performing benchmark) while selecting a small number of candidate kernels in comparable runtime. With warm starting, our method outperforms the best performing benchmark's out of sample prediction accuracy on average by 4.05 percentage points. Our convex relaxations provide a certificate that in several cases, the solution returned by our best response algorithm is the globally optimal solution.

## 1 Introduction

Given a dataset  $\{(\mathbf{x}_i, y_i)\}_{i=1}^n$  where  $\mathbf{x}_i \in \mathbb{R}^m$  are  $m$ -dimensional features and  $y_i \in \{-1, 1\}$  are binary labels, the Kernel Support Vector Machines problem (KSVM) seeks to select a possibly infinite dimensional feature map  $\Phi : \mathbb{R}^m \rightarrow \mathbb{R}^d$  and learn a linear classification rule  $\hat{y} = \text{sgn}(\mathbf{w}^T \Phi(\mathbf{x}) + b)$  that generalizes well to unseen data. For a fixed feature map  $\Phi$ , the learning problem can be written in primal form as:

$$\min_{\mathbf{w} \in \mathbb{R}^d, \boldsymbol{\eta} \in \mathbb{R}_+^n, b \in \mathbb{R}} \frac{1}{2} \|\mathbf{w}\|_2^2 + C \sum_{i=1}^n \eta_i, \quad \text{s.t.} \quad y_i(\mathbf{w}^T \Phi(\mathbf{x}_i) + b) \geq 1 - \eta_i \quad \forall i, \quad (1)$$

where  $C \in \mathbb{R}_+$  is a hyperparameter that is to be cross-validated by minimizing a validation metric (see, e.g., Owen & Perry, 2009) to obtain strong out-of-sample performance in theory and practice (Bousquet & Elisseeff, 2002). The first term in the objective function of Problem 1 encourages minimal norm solutions as a regularity condition that reduces overfitting while the second term in the objective function penalizes misclassified examples in the training dataset.

The KSVM learning problem given by equation 1 emits a dual formulation that is often the preferred formulation to consider as it avoids requiring explicit construction of the feature map  $\Phi$  (which may be prohibitive for large dimensional feature maps or simply impossible in the case of infinite dimensional feature maps). Specifically, the dual form of KSVM is given by

$$\max_{\alpha \in \mathbb{R}_+^n} \sum_{i=1}^n \alpha_i - \frac{1}{2}(\mathbf{y} \circ \alpha)^T \mathbf{K}(\mathbf{y} \circ \alpha), \quad \text{s.t.} \quad \sum_{i=1}^n \alpha_i y_i = 0, \quad C \geq \alpha_i \geq 0 \quad \forall i, \quad (2)$$

where  $\mathbf{K}$  is the kernel matrix corresponding to the feature map  $\Phi$  (meaning that  $K_{ij} = \Phi(\mathbf{x}_i)^T \Phi(\mathbf{x}_j)$ ).

In Vanilla KSVM, the choice of kernel matrix is typically selected via cross-validation. This however is inefficient and significantly limits the number of kernels that are considered. Moreover, it can be natural to use different kernels for different features. This might occur, for instance, if each data point consists of both text-based and visual features. Multiple kernel learning (MKL) addresses this limitation by considering a prespecified collection of  $q$  candidate kernels  $\{\mathbf{K}_i\}_{i=1}^q$  and searching over the space of convex combinations of said kernels (note that some treatments of MKL consider general conic combinations of candidate kernels rather than only convex combinations). Explicitly, MKL is formulated as:

$$\begin{aligned} \min_{\beta \in \mathbb{R}_+^q} \max_{\alpha \in \mathbb{R}_+^n} \quad & \sum_{i=1}^n \alpha_i - \frac{1}{2}(\mathbf{y} \circ \alpha)^T \left[ \sum_{i=1}^q \beta_i \mathbf{K}_i \right] (\mathbf{y} \circ \alpha) \\ \text{s.t.} \quad & \sum_{i=1}^n \alpha_i y_i = 0, \quad C \geq \alpha_i \geq 0 \quad \forall i, \quad \|\beta\|_1 = 1. \end{aligned} \quad (3)$$

Since any symmetric positive semidefinite matrix is a valid kernel (meaning that it corresponds to some possibly infinite dimensional feature map), the matrix  $\mathbf{K}(\beta) = \sum_{i=1}^q \beta_i \mathbf{K}_i$  is a valid kernel matrix for any  $\beta \in \mathbb{R}_+^q$ . In words, Problem 3 seeks to find the kernel in the convex hull of the  $q$  kernels  $\{\mathbf{K}_i\}_{i=1}^q$  that results in the lowest training objective value when used to fit KSVM. The inclusion of the constraint  $\|\beta\|_1 = 1$  in Problem 3 is typically motivated as inducing sparsity in  $\beta$  as a convex surrogate to the  $\ell_0$  norm function. In this work, we take the approach of directly imposing a cardinality constraint on  $\beta$  and demonstrate that this yields a tractable formulation that produces superior solutions to MKL. Explicitly, given a target sparsity level  $k \in \mathbb{N}$ , we consider:

$$\begin{aligned} \min_{\beta \in \mathbb{R}_+^q} \max_{\alpha \in \mathbb{R}_+^n} \quad & \sum_{i=1}^n \alpha_i - \frac{1}{2}(\mathbf{y} \circ \alpha)^T \left[ \sum_{i=1}^q \beta_i \mathbf{K}_i \right] (\mathbf{y} \circ \alpha) + \lambda \|\beta\|_2^2 \\ \text{s.t.} \quad & \sum_{i=1}^n \alpha_i y_i = 0, \quad C \geq \alpha_i \geq 0 \quad \forall i, \quad \|\beta\|_1 = 1, \quad \|\beta\|_0 \leq k_0, \end{aligned} \quad (4)$$

where in addition to the cardinality constraint on  $\beta$  we have introduced an  $\ell_2$  regularization term in the objective function to encourage robustness.

## 1.1 Contribution and Structure

In this paper, we present an alternating best response algorithm that produces high quality solutions to Problem 4 efficiently and outperforms state-of-the-art benchmark methods. We exactly reformulate Problem 4 as a mixed integer semidefinite program which emits a natural hierarchy of semidefinite convex relaxations that produce lower bounds on the optimal objective value of Problem 4. We employ these lower bounds to produce certificates of near optimality of the feasible solutions returned by our alternating best response algorithm and additionally as warm starts for our algorithmic approach. We present rigorous numerical results across UCI benchmark datasets. On average, our approach outperforms the best performing benchmark method in out of sample prediction accuracy by 3.34 percentage points while selecting a small number of candidate kernels in comparable runtime when using a random initialization. With warm starting, our approach outperforms the best performing benchmark's out of sample prediction accuracy by 4.05 percentage points on average. In many cases, the lower bounds produced by our semidefinite relaxations certify that the solution returned by our alternating best response heuristic is the globally optimal solution to Problem 4.

The rest of the paper is laid out as follows. In Section 2, we review previous work that is closely related to SMKL. In Section 3, we study the two subproblems that emerge naturally from Problem 4 and we present our

alternating best response algorithm. We reformulate Problem 4 as a mixed integer semidefinite optimization problem in Section 4 and present a family of convex relaxations. Finally, in Section 5 we investigate the performance of our algorithm against benchmark methods on real world UCI datasets.

**Notation:** We let nonbold face characters such as  $b$  denote scalars, lowercase bold faced characters such as  $\mathbf{x}$  denote vectors, uppercase bold faced characters such as  $\mathbf{X}$  denote matrices, and calligraphic uppercase characters such as  $\mathcal{Z}$  denote sets. We let  $[n]$  denote the set of running indices  $\{1, \dots, n\}$ . We let  $\mathbf{0}_n$  denote an  $n$ -dimensional vector of all 0's,  $\mathbf{0}_{n \times m}$  denote an  $n \times m$ -dimensional matrix of all 0's, and  $\mathbf{I}_n$  denote the  $n \times n$  identity matrix. We let  $\mathcal{S}^n$  denote the cone of  $n \times n$  symmetric matrices and  $\mathcal{S}_+^n$  denote the cone of  $n \times n$  positive semidefinite matrices.

## 2 Literature Review and Preliminaries

In this section, we present a non-exhaustive literature review to highlight key ideas from the MKL literature that we build upon with a particular focus on open source methods that are used as benchmarks in this work. We additionally review commonly used kernel families that we leverage in this work.

### 2.1 Literature Review

In this subsection, we present a brief overview of the MKL literature and highlight a collection of open source approaches which will be used as benchmark methods in this work. As an exhaustive review of the literature is outside the scope of this paper, we refer the interested reader to Gönen & Alpaydm (2011).

#### 2.1.1 MKL Foundations

Multiple Kernel Learning emerged in the early 2000s as a framework to learn an optimal combination of prespecified base kernels rather than relying on a single preselected kernel. Lanckriet et al. (2004) first formulated MKL as a semidefinite optimization problem that learned nonnegative kernel weights alongside an SVM classifier. The authors demonstrated that combining kernels from heterogeneous data sources can significantly improve accuracy over the best single kernel. As off the shelf interior point methods suffer from scalability challenges for semidefinite programming, several subsequent efforts sought to improve the scalability of MKL by leveraging alternate formulations and algorithmic methods. To this end, Bach et al. (2004) reformulated MKL as a quadratically constrained quadratic program and designed a sequential minimization based solution method. In order to solve MKL problems with large datasets, Sonnenburg et al. (2006) developed a column generation algorithm to solve a semi-infinite linear programming reformulation of MKL and Rakotomamonjy et al. (2008) developed the widely used SimpleMKL solver which leverages an efficient gradient descent based algorithm.

Recognizing that the number of base kernels to consider may often be considerably large, several authors sought to develop techniques that produced a sparse kernel meaning that the selected kernel was formed using only a small subset of the input base kernels. As directly imposing a cardinality constraint on the kernel combination vector  $\beta$  produces an NP hard problem, many approaches leverage the  $\ell_1$  norm function as a convex surrogate for the  $\ell_0$  norm function analogous to the approach taken by Lasso to solve sparse regression (Tibshirani, 1996). To this end, Bach (2008) proposed a MKL formulation that sought to enforce group-level sparsity through the inclusion of a weighted  $\ell_1$ -norm regularizer in the SVM dual which was shown to have desirable consistency properties. Several other convex approximations have been employed in the literature. For example, Koltchinskii & Yuan (2010) attempt to induce sparsity through a combination of an empirical  $\ell_2$  norm and the reproducing kernel Hilbert space norms induced by the base kernels. In an alternate direction, Kloft et al. (2011) explore  $\ell_p$  norm constraints (for  $p \geq 1$ ) and theoretically explore their influence on sparsity.

Beyond  $\ell_p$  surrogates, some authors have explored non-convex penalties that more closely approximate the  $\ell_0$  norm. For example, Subrahmanya & Shin (2010) propose a log-penalised MKL formulation that attains competitive accuracy on cancer-detection and hyperspectral-image benchmarks while selecting markedly

fewer kernels than other multiple kernel methods. Nevertheless, the sparsity in their approach is still induced through a smooth relaxation rather than an explicit cardinality constraint.

Despite their empirical success, smooth or convex surrogates of  $\ell_0$  do not guarantee truly sparse solutions. Lasso-type penalties, for instance, produce robust estimators but not sparse estimators and can be outperformed by methods that exactly enforce sparsity (Bertsimas & Copenhaver, 2018). This insight has motivated a recent line of work that revisits direct cardinality constraints in MKL. Xue & Song (2019) propose a non-convex formulation of sparse MKL which they solve using techniques from difference-of-convex programming. However, their approach considers using various smooth approximations to the  $\ell_0$  norm (exponential, logarithmic, capped  $\ell_1$  and SCAD). In this work, we consider an explicit  $\ell_0$  norm constraint which we enforce directly without approximation.

### 2.1.2 Open Source MKL Benchmarks

In this work, we benchmark our methods against three well studied and widely used MKL algorithms that are implemented in the open-source **MKLpy** library (Lauriola & Aioli, 2020): *AverageMKL*, *EasyMKL* (Aioli & Donini, 2015), and *CKA* (Cortes et al., 2012). We briefly outline each of these approaches below.

**AverageMKL** This baseline sets all kernel weights uniformly as  $\beta_i = \frac{1}{q}$  for all  $i$ . This produces the following combined kernel:

$$\mathbf{K}(\boldsymbol{\beta}) = \frac{1}{q} \sum_{i=1}^q \mathbf{K}_i.$$

Although AverageMKL trivially determines kernel weights, this method is surprisingly competitive when most kernels are informative (Lewis et al., 2006). Its parameter-free design and limited training cost make it a useful sanity check and benchmark for evaluating the performance of more sophisticated MKL methods.

**EasyMKL** This method is a scalable MKL algorithm that learns a combined kernel without alternating between SVM training and kernel weight updates. Explicitly, EasyMKL seeks a combined kernel  $\mathbf{K}(\boldsymbol{\beta}) = \sum_{i=1}^q \beta_i \mathbf{K}_i$  where  $\boldsymbol{\beta} \geq 0$  and  $\|\boldsymbol{\beta}\|_2^2 = 1$ . This is achieved by introducing a distribution vector  $\boldsymbol{\gamma} \in \Gamma$  over the training examples (separately normalized on positive and negative classes), and letting  $\mathbf{Y} = \text{diag}(\mathbf{y})$ . Here,  $\Gamma$  denotes the set of possible distributions. Define  $\mathbf{d}(\boldsymbol{\gamma}) \in \mathbb{R}^q$  by letting  $d_i(\boldsymbol{\gamma}) = \boldsymbol{\gamma}^\top \mathbf{Y} \mathbf{K}_i \mathbf{Y} \boldsymbol{\gamma}$ , the (squared) margin of  $\mathbf{K}_i$  under weighting  $\boldsymbol{\gamma}$ . EasyMKL casts MKL as the saddle-point problem

$$\min_{\boldsymbol{\gamma} \in \Gamma} \max_{\|\boldsymbol{\beta}\|_2=1} [(1-\lambda) \boldsymbol{\beta}^\top \mathbf{d}(\boldsymbol{\gamma}) + \lambda \|\boldsymbol{\gamma}\|_2^2], \quad (5)$$

where  $\lambda \in [0, 1]$  trades off margin maximization against variance regularization. For fixed  $\boldsymbol{\gamma}$ , the inner maximization admits the closed-form solution  $\boldsymbol{\beta}^* = \frac{\mathbf{d}(\boldsymbol{\gamma})}{\|\mathbf{d}(\boldsymbol{\gamma})\|_2}$  which reduces Problem 5 to

$$\min_{\boldsymbol{\gamma} \in \Gamma} [(1-\lambda) \|\mathbf{d}(\boldsymbol{\gamma})\|_2 + \lambda \|\boldsymbol{\gamma}\|_2^2]. \quad (6)$$

The authors then replace  $\|\mathbf{d}(\boldsymbol{\gamma})\|_2$  in Problem 6 with  $\|\mathbf{d}(\boldsymbol{\gamma})\|_1$ , which produces a more tractable optimization problem and empirically is sparsity inducing.

**CKA (Centered Kernel Alignment)** This approach selects kernel weights by directly maximizing the alignment between a convex combination of base kernels and the target label kernel. Let  $\mathbf{K}_0 = \mathbf{y}\mathbf{y}^\top$  and denote by  $\mathbf{K}_i^c = \mathbf{C} \mathbf{K}_i \mathbf{C}$  the centered version of kernel  $\mathbf{K}_i$  where  $\mathbf{C} = \mathbf{I} - \frac{1}{n} \mathbf{1}\mathbf{1}^\top$ . Let  $\mathbf{a} = (\langle \mathbf{K}_1^c, \mathbf{K}_0^c \rangle, \dots, \langle \mathbf{K}_q^c, \mathbf{K}_0^c \rangle)^\top$  and let  $M_{ij} = \langle \mathbf{K}_i^c, \mathbf{K}_j^c \rangle$ . Then CKA solves the generalized Rayleigh quotient

$$\max_{\boldsymbol{\beta} \geq 0, \|\boldsymbol{\beta}\|_2=1} \frac{(\boldsymbol{\beta}^\top \mathbf{a})^2}{\boldsymbol{\beta}^\top \mathbf{M} \boldsymbol{\beta}},$$

whose solution admits the closed-form  $\boldsymbol{\beta}^* = \mathbf{M}^{-1} \mathbf{a} / \|\mathbf{M}^{-1} \mathbf{a}\|_2$ . While CKA typically produces dense weight vectors, its efficiency and strong empirical performance make it a valuable baseline for MKL.

In summary, AverageMKL serves as a simple and effective baseline, EasyMKL provides a sparsity inducing and scalable convex method grounded in margin maximization, and CKA offers a fast alignment-based solution. The strong empirical performance and reproducibility of these methods make them well-suited as benchmark methods for evaluating sparse MKL algorithms.

## 2.2 Common Kernel Families

In this subsection, we briefly review commonly used families of kernel functions that we consider in this work.

### 2.2.1 Linear

The linear kernel is given by  $K(\mathbf{x}_i, \mathbf{x}_j) = \mathbf{x}_i^\top \mathbf{x}_j$  for vectors  $\mathbf{x}_i, \mathbf{x}_j \in \mathbb{R}^n$ . This is the most primitive kernel which corresponds to the identity feature map. This kernel works well when data is approximately linearly separable in its latent space. Linear kernels are especially popular in high-dimensional settings such as document classification, where sparse bag-of-words representations often admit an effective linear decision boundary and enable efficient training (Joachims, 1998).

### 2.2.2 Polynomial

The polynomial kernel of degree  $d \in \mathbb{N}$  is given by  $K(\mathbf{x}_i, \mathbf{x}_j) = (\alpha \mathbf{x}_i^\top \mathbf{x}_j + c)^d$  for vectors  $\mathbf{x}_i, \mathbf{x}_j \in \mathbb{R}^n$  where  $\alpha > 0, c \in \mathbb{R}$ . This kernel allows learning nonlinear relationships by mapping inputs into the feature space of monomials up to degree  $d$ . Polynomial kernels were featured in early SVM work to demonstrate the kernel trick Boser et al. (1992); Cortes & Vapnik (1995). They can be useful for data where feature interactions of a known order are important, though choosing a very large degree  $d$  increases the risk of overfitting.

### 2.2.3 Radial Basis Function (RBF)

The RBF kernel, sometimes referred to as the Gaussian kernel, is given by  $K(\mathbf{x}_i, \mathbf{x}_j) = \exp(-\gamma \|\mathbf{x}_i - \mathbf{x}_j\|^2)$  for vectors  $\mathbf{x}_i, \mathbf{x}_j \in \mathbb{R}^n$  where  $\gamma > 0$ . The RBF kernel corresponds to an infinite-dimensional feature map and is universal, meaning that its reproducing kernel Hilbert space can approximate any continuous function arbitrarily well on compact subsets of  $\mathbb{R}^n$  (Steinwart, 2001). This kernel has become one of the most commonly used kernels in SVM practice (Schölkopf & Smola, 2002). The RBF kernel is governed by a single bandwidth parameter  $\gamma$ , simplifying model selection relative to, for example, polynomial kernels. Thanks to its flexibility and strong empirical performance, it is the default choice in many SVM packages (e.g., LIBSVM uses the RBF kernel by default (Chang & Lin, 2011a)).

### 2.2.4 Sigmoid

The sigmoid kernel is given by  $K(\mathbf{x}_i, \mathbf{x}_j) = \tanh(\gamma \mathbf{x}_i^\top \mathbf{x}_j + r)$  for vectors  $\mathbf{x}_i, \mathbf{x}_j \in \mathbb{R}^n$  where  $\gamma > 0, r \in \mathbb{R}$ . Also known as the neural network kernel, this kernel arises from the activation function of a two-layer perceptron. It gained some popularity in early SVM experiments due to its connection to neural networks (Burges, 1998). In practice, the sigmoid kernel can behave like the RBF kernel for certain parameter settings (Schölkopf & Smola, 2002).

### 2.2.5 Laplacian

The Laplacian kernel is given by  $K(\mathbf{x}_i, \mathbf{x}_j) = \exp(-\gamma \|\mathbf{x}_i - \mathbf{x}_j\|_1)$  for vectors  $\mathbf{x}_i, \mathbf{x}_j \in \mathbb{R}^n$  where  $\gamma > 0$ . The Laplacian kernel is similar to the RBF kernel but uses the  $\ell_1$  distance as a measure of similarity in place of using the  $\ell_2$  distance. Like the RBF kernel, the Laplacian kernel has a single width parameter  $\gamma$  that controls the decay of the similarity measure.

## 3 An Alternating Best Response Algorithm

In this section, we present an alternating best response algorithm that produces high quality feasible solutions to Problem 4. The presence of the  $\ell_0$  constraint on  $\beta$  makes the problem non convex and therefore not

amenable to commonly used gradient descent ascent style algorithms for minimax optimization. Nevertheless, we show in Sections 3.1 and 3.2 that both the maximization over  $\alpha$  for a fixed value of  $\beta$  (which we refer to as the  $\alpha$ -subproblem) and the minimization over  $\beta$  for a fixed value of  $\alpha$  (which we refer to as the  $\beta$ -subproblem) can be solved efficiently.

### 3.1 The $\alpha$ -Subproblem

For a fixed value  $\beta \in \mathbb{R}_+^q$ , define the combined kernel  $\mathbf{K}(\beta)$  as  $\mathbf{K}(\beta) := \sum_{j=1}^q \beta_j \mathbf{K}_j$ . Then, the  $\alpha$ -subproblem is given by:

$$\max_{\substack{\alpha \in [0, C]^n \\ \sum_{i=1}^n y_i \alpha_i = 0}} \left\{ \sum_{i=1}^n \alpha_i - \frac{1}{2} (\mathbf{y} \circ \alpha)^\top \mathbf{K}(\beta) (\mathbf{y} \circ \alpha) \right\}. \quad (7)$$

This formulation is equivalent to the standard Kernel SVM dual problem, with the only difference being that the kernel matrix is given by  $\mathbf{K}(\beta)$ . As a convex quadratic program, Problem 7 can be solved to global optimality using any one of several well-established solvers. In the implementation of our approach, we leverage the LIBSVM solver (Chang & Lin, 2011b) which employs the Sequential Minimal Optimization (SMO) algorithm. SMO decomposes the large-scale quadratic program into a series of analytically solvable two-dimensional subproblems. This coordinate descent strategy is both efficient and robust in practice.

### 3.2 The $\beta$ -Subproblem

For a fixed value of  $\alpha \in \mathbb{R}^n$  that is feasible to Problem 7, the  $\beta$ -subproblem is given by

$$\min_{\beta \in \mathbb{R}_+^q} -\frac{1}{2} (\mathbf{y} \circ \alpha)^\top \left[ \sum_{i=1}^q \beta_i \mathbf{K}_i \right] (\mathbf{y} \circ \alpha) + \lambda \|\beta\|_2^2 \quad \text{s.t.} \quad \|\beta\|_1 = 1, \quad \|\beta\|_0 \leq k_0. \quad (8)$$

We define  $d_i := (\mathbf{y} \circ \alpha)^\top \mathbf{K}_i (\mathbf{y} \circ \alpha)$  for every  $i \in [q]$ . Given this, we can express the objective function of Problem 8 as:

$$f(\beta) = \sum_{i=1}^q \left[ -\frac{1}{2} \beta_i d_i + \lambda \beta_i^2 \right].$$

For any given  $\beta_i$ , completing the square shows that we have

$$-\frac{1}{2} \beta_i d_i + \lambda \beta_i^2 = \lambda \left( \beta_i - \frac{d_i}{4\lambda} \right)^2 - \frac{d_i^2}{16\lambda}.$$

Observe that if Problem 8 were unconstrained, the optimal solution  $\beta^*$  would be given by  $\beta_i^* = \frac{d_i}{4\lambda}$ . Let  $\mathbf{w} := \frac{\mathbf{d}}{4\lambda}$  and define the sets  $\Sigma_{k_0} = \{\beta \in \mathbb{R}^q : \|\beta\|_0 \leq k_0\}$  and  $\Delta^+ = \{\beta \in \mathbb{R}_+^q : \sum_{j=1}^q \beta_j = 1\}$ . Problem 8 has the same optimal solution as:

$$\min_{\beta \in \Sigma_{k_0} \cap \Delta^+} \|\beta - \mathbf{w}\|_2^2. \quad (9)$$

In words, solving Problem 9 gives the euclidean projection of the vector  $\mathbf{w}$  onto the intersection of the sets  $\Sigma_{k_0}$  and  $\Delta^+$ . We obtain the solution to Problem 9 by employing the Greedy Selector and Simplex Projector (GSSP) algorithm (Kyrillidis et al., 2013).

GSSP proceeds in three steps:

1. **Top- $k$  truncation**  $\text{PL}_k$ . Let  $S \subset \{1, \dots, q\}$  index the  $k$  largest entries of  $\mathbf{w}$ . Define

$$(\text{PL}_k(\mathbf{w}))_i = \begin{cases} w_i, & i \in S, \\ 0, & i \notin S. \end{cases}$$

This retains only the top  $k$  components of  $\mathbf{w}$ .

2. **Simplex projection**  $P^+$ . Restrict  $\mathbf{w}$  to the active entries  $\{w_i\}_{i \in S}$ , sort them into descending order  $w_{(1)} \geq w_{(2)} \geq \dots \geq w_{(k)}$ , and compute

$$\tau = \frac{1}{\rho} \left( \sum_{j=1}^{\rho} w_{(j)} - 1 \right), \quad \rho = \max \left\{ j : w_{(j)} > \frac{1}{j} \left( \sum_{m=1}^j w_{(m)} - 1 \right) \right\}.$$

Then form

$$(P^+(\mathbf{w}))_i = \max(w_i - \tau, 0) \quad \text{for } i \in S,$$

and set  $(P^+(\mathbf{w}))_i = 0$  for  $i \notin S$ . This enforces nonnegativity and the unit  $\ell_1$  norm constraint.

3. **Reconstruction.** The final projected vector is

$$\beta_i = (P^+(\mathbf{w}))_i, \quad i = 1, \dots, q,$$

which has support  $S$  and sums to 1.

By Theorem 1 of Kyrillidis et al. (2013), these steps yield the exact minimizer of Problem 9 in quasilinear time.

### 3.3 Alternating Best Response

We now describe our complete alternating best response procedure used to obtain high quality solutions to the sparse multiple kernel learning problem given by equation 4. Our approach iteratively alternates between optimizing the SVM dual variables  $\boldsymbol{\alpha}$  and the kernel weights  $\boldsymbol{\beta}$ . We will see in Section 5 that this heuristic performs well in practice by providing both strong predictive accuracy and kernel interpretability.

Each iteration  $t$  of Algorithm 1 consists of two steps. First, we fix  $\boldsymbol{\beta} = \boldsymbol{\beta}^{(t-1)}$  and solve the standard SVM dual problem using the kernel matrix  $\mathbf{K}(\boldsymbol{\beta}^{(t-1)}) = \sum_{j=1}^q \beta_j^{(t-1)} \mathbf{K}_j$ . This subproblem is convex and can be solved to global optimality using either SMO-based solvers (e.g., LIBSVM), or quadratic programming solvers (e.g., Gurobi). This produces an updated vector of dual coefficients  $\boldsymbol{\alpha}^{(t)}$ . Next, we fix  $\boldsymbol{\alpha} = \boldsymbol{\alpha}^{(t)}$  and update the kernel weights  $\boldsymbol{\beta}$  by solving Problem 9.

**Initialization Strategies.** Before the first iteration, we initialize the kernel weights  $\boldsymbol{\beta}^{(0)}$  using one of the following strategies:

- *k-Sparse Random Initialization:* A subset  $S \subseteq \{1, \dots, q\}$  with  $|S| = k_0$  is drawn uniformly at random, and  $\beta_j^{(0)} = 1/k_0$  for  $j \in S$ , zero otherwise.
- *Warm Start:* A feasible vector  $\boldsymbol{\beta}^{(0)}$  produced by solving a separate optimization problem (typically a convex lower bound) is used.

**Objective Tracking.** At each iteration, we compute the value of the dual objective associated with the current kernel weights and dual variables. Let  $J^{(t)}$  denote the value of the regularized SVM dual objective at iteration  $t$ , given by

$$J^{(t)} = \sum_{i=1}^n \alpha_i^{(t)} - \frac{1}{2} \sum_{i,j=1}^n y_i y_j \alpha_i^{(t)} \alpha_j^{(t)} \mathbf{K}_{ij}^{(t)} + \lambda \|\boldsymbol{\beta}^{(t)}\|_2^2,$$

where  $\mathbf{K}^{(t)} = \sum_{j=1}^q \beta_j^{(t)} \mathbf{K}_j$  is the kernel matrix induced by the current weights  $\boldsymbol{\beta}^{(t)}$ , and the final term corresponds to the  $\ell_2$  regularization of the kernel combination. This quantity is used to monitor convergence and trigger early termination when progress stalls.

**Stopping Criteria.** The algorithm terminates when either (i) a maximum number of iterations  $T$  is reached, or (ii) the absolute improvement in the dual objective falls below a tolerance  $\epsilon$  for  $M$  consecutive iterations.

---

**Algorithm 1:** Alternating Best Response for Sparse MKL

---

**Data:**  $\{(x_i, y_i)\}_{i=1}^n, \{\mathbf{K}_j\}_{j=1}^q, C, \lambda, k_0, \epsilon, M, T$

**Result:**  $\alpha, \beta$

**Init:** choose  $\beta^{(0)}$ ;  $\text{non\_decrease} \leftarrow 0, \text{obj\_best} \leftarrow \infty$

```

for  $t = 1, \dots, T$  do                                     // max  $T$  iterations
     $\mathbf{K}^{(t)} \leftarrow \sum_j \beta_j^{(t-1)} \mathbf{K}_j$ 
     $\alpha^{(t)} \leftarrow \arg \max \text{SVM dual on } \mathbf{K}^{(t)}$ 
    for  $j = 1, \dots, q$  do
         $d_j \leftarrow (\mathbf{y} \circ \alpha^{(t)})^\top \mathbf{K}_j (\mathbf{y} \circ \alpha^{(t)})$ 
         $w_j \leftarrow d_j / (4\lambda)$ 
     $\beta^{(t)} \leftarrow \text{GSSP}(\{w_j\}, k_0)$ 
    Compute objective  $J^{(t)}$ 
    if  $\text{obj\_best} - J^{(t)} < \epsilon$  then
         $\text{non\_decrease} \leftarrow \text{non\_decrease} + 1$ 
    else
         $\text{non\_decrease} \leftarrow 0$ 
         $\text{obj\_best} \leftarrow J^{(t)}$ 
        Save  $\alpha^{(t)}, \beta^{(t)}$ 
    if  $\text{non\_decrease} \geq M$  then                             // early stopping
        break

```

---

## 4 An Exact Reformulation and Convex Relaxations

In this section, we reformulate Problem 4 exactly as a mixed integer semidefinite optimization problem and derive several convex relaxations that will be used in Section 5 to both produce lower bounds for solutions returned by Algorithm 1 and to warm start it.

### 4.1 A Mixed Integer Semidefinite Reformulation

We begin by exactly reformulating Problem 4 as a mixed integer semidefinite optimization problem. Consider the alpha subproblem given by Problem 7. Its dual is given by:

$$\begin{aligned}
 \min_{\eta \in \mathbb{R}, \sigma \in \mathbb{R}_+^n, \gamma \in \mathbb{R}^n} \quad & C \sum_{i=1}^n \sigma_i + \frac{1}{2} \gamma^\top \left[ \sum_{i=1}^q \beta_i \mathbf{K}_i \right]^\dagger \gamma \\
 \text{s.t.} \quad & 1 - \sigma_i \leq y_i(\eta + \gamma_i) \quad \forall i \in [n], \\
 & \left( \mathbf{I}_n - \left[ \sum_{i=1}^q \beta_i \mathbf{K}_i \right] \left[ \sum_{i=1}^q \beta_i \mathbf{K}_i \right]^\dagger \right) \gamma = 0.
 \end{aligned} \tag{10}$$

Since Problem 10 contains only affine constraints, Slater's condition is trivially satisfied implying that strong duality holds between Problem 7 and Problem 10. Problem 10 is a convex quadratic optimization problem since for any  $\beta$  that is feasible in Problem 4, the matrix  $\left[ \sum_{i=1}^q \beta_i \mathbf{K}_i \right]^\dagger$  will be positive semidefinite. We seek to leverage Problem 10 to replace the inner maximization problem in Problem 4 which would yield a single minimization problem.

Directly substituting the inner maximization of Problem 4 with Problem 10 is not desirable since  $\beta$  would be an optimization variable that appears within matrix pseudo inverse terms and is multiplied by  $\gamma$ , another



optimization variable, in the objective function and in the last constraint. Consequently, we proceed by introducing a semidefinite constraint that allows Problem 10 to be reformulated as a semidefinite optimization problem without products between  $\gamma$  and  $\beta$  or matrix pseudo inverse terms. Consider the optimization problem given by:

$$\begin{aligned} \min_{\eta, \theta \in \mathbb{R}, \sigma \in \mathbb{R}_+^n, \gamma \in \mathbb{R}^n} \quad & C \sum_{i=1}^n \sigma_i + \frac{1}{2} \theta \\ \text{s.t.} \quad & 1 - \sigma_i \leq y_i(\eta + \gamma_i) \quad \forall i \in [n], \\ & \begin{pmatrix} \theta & \gamma^T \\ \gamma & [\sum_{i=1}^q \beta_i \mathbf{K}_i] \end{pmatrix} \succeq 0. \end{aligned} \tag{11}$$

It follows from the Generalized Schur Complement Lemma (see Boyd et al. (1994), Equation 2.41) that problems Problem 10 and Problem 11 achieve the same optimal objective value. We state this equivalence formally in Proposition 1.

**Proposition 1** *Problem 11 is a valid reformulation of Problem 10.*

**Proof** We show that given a feasible solution to Problem 10, we can construct a feasible solution to Problem 11 that achieves the same objective value and that given a feasible solution to Problem 11, we can construct a feasible solution to Problem 10 that achieves an objective that is no greater than the objective value achieved by the feasible solution in Problem 11.

First, consider an arbitrary solution  $(\bar{\eta}, \bar{\sigma}, \bar{\gamma})$  that is feasible to Problem 10. Let

$$\tilde{\theta} := \bar{\gamma}^T \left[ \sum_{i=1}^q \beta_i \mathbf{K}_i \right]^\dagger \bar{\gamma}.$$

Clearly,  $(\bar{\eta}, \tilde{\theta}, \bar{\sigma}, \bar{\gamma})$  achieves the same objective value in Problem 11 as  $(\bar{\eta}, \bar{\sigma}, \bar{\gamma})$  achieves in Problem 10. To see that  $(\bar{\eta}, \tilde{\theta}, \bar{\sigma}, \bar{\gamma})$  is feasible to Problem 11, we need only verify that semidefinite constraint in Problem 11 is satisfied. Recall that for any  $\beta$  feasible in Problem 4, the matrix  $\sum_{i=1}^q \beta_i \mathbf{K}_i$  will be positive semidefinite. Furthermore, from the definition of  $\tilde{\theta}$  it follows immediately that  $\tilde{\theta} - \bar{\gamma}^T [\sum_{i=1}^q \beta_i \mathbf{K}_i]^\dagger \bar{\gamma} \succeq 0$ , and from the feasibility of  $\bar{\gamma}$  in Problem 10 we have  $(\mathbf{I}_n - [\sum_{i=1}^q \beta_i \mathbf{K}_i] [\sum_{i=1}^q \beta_i \mathbf{K}_i]^\dagger) \bar{\gamma} = 0$ . Therefore, by the Generalized Schur Complement Lemma, we have  $\begin{pmatrix} \tilde{\theta} & \bar{\gamma}^T \\ \bar{\gamma} & [\sum_{i=1}^q \beta_i \mathbf{K}_i] \end{pmatrix} \succeq 0$ .

Now, consider an arbitrary solution  $(\bar{\eta}, \bar{\theta}, \bar{\sigma}, \bar{\gamma})$  that is feasible to Problem 11. We claim that  $(\bar{\eta}, \bar{\sigma}, \bar{\gamma})$  is feasible in Problem 10 and achieves an objective value no greater than the objective value achieved by  $(\bar{\eta}, \bar{\theta}, \bar{\sigma}, \bar{\gamma})$  in Problem 11. To see this, it suffices to note that from the Generalized Schur Complement Lemma, feasibility of  $(\bar{\eta}, \bar{\theta}, \bar{\sigma}, \bar{\gamma})$  in Problem 11 implies that we have  $(\mathbf{I}_n - [\sum_{i=1}^q \beta_i \mathbf{K}_i] [\sum_{i=1}^q \beta_i \mathbf{K}_i]^\dagger) \bar{\gamma} = 0$  and that  $\bar{\theta} \geq \bar{\gamma}^T [\sum_{i=1}^q \beta_i \mathbf{K}_i]^\dagger \bar{\gamma}$ . This completes the proof. ■

We can now obtain a more tractable reformulation of Problem 4 by substituting its inner maximization problem with the equivalent minimization Problem 11. Consider the optimization problem given by:

$$\begin{aligned}
& \min_{\substack{\eta, \theta \in \mathbb{R}, \sigma \in \mathbb{R}_+^n, \\ \gamma \in \mathbb{R}^n, \beta \in \mathbb{R}_+^q}} C \sum_{i=1}^n \sigma_i + \frac{1}{2} \theta + \lambda \|\beta\|_2^2, \\
& \text{s.t.} \quad 1 - \sigma_i \leq y_i(\eta + \gamma_i) \quad \forall i \in [n], \\
& \quad \begin{pmatrix} \theta & \gamma^T \\ \gamma & [\sum_{i=1}^q \beta_i \mathbf{K}_i] \end{pmatrix} \succeq 0, \\
& \quad \sum_{i=1}^q \beta_i = 1, \|\beta\|_0 \leq k.
\end{aligned} \tag{12}$$

Problem 12 consists of a non-convex minimization problem where the nonconvexity is entirely captured by the cardinality constraint  $\|\beta\|_0 \leq k$ . We further reformulate Problem 12 by introducing binary variables  $\mathbf{z} \in \{0, 1\}^q$  and imposing the nonlinear constraint  $\beta_i = z_i \beta_i$  to model the sparsity pattern of  $\beta$ . Explicitly, we characterize the set of  $q$ -dimensional vectors with cardinality at most  $k$  using the following equivalence:

$$\{\beta \in \mathbb{R}^q : \|\beta\|_0 \leq k\} = \{\beta \in \mathbb{R}^q : \exists \mathbf{z} \in \{0, 1\}^q, \sum_{i=1}^q z_i \leq k, \beta_i = z_i \beta_i \forall i\}.$$

As  $\mathbf{z}$  and  $\beta$  will both be optimization variables in our relaxation, the constraints  $\beta_i = z_i \beta_i$  are complicating as they are non-convex in the decision variables. Accordingly, we invoke the perspective reformulation (Günlük & Linderoth, 2012) to model these constraints in a convex manner. Specifically, we introduce a variable  $\omega \in \mathbb{R}^n$  where  $\omega_i$  models  $\beta_i^2$  and we introduce constraints  $\omega_i z_i \geq \beta_i^2$  which are second-order cone representable. This results in the following exact reformulation of Problem 4 and Problem 12:

$$\begin{aligned}
& \min_{\substack{\eta, \theta \in \mathbb{R}, \sigma \in \mathbb{R}_+^n, \gamma \in \mathbb{R}^n, \\ \beta, \omega \in \mathbb{R}_+^q, \mathbf{z} \in \{0, 1\}^q}} C \sum_{i=1}^n \sigma_i + \frac{1}{2} \theta + \lambda \sum_{i=1}^q \omega_i, \\
& \text{s.t.} \quad 1 - \sigma_i \leq y_i(\eta + \gamma_i) \quad \forall i \in [n], \\
& \quad \begin{pmatrix} \theta & \gamma^T \\ \gamma & [\sum_{i=1}^q \beta_i \mathbf{K}_i] \end{pmatrix} \succeq 0, \\
& \quad \sum_{i=1}^q \beta_i = 1, \quad \sum_{i=1}^q z_i \leq k, \quad \beta_i^2 \leq z_i \omega_i \quad \forall i \in [q].
\end{aligned} \tag{13}$$

## 4.2 Positive Semidefinite Cone Relaxations

Problem 13 is an exact mixed integer semidefinite optimization reformulation of Problem 4 where the problem's non-convexity is entirely captured through the binary condition  $\mathbf{z} \in \{0, 1\}^q$ . We now obtain a semidefinite relaxation by solving Problem 13 with  $\mathbf{z} \in \text{conv}(\{0, 1\}^q) = [0, 1]^q$ . This yields the following convex optimization problem, which we refer to as the *Full SDP* relaxation:

$$\begin{aligned}
& \min_{\substack{\eta, \theta \in \mathbb{R}, \sigma \in \mathbb{R}_+^n, \gamma \in \mathbb{R}^n, \\ \beta, \omega \in \mathbb{R}_+^q, \mathbf{z} \in [0, 1]^q}} C \sum_{i=1}^n \sigma_i + \frac{1}{2} \theta + \lambda \sum_{i=1}^q \omega_i, \\
& \text{s.t.} \quad 1 - \sigma_i \leq y_i(\eta + \gamma_i) \quad \forall i \in [n], \\
& \quad \begin{pmatrix} \theta & \gamma^T \\ \gamma & [\sum_{i=1}^q \beta_i \mathbf{K}_i] \end{pmatrix} \succeq 0, \\
& \quad \sum_{i=1}^q \beta_i = 1, \quad \sum_{i=1}^q z_i \leq k, \quad \beta_i^2 \leq z_i \omega_i \quad \forall i \in [q].
\end{aligned} \tag{14}$$

**Theorem 2** *Problem 14 is a valid convex relaxation of Problem 4.*

**Proof** By the duality between the alpha subproblem of Problem 4 (given by Problem 7) and Problem 10 coupled with Proposition 1, it suffices to argue that Problem 14 is a valid convex relaxation of Problem 12 to establish Theorem 2. Clearly, Problem 14 is a convex optimization problem. We will show that given any feasible solution to Problem 12, we can construct a feasible solution to Problem 14 that achieves the same objective value.

Let  $(\bar{\eta}, \bar{\theta}, \bar{\sigma}, \bar{\gamma}, \bar{\beta}) \in \mathbb{R} \times \mathbb{R} \times \mathbb{R}_+^n \times \mathbb{R}^n \times \mathbb{R}_+^q$  denote an arbitrary feasible solution to Problem 12. Define  $\tilde{\omega}_i = \bar{\beta}_i^2$  and  $\tilde{z}_i = 1$  if  $\bar{\beta}_i > 0$ , otherwise  $\tilde{z}_i = 0$  for all  $i \in [q]$ . Observe that  $(\bar{\eta}, \bar{\theta}, \bar{\sigma}, \bar{\gamma}, \bar{\beta}, \tilde{\omega}, \tilde{z})$  is feasible to Problem 14. To see this, note that we have  $\tilde{z}_i \tilde{\omega}_i = \bar{\beta}_i^2$  and  $\sum_{i=1}^q \tilde{z}_i = \|\bar{\beta}\|_0 \leq k$ . Moreover, we have  $\sum_{i=1}^q \tilde{\omega}_i = \sum_{i=1}^q \bar{\beta}_i^2 = \|\bar{\beta}\|_2^2$  which shows that  $(\bar{\eta}, \bar{\theta}, \bar{\sigma}, \bar{\gamma}, \bar{\beta}, \tilde{\omega}, \tilde{z})$  achieves the same objective value in Problem 14 as  $(\bar{\eta}, \bar{\theta}, \bar{\sigma}, \bar{\gamma}, \bar{\beta})$  achieves in Problem 12. This completes the proof. ■

In Section 5, we use Problem 14 to produce lower bounds for feasible solutions returned by Algorithm 1. Note that Algorithm 1 and Problem 14 can be leveraged to design a custom branch and bound algorithm in the sense of Little (1966); Land & Doig (2010) to solve Problem 4 to global optimality. Such an approach would involve constructing an enumeration tree that branches on the entries of the vector  $\mathbf{z}$ . Each node in the enumeration tree would be defined by two disjoint collections of indices  $\mathcal{I}_0, \mathcal{I}_1 \subseteq [q]$  where  $\mathcal{I}_i$  denotes indices of  $\mathbf{z}$  constrained to take value  $i$ . At a given node, feasible solutions and lower bounds would be computed using slightly modified versions of Algorithm 1 and Problem 14 respectively.

As a brief aside, we note that in the special case where all of the input kernels  $\{\mathbf{K}_i\}_{i=1}^q$  are simultaneously diagonalizable, Problem 14 can be reduced to a second order cone problem. Explicitly, consider the following optimization problem:

$$\begin{aligned}
& \min_{\substack{\eta, \theta \in \mathbb{R}, \sigma \in \mathbb{R}_+^n, \gamma, \tau \in \mathbb{R}^n, \\ \beta, \omega \in \mathbb{R}_+^q, \mathbf{z} \in [0,1]^q}} C \sum_{i=1}^n \sigma_i + \frac{1}{2} \theta + \lambda \sum_{i=1}^q \omega_i, \\
& \text{s.t.} \quad 1 - \sigma_i \leq y_i(\eta + \gamma_i) \quad \forall i \in [n], \\
& \quad \tau_j \cdot \sum_{i=1}^q \beta_i [\mathbf{D}_i]_{jj} \geq (\gamma^T \mathbf{u}_j)^2 \quad \forall j \in [n], \\
& \quad \theta \geq \sum_{j=1}^n \tau_j, \quad \sum_{i=1}^q \beta_i = 1, \quad \sum_{i=1}^q z_i \leq k, \quad \beta_i^2 \leq z_i \omega_i \quad \forall i \in [q],
\end{aligned} \tag{15}$$

where  $\mathbf{K}_i = \mathbf{U} \mathbf{D}_i \mathbf{U}^T$  for all  $i \in [q]$  are spectral decompositions of the candidate kernels which are assumed to be simultaneously diagonalizable. Note that  $\mathbf{u}_j \in \mathbb{R}^n$  denotes the  $j^{\text{th}}$  column of  $\mathbf{U} \in \mathbb{R}^{n \times n}$ . We now have the following result:

**Theorem 3** *If the candidate kernels  $\{\mathbf{K}_i\}_{i=1}^q$  are simultaneously diagonalizable, meaning that we have  $\mathbf{K}_i = \mathbf{U} \mathbf{D}_i \mathbf{U}^T$  for all  $i \in [q]$  where  $\mathbf{D}_i \in \mathbb{R}^{n \times n}$  are diagonal matrices and  $\mathbf{U} \in \mathbb{R}^{n \times n}$  satisfies  $\mathbf{U} \mathbf{U}^T = \mathbf{U}^T \mathbf{U} = \mathbf{I}_n$ , then Problem 14 is equivalent to Problem 15.*

**Proof**

Define the sets  $\mathcal{A}$  and  $\mathcal{B}$  as:

$$\begin{aligned}
\mathcal{A} &:= \left\{ (\theta, \gamma, \beta) \in \mathbb{R} \times \mathbb{R}^n \times \mathbb{R}_+^q : \begin{pmatrix} \theta & \gamma^T \\ \gamma & [\sum_{i=1}^q \beta_i \mathbf{K}_i] \end{pmatrix} \succeq 0 \right\}, \\
\mathcal{B} &:= \left\{ (\theta, \gamma, \beta) \in \mathbb{R} \times \mathbb{R}^n \times \mathbb{R}_+^q : \exists \tau \in \mathbb{R}^n \mid \theta \geq \sum_{j=1}^n \tau_j, \tau_j \cdot \sum_{i=1}^q \beta_i [\mathbf{D}_i]_{jj} \geq (\gamma^T \mathbf{u}_j)^2 \quad \forall j \right\}.
\end{aligned}$$

It suffices to show that we have  $\mathcal{A} = \mathcal{B}$ . Recall that for any  $\beta \in \mathbb{R}_+^q$ , the combined kernel  $\sum_{i=1}^q \beta_i \mathbf{K}_i$  will be positive semidefinite so by the generalized schur complement the matrix  $\begin{pmatrix} \theta & \gamma^T \\ \gamma & [\sum_{i=1}^q \beta_i \mathbf{K}_i] \end{pmatrix}$  will be positive semidefinite if and only if we have  $\theta \geq \gamma^T [\sum_{i=1}^q \beta_i \mathbf{K}_i]^\dagger \gamma$ . Since the candidate kernels are simultaneously diagonalizable, we can express the pseudo inverse of the combined kernel as:

$$\left[ \sum_{i=1}^q \beta_i \mathbf{K}_i \right]^\dagger = \left[ \mathbf{U} \left[ \sum_{i=1}^q \beta_i \mathbf{D}_i \right] \mathbf{U}^T \right]^\dagger = \mathbf{U} \mathbf{\Lambda} \mathbf{U}^T,$$

where  $\mathbf{\Lambda} \in \mathbb{R}^{n \times n}$  is a diagonal matrix with diagonal entries defined as  $\Lambda_{jj} = \frac{1}{\sum_{i=1}^q \beta_i [\mathbf{D}_i]_{jj}}$  if  $\sum_{i=1}^q \beta_i [\mathbf{D}_i]_{jj} > 0$ , otherwise  $\Lambda_{jj} = 0$ . Thus, we have

$$\gamma^T \left[ \sum_{i=1}^q \beta_i \mathbf{K}_i \right]^\dagger \gamma = \gamma^T \mathbf{U} \mathbf{\Lambda} \mathbf{U}^T \gamma = \sum_{j=1}^n \Lambda_{jj} (\gamma^T \mathbf{u}_j)^2.$$

Therefore, the condition  $\theta \geq \gamma^T [\sum_{i=1}^q \beta_i \mathbf{K}_i]^\dagger \gamma$  is equivalent to the existence of a vector  $\boldsymbol{\tau} \in \mathbb{R}^n$  such that  $\theta \geq \sum_{j=1}^n \tau_j$  and  $\tau_j \geq \Lambda_{jj} (\gamma^T \mathbf{u}_j)^2$  for each  $j$ . Finally, recalling the definition of  $\mathbf{\Lambda}$ , we can rewrite the inequality  $\tau_j \geq \Lambda_{jj} (\gamma^T \mathbf{u}_j)^2$  as  $\tau_j \cdot \sum_{i=1}^q \beta_i [\mathbf{D}_i]_{jj} \geq (\gamma^T \mathbf{u}_j)^2$ . This completes the proof. ■

#### 4.2.1 Sparse Positive Semidefinite Cone Approximations

We will see in Section 5 that Problem 14 empirically produces strong lower bounds. However, solving Problem 14 becomes computationally difficult as the number of training examples  $n$  grows because even the most efficient interior point solvers for semidefinite optimization problems exhibit poor scaling with the size of the positive semidefinite constraint which in the case of Problem 14 has dimension  $(n+1) \times (n+1)$ . In order to produce lower bounds efficiently for large problem instances, we are interested in developing approximations to Problem 14 that can be computed more efficiently as the problem dimension scales.

To this end, we leverage a common approximation approach that involves relaxing the global positive semidefinite constraint and instead, for a given value  $\ell \in \mathbb{N}$ , constraining all  $\ell \times \ell$  principal submatrices to be positive semidefinite. This technique is often called the sparse SDP relaxation and is known to have desirable approximation properties in both theory and practice (Blekherman et al., 2022; Song & Parrilo, 2023). Recall that  $\mathbb{S}_+^n$  denotes the set of all  $n \times n$  real symmetric positive semidefinite matrices and let  $\mathbb{S}_\ell^n$  denote the set of all  $n \times n$  real symmetric matrices where all  $\ell \times \ell$  principal submatrices are positive semidefinite. It is clear that for any  $\ell \in \mathbb{N}$ , we have  $\mathbb{S}_+^n \subseteq \mathbb{S}_{\ell+1}^n \subseteq \mathbb{S}_\ell^n$ .

In general, for a matrix  $\mathbf{X} \in \mathbb{R}^{(n+1) \times (n+1)}$ , substituting the constraint  $\mathbf{X} \in \mathbb{S}_+^{n+1}$  by  $\mathbf{X} \in \mathbb{S}_\ell^{n+1}$  would involve replacing a  $(n+1) \times (n+1)$  semidefinite constraint with  $\binom{n+1}{\ell} \ell \times \ell$  semidefinite constraints. However, recall that the  $n \times n$  sub block  $\sum_{i=1}^q \beta_i \mathbf{K}_i$  of the  $(n+1) \times (n+1)$  matrix involved in the semidefinite constraint of Problem 14 is itself guaranteed to be positive semidefinite for any feasible  $\beta$ . Accordingly, any principal submatrix that is entirely contained within this sub block is guaranteed to be positive semidefinite which implies we need only concern ourselves with principal submatrices that involve the first row and column when defining our approximation. Thus, letting  $\mathbf{X} = \begin{pmatrix} \theta & \gamma^T \\ \gamma & [\sum_{i=1}^q \beta_i \mathbf{K}_i] \end{pmatrix}$ , imposing the constraint  $\mathbf{X} \in \mathbb{S}_\ell^{n+1}$  requires only  $\binom{n}{\ell-1} \ell \times \ell$  semidefinite constraints rather than  $\binom{n+1}{\ell} \ell \times \ell$  many. This equates to using  $\binom{n}{\ell-1} \times \binom{n}{\ell} - 1$  fewer constraints.

As we report results for  $\ell = 2$  and  $\ell = 3$  in Section 5, we explicitly state the corresponding relaxations below. For  $\ell = 2$ , the approximation to Problem 14 is given by:

$$\begin{aligned}
& \min_{\substack{\eta, \theta \in \mathbb{R}, \sigma \in \mathbb{R}_+^n, \gamma \in \mathbb{R}^n, \\ \beta, \omega \in \mathbb{R}_+^q, \mathbf{z} \in [0,1]^q}} C \sum_{i=1}^n \sigma_i + \frac{1}{2} \theta + \lambda \sum_{i=1}^q \omega_i, \\
& \text{s.t.} \quad 1 - \sigma_i \leq y_i(\eta + \gamma_i) \quad \forall i \in [n], \\
& \quad \theta \cdot \sum_{i=1}^q \beta_i [\mathbf{K}_i]_{jj} \geq \gamma_j^2 \quad \forall j \in [n], \\
& \quad \sum_{i=1}^q \beta_i = 1, \quad \sum_{i=1}^q z_i \leq k, \quad \beta_i^2 \leq z_i \omega_i \quad \forall i \in [q].
\end{aligned} \tag{16}$$

Problem 16 is a second order cone relaxation of Problem 14. In Section 5, we refer to this relaxation as *Second Order Cone Relaxation with basis vectors*. Moreover, we can strengthen this second order cone relaxation by randomly sampling unit vectors  $\mathbf{x} \in \mathbb{R}^n$  and imposing the additional second order cone representable constraints given by  $\theta \cdot \sum_{i=1}^q \beta_i (\mathbf{x}^T \mathbf{K}_i \mathbf{x}) \geq (\mathbf{x}^T \boldsymbol{\gamma})^2$  for each sampled unit vector  $\mathbf{x}$ . To see why these constraints are valid for Problem 14, notice that by the schur complement lemma the semidefinite constraint in Problem 14 implies that the matrix  $\sum_{i=1}^q \beta_i \mathbf{K}_i - \frac{1}{\theta} \boldsymbol{\gamma} \boldsymbol{\gamma}^T$  must also be positive semidefinite. This condition can equivalently be expressed as  $\theta \cdot \sum_{i=1}^q \beta_i \mathbf{K}_i \succeq \boldsymbol{\gamma} \boldsymbol{\gamma}^T$ . Thus, for any vector  $\mathbf{x}$  (whether a unit vector or not), we have  $\theta \cdot \sum_{i=1}^q \beta_i (\mathbf{x}^T \mathbf{K}_i \mathbf{x}) \geq (\mathbf{x}^T \boldsymbol{\gamma})^2$ . Notice that if we generate these constraints letting  $\mathbf{x}$  be the standard basis vectors, we obtain the second order cone constraints in Problem 16. Explicitly, let  $\mathcal{L}$  denote a collection of randomly sampled unit vectors. Consider the optimization problem given by:

$$\begin{aligned}
& \min_{\substack{\eta, \theta \in \mathbb{R}, \sigma \in \mathbb{R}_+^n, \gamma \in \mathbb{R}^n, \\ \beta, \omega \in \mathbb{R}_+^q, \mathbf{z} \in [0,1]^q}} C \sum_{i=1}^n \sigma_i + \frac{1}{2} \theta + \lambda \sum_{i=1}^q \omega_i, \\
& \text{s.t.} \quad 1 - \sigma_i \leq y_i(\eta + \gamma_i) \quad \forall i \in [n], \\
& \quad \theta \cdot \sum_{i=1}^q \beta_i \mathbf{x}^T \mathbf{K}_i \mathbf{x} \geq (\mathbf{x}^T \boldsymbol{\gamma})^2 \quad \forall \mathbf{x} \in \mathcal{L} \cup \{\mathbf{e}_j\}_{j=1}^n, \\
& \quad \sum_{i=1}^q \beta_i = 1, \quad \sum_{i=1}^q z_i \leq k, \quad \beta_i^2 \leq z_i \omega_i \quad \forall i \in [q].
\end{aligned} \tag{17}$$

Problem 17 is a second order cone relaxation of Problem 14. In Section 5, we refer to this relaxation as *Second Order Cone Relaxation with randomized vectors*. Moreover it should be clear that Problem 17 is a stronger relaxation than Problem 16 since the feasible set of Problem 17 is a subset of that of Problem 16. Moreover, Problem 17 becomes an increasingly strong approximation to Problem 14 as the cardinality of  $\mathcal{L}$  increases. We will explore this further in Section 5. The final relaxation we introduce is the sparse SDP relaxation with  $\ell = 3$  for Problem 14 which is given by:

$$\begin{aligned}
& \min_{\substack{\eta, \theta \in \mathbb{R}, \sigma \in \mathbb{R}_+^n, \gamma \in \mathbb{R}^n, \\ \beta, \omega \in \mathbb{R}_+^q, \mathbf{z} \in [0,1]^q}} C \sum_{i=1}^n \sigma_i + \frac{1}{2} \theta + \lambda \sum_{i=1}^q \omega_i, \\
& \text{s.t.} \quad 1 - \sigma_i \leq y_i(\eta + \gamma_i) \quad \forall i \in [n], \\
& \quad \begin{pmatrix} \theta & \gamma_j & \gamma_k \\ \gamma_j & \sum_{i=1}^q \beta_i [\mathbf{K}_i]_{jj} & \sum_{i=1}^q \beta_i [\mathbf{K}_i]_{jk} \\ \gamma_k & \sum_{i=1}^q \beta_i [\mathbf{K}_i]_{kj} & \sum_{i=1}^q \beta_i [\mathbf{K}_i]_{kk} \end{pmatrix} \succeq 0 \quad \forall 1 \leq j, k \leq n, j \neq k, \\
& \quad \sum_{i=1}^q \beta_i = 1, \quad \sum_{i=1}^q z_i \leq k, \quad \beta_i^2 \leq z_i \omega_i \quad \forall i \in [q].
\end{aligned} \tag{18}$$

In Section 5, we refer to this relaxation with  $\ell = 3$  as *3x3-SDP* relaxation.

## 5 Computational Results

In this section, we evaluate the performance of Algorithm 1 and the relaxations introduced in Section 4. We benchmark against open source methods implemented in `MKLpy` library (Lauriola & Aioli, 2020) that we introduced in Section 2.1.2. All experiments were run using Julia 1.10.1 and Python 3.10.14; semidefinite programs were solved with `MOSEK 10.1.31`, while all other optimization relied on Julia’s `LIBSVM v0.8.0` and `MKLpy 0.6`. We perform experiments on MIT’s Supercloud Cluster (Reuther et al., 2018), which hosts Intel Xeon Platinum 8260 processors. To bridge the gap between theory and practice, we have made our code freely available at <https://github.com/XXXX/SparseMKL>.

We seek to answer three questions:

- Q1** *Predictive accuracy and runtime.* How does Algorithm 1 compare to other widely used open-source MKL solvers—EasyMKL, AverageMKL, and CKA—implemented in `MKLpy` (Lauriola & Aioli, 2020), in terms of prediction accuracy and computational efficiency?
- Q2** *Warm starts.* Does initializing Algorithm 1 with solutions from our convex relaxations enhance the quality of the returned solution?
- Q3** *Optimality.* How close is the objective value achieved by the output of Algorithm 1 to the lower bounds produced by the SDP relaxations of Section 4?

We address Q1–Q3 in the following sections.

### 5.1 Experimental design

**Data sets.** We evaluate all methods on ten benchmark binary classification tasks drawn from the UCI Machine Learning Repository (Dua & Graff, 2017). Seven of these—BREASTCANCER (Wisconsin Diagnostic), IONOSPHERE, SPAMBASE, BANKNOTE, HABERMAN, MAMMOGRAPHIC, and PARKINSONS—are provided in their original binary form. The remaining three—IRIS, WINE, and HEART (Cleveland)—are originally multiclass and were converted into binary tasks.

Each data set was parsed into numeric features, with any categorical variables one-hot encoded while omitting one dimension to avoid collinearity. To ensure reproducibility, we fixed the random seed, shuffled the rows, and split 80% of the examples into a training set and the remaining 20% into a test set. Numeric features were standardized to zero mean and unit variance using statistics computed solely on the training set; the same centering and scaling parameters were then applied to the test set. Table 1 summarizes, for each task, the original number of examples, the train set size  $n_{\text{tr}}$ , the test set size  $n_{\text{te}}$ , and the proportion of positive labels.

Table 1: Summary of binary classification tasks used in our experiments. The table reports, for each dataset, the number of examples  $n$ , the training size  $n_{\text{tr}}$ , the test size  $n_{\text{te}}$ , and the proportion of positive labels.

Dataset	$n$	$n_{\text{tr}}$	$n_{\text{te}}$	% Positive (+1)
IRIS	150	120	30	33.3%
WINE	178	142	36	33.1%
BREASTCANCER	569	455	114	37.3%
IONOSPHERE	351	280	71	64.1%
SPAMBASE	4601	3680	921	39.4%
BANKNOTE	1372	1097	275	44.5%
HEART	303	242	61	45.9%
HABERMAN	306	244	62	73.5%
MAMMOGRAPHIC	961	768	193	46.3%
PARKINSONS	195	156	39	75.4%

**Set of Base Kernels.** All methods are evaluated on the same predefined set of ten base kernels. Our goal in assembling this collection was to include a diverse yet representative sample of the similarity measures introduced in Section 2.2, so as to provide each model with flexibility across different feature space geometries.

Concretely, our set comprises a linear kernel (to capture simple inner-product structure), three polynomial kernels of degree  $d \in \{2, 3, 5\}$  with offset  $c_0 = 1$  and  $\alpha = 0.01$ , and three Gaussian radial-basis-function (RBF) kernels with parameters  $\gamma \in \{0.5, 0.3, 0.1\}$ . We further include two sigmoid kernels, each with offset  $c_0 = 1$  and slopes  $\gamma \in \{0.5, 0.7\}$ , and a single Laplacian kernel with  $\gamma = 0.3$ .

Each kernel matrix is precomputed once and then symmetrized via the transformation  $\mathbf{K} \leftarrow \frac{1}{2}(\mathbf{K} + \mathbf{K}^\top)$  to guard against numerical asymmetry and guarantee positive semidefiniteness. Finally, we add  $10^{-6} \cdot \mathbf{I}$  to every kernel matrix to improve conditioning and ensure stable solver performance.

**Hyperparameter tuning.** For Algorithm 1, we perform 10-fold cross-validation over the training set to select the regularization and sparsity parameters from the grid

$$C \in \{5, 10, 50, 100\}, \quad \lambda \in \{0.01, 0.1, 1, 10, 100\}, \quad k_0 \in \{1, 2, 3, 4, 5\}.$$

The same folds are reused to tune EasyMKL’s margin-variance trade-off parameter  $\lambda_{\text{Easy}}$  over 25 values in  $[10^{-4}, 1]$ . AverageMKL and CKA require no parameter tuning. Once the best hyperparameters are identified, each model is retrained on the full training split and evaluated on the withheld test set.

**Performance metrics.** Classification quality is measured by test accuracy (Table 2); solution quality by the MKL objective value (Table 6); optimality by the percentage gap with respect to the tightest available lower bound (Table 7); and computational burden by the training time (Table 4).

## 5.2 Predictive Accuracy and Runtime Compared to Open-Source MKL Solvers (Q1)

We begin our computational evaluation by assessing the predictive performance and runtime of Algorithm 1 in comparison with existing open-source MKL solvers.

Table 2 reports the test set accuracy achieved by EasyMKL, AverageMKL, CKA, and Algorithm 1 with random initialization. We find that Algorithm 1 matches or outperforms the best open-source baseline on all tasks: it strictly improves over the strongest competitor in eight cases and ties in the remaining two. The largest absolute gains occur on IONOSPHERE (+7.1 percentage points), HEART (+6.5 points), and PARKINSONS (+7.6 points), with additional consistent improvements observed on WINE (+2.8 points), BREASTCANCER (+3.6 points), HABERMAN (+1.6 points), and MAMMOGRAPHIC (+3.7 points). On the perfectly separable tasks IRIS and BANKNOTE, Algorithm 1 attains 100% test accuracy, matching the best-performing solvers. Averaged across datasets, Algorithm 1 outperforms the strongest baseline by 3.34 percentage points.

Dataset	Train Size	EasyMKL	AverageMKL	CKA	Algorithm 1
iris	120	<b>100.0</b>	<b>100.0</b>	96.7	<b>100.0</b>
wine	142	97.2	97.2	91.7	<b>100.0</b>
breastcancer	455	93.0	92.1	94.7	<b>98.3</b>
ionosphere	280	73.2	74.6	85.9	<b>93.0</b>
spambase	3680	90.4	87.6	81.0	<b>90.9</b>
banknote	1097	<b>100.0</b>	<b>100.0</b>	85.1	<b>100.0</b>
heart	242	85.2	85.2	86.9	<b>93.4</b>
haberman	244	61.3	62.9	66.1	<b>67.7</b>
mammographic	768	80.8	79.3	75.1	<b>84.5</b>
parkinsons	156	82.1	82.1	74.4	<b>89.7</b>

Table 2: Test accuracy (%) of EasyMKL, AverageMKL, CKA, and Algorithm 1 with random initialization.

Dataset	EasyMKL	AverageMKL	CKA	Algorithm 1
iris	10	10	9	<b>1</b>
wine	4	10	7	<b>2</b>
breastcancer	3	10	7	<b>2</b>
ionosphere	3	10	6	<b>1</b>
spambase	10	10	7	<b>5</b>
banknote	10	10	7	<b>1</b>
heart	<b>3</b>	10	7	4
haberman	8	10	8	<b>1</b>
mammographic	7	10	7	<b>1</b>
parkinsons	<b>2</b>	10	7	<b>2</b>

Table 3: Number of non-zero coefficients ( $|\beta_j| > 10^{-3}$ ) for each algorithm on every dataset.

Crucially, these gains are achieved while selecting sparse kernel combinations as illustrated by Table 3. In contrast, AverageMKL uniformly assigns nonzero weight to all kernels, EasyMKL encourages sparsity to some extent via regularization but does not allow explicit control over the sparsity level; and CKA typically yields dense combinations of all kernels. We stress that the randomly initialized variant of Algorithm 1 achieves these improvements without solving any semidefinite programming (SDP) relaxations and in only a modest number of alternating best response steps. Consequently, the runtime of Algorithm 1 is competitive with standard MKL solvers as illustrated in Table 4.

Dataset	Train Size	EasyMKL	AverageMKL	CKA	Algorithm 1
iris	120	0.030	<b>0.002</b>	0.026	0.022
wine	142	0.045	<b>0.002</b>	0.029	0.005
breastcancer	455	0.244	<b>0.004</b>	0.061	0.110
ionosphere	280	0.072	<b>0.003</b>	0.036	0.070
spambase	3680	33.671	<b>0.631</b>	8.859	8.330
banknote	1097	0.337	<b>0.006</b>	0.301	0.519
heart	242	0.062	<b>0.003</b>	0.035	0.188
haberman	244	0.029	1.907	0.033	<b>0.025</b>
mammographic	768	0.228	13.027	<b>0.125</b>	0.436
parkinsons	156	0.038	<b>0.002</b>	0.028	0.245

Table 4: Training time (seconds) of EasyMKL, AverageMKL, CKA, and Algorithm 1 with random initialization. The fastest runtime in each row is bolded.

It should come as no surprise that AverageMKL has the fastest runtime in eight of the ten benchmarks since this method just takes a naive average of the input kernels. Among the three weight-learning algorithms, performance is more nuanced. CKA posts the shortest times on a majority of datasets, yet its advantage over Algorithm 1 seldom exceeds a few hundredths of a second and disappears entirely on the largest workload: on SPAMBASE (3,680 data points) Algorithm 1 is the fastest learner, running 6% faster than CKA and over four times faster than EasyMKL. Algorithm 1 also takes the lead on IRIS, WINE, and HABERMAN. EasyMKL offers savings relative to Algorithm 1 on some mid-sized tasks (e.g., PARKINSONS, MAMMOGRAPHIC), but it is the slowest option on the largest problem.

These results highlight Algorithm 1’s favorable trade-off between predictive accuracy, model sparsity, and computational efficiency in the case of random initialization. This variant of Algorithm 1 suffices to achieve state-of-the-art performance across a broad range of benchmarks, offering strong predictive accuracy, interpretable sparsity, and practical training times.



### 5.3 Effect of Warm Starts (Q2)

We next evaluate the practical benefits of initializing Algorithm 1 with solutions of SDP relaxations. We consider three variants:

- **Random:** Employ the k-sparse random initialization technique described in Section 3.3;
- **$3 \times 3$ -SDP:** Solve our sparse semidefinite relaxation for  $\ell = 3$ , given by Problem 18. Select the  $k_0$  largest entries of  $\mathbf{z}$  as the support of  $\beta$  and scale  $\beta$  (restricted to these entries) to have unit  $\ell_1$  norm;
- **Full SDP:** Solve the dense semidefinite relaxation give by Problem 14. Select the  $k_0$  largest entries of  $\mathbf{z}$  as the support of  $\beta$  and scale  $\beta$  (restricted to these entries) to have unit  $\ell_1$  norm.

Table 5 reports test set accuracy for these variants. Warm starts occasionally yield improvements, but their effect is heterogeneous across datasets. In some cases, warm initialization produces significant gains: on PARKINSONS, the full-warm start improves accuracy from 89.7% to 94.9%, and on IONOSPHERE, the  $3 \times 3$ -warm start increases accuracy from 93.0% to 94.4%.

Dataset	Train Size	Best Baseline (max of benchmarks)	Algorithm 1 Variants		
			(random)	(warm $3 \times 3$ )	(warm full)
iris	120	<b>100.0</b>	<b>100.0</b>	<b>100.0</b>	<b>100.0</b>
wine	142	97.2	<b>100.0</b>	<b>100.0</b>	<b>100.0</b>
breastcancer	455	94.7	<b>98.3</b>	97.4	—
ionosphere	280	85.9	93.0	<b>94.4</b>	90.1
spambase	3680	90.4	90.9	<b>93.9<sup>†</sup></b>	—
banknote	1097	<b>100.0</b>	<b>100.0</b>	<b>100.0</b>	—
heart	242	86.9	<b>93.4</b>	<b>93.4</b>	<b>93.4</b>
haberman	244	66.1	<b>67.7</b>	66.1	66.1
mammographic	768	80.8	<b>84.5</b>	<b>84.5</b>	—
parkinsons	156	82.1	89.7	89.7	<b>94.9</b>

Table 5: Best baseline accuracy (highest of EasyMKL, AverageMKL, CKA) versus Algorithm 1 variants. “—” denotes unavailable results due to memory limits. <sup>†</sup>For SPAMBASE, the  $3 \times 3$  warm start uses the randomized SOC relaxation due to memory limits.

In other cases, however, the benefits are marginal or nonexistent: on HABERMAN, for instance, the random start already achieves 67.7% accuracy, and both warm starts slightly reduce performance to 66.1%. Overall, while warm starts sometimes improve the quality of the solution, they do not consistently dominate the random start variant across all benchmarks. Averaged across datasets, the best warm-started variant of Algorithm 1 improves upon the strongest baseline by 4.05 percentage points, where for each dataset the improvement is defined as the difference between the maximum result of the full SDP and  $3 \times 3$  minor warm-starts and the best baseline. This exceeds the 3.34-point gain of the random-start variant.

It is important to note that for SPAMBASE, the  $3 \times 3$ -SDP warm entry in Table 5 does *not* correspond to the true  $3 \times 3$ -SDP, but instead uses a randomized second-order cone (SOC) relaxation given by Problem 17. This substitution was necessary because both the full SDP and the  $3 \times 3$  SDP exceeded available memory on this large dataset. Interestingly, even though the SOC relaxation provides an extremely loose lower bound—far from the heuristic objective value (0.2 vs. over 700 as shown in Table 6 in Section 5.4)—warm-starting from the SOC solution still improves the test accuracy substantially, from 90.9% to 93.9%. This highlights that even relatively weak relaxations can offer valuable guidance for initialization when solving large-scale SMKL problems with Algorithm 1.

While warm starts can help improve predictive performance, they impose substantial computational burdens. Both relaxations become impractical for large problems because of the poor memory scalability exhibited

by interior point methods when used to solve semidefinite programs. In contrast, the randomly initialized variant of Algorithm 1 requires only a modest number of alternating best response steps, scaling efficiently even for datasets with thousands of examples. Despite its simplicity, the random start variant achieves competitive or superior predictive performance across all benchmarks without relying on solving costly semidefinite programs.

#### 5.4 Optimality certification via SDP relaxations (Q3)

Using the semidefinite relaxations introduced in Section 4, we can certify how close the solutions returned by Algorithm 1 are to a global optimum. We consider three levels of relaxation:

- **Second Order Cone (SOC):** Solve our sparse semidefinite relaxation with  $\ell = 2$  with standard basis and randomized vectors, given by Problem 17;
- **$3 \times 3$ -SDP:** Solve our sparse semidefinite relaxation for  $\ell = 3$ , given by Problem 18;
- **Full SDP:** Solve the dense semidefinite relaxation given by Problem 14.

Each of these yields a *lower bound*  $f_{LB}$  on the true optimal MKL objective. Given any SMKL feasible solution  $(\hat{\beta}, \hat{\alpha})$ , we compute its objective value  $F(\hat{\beta}, \hat{\alpha})$  and report the *relative optimality gap*

$$\text{Gap}(\hat{\beta}, \hat{\alpha}) = \frac{F(\hat{\beta}, \hat{\alpha}) - f_{LB}}{F(\hat{\beta}, \hat{\alpha})} \times 100\%,$$

which is an upper bound on how far  $(\hat{\beta}, \hat{\alpha})$  can be from a global minimizer.

Table 6 presents, for each dataset, the lower bounds from the SOC (standard basis and randomized), the  $3 \times 3$ -SDP and the full SDP relaxations, alongside the objective achieved by Algorithm 1 under three initialization strategies: random start, the  $3 \times 3$ -SDP warm start, and the full SDP warm start. As expected, the SOC bounds are weakest, the  $3 \times 3$ -SDP bounds are substantially tighter, and the full SDP bounds are the strongest wherever they could be computed. Moreover, warm-start initializations consistently reduce the Algorithm 1 objective: for example, on IONOSPHERE the random start gives 169.77, whereas the  $3 \times 3$  and full SDP starts lower it to 125.36 and 121.83, respectively. In small problems like WINE, where Algorithm 1 with random initialization is already nearly optimal (7.71 vs. a 7.70 full-SDP lower bound), these gains are marginal. On the largest datasets (SPAMBASE, MAMMOGRAPHIC), memory constraints precluded the  $3 \times 3$  and full SDP, so only the randomized SOC bounds are reported.

Table 6: Objective Values by Dataset. The line separate the relaxations (lower bounds) from the heuristic solutions (upper bounds).

Dataset	Train Size	SOC (basis)	SOC (rand.)	SDP (3x3)	SDP (full)	Alg 1 (rand.)	Alg 1 (3x3)	Alg 1 (full)
iris	120	0.50	1.04	7.68	<b>14.69</b>	<b>16.78</b>	<b>16.78</b>	<b>16.78</b>
wine	142	5.19	5.26	5.65	<b>7.70</b>	7.71	7.71	<b>7.70</b>
breastcancer	455	0.15	0.15	<b>1.06</b>	–	<b>70.86</b>	83.23	–
ionosphere	280	50.25	50.25	51.31	<b>80.39</b>	169.77	125.36	<b>121.83</b>
spambase	3680	<b>0.20</b>	<b>0.20</b>	–	–	<b>759.10</b>	–	–
banknote	1097	0.50	0.74	<b>6.08</b>	–	<b>42.13</b>	<b>42.13</b>	–
heart	242	0.19	0.33	1.75	<b>69.95</b>	<b>98.75</b>	<b>98.75</b>	<b>98.75</b>
haberman	244	0.51	1.56	63.15	<b>431.97</b>	476.80	<b>431.97</b>	<b>431.97</b>
mammographic	768	0.50	0.58	<b>450.26</b>	–	<b>1092.47</b>	1173.24	–
parkinsons	156	0.19	0.23	1.54	<b>25.69</b>	37.06	37.06	<b>26.31</b>

Table 7 reports the relative optimality gap for those datasets where the full SDP was tractable. On WINE and HABERMAN, the full SDP start drives the gap to zero (within solver tolerance), thereby certifying global

optimality. On PARKINSONS, the gap shrinks from 44.3% (random start) to 2.4% (full SDP start), a dramatic improvement. Even for the most challenging instance (IONOSPHERE), the  $3 \times 3$ -SDP halves the gap (111.2% to 55.9%), and the full SDP reduces it further to 51.6%. This demonstrates that our SDP hierarchy furnishes meaningful certificates of near-optimality for Algorithm 1 on real benchmarks.

Table 7: Worst-Case Optimality Gap (%) per Algorithm 1 Variant and Best Estimate

Dataset	Train Size	Random	3x3 SDP warm start	Full SDP warm start	Best gap estimate
iris	120	14.23	14.23	14.23	14.23
wine	142	0.13	0.13	0	0
ionosphere	280	111.18	55.94	51.55	51.55
heart	242	41.17	41.17	41.17	41.17
haberman	244	10.38	0	0	0
parkinsons	156	44.26	44.26	2.41	2.41

## 5.5 Summary of Findings

We are now in a position to answer the three questions introduced at the start of this section.

First, concerning **predictive accuracy (Q1)**, Algorithm 1 achieves high out of sampling accuracy while enforcing exact sparsity constraints. Across all ten UCI benchmarks, the randomly initialized variant of Algorithm 1 matches or outperforms the strongest open-source MKL baseline—EasyMKL, AverageMKL, and CKA—while selecting a small number of kernels from the candidate pool. On average, Algorithm 1 improves test set accuracy by 3.34 percentage points compared to the best competing method, with absolute gains reaching up to 7.6 percentage points on challenging datasets such as IONOSPHERE, HEART, and PARKINSONS. These results highlight that exact cardinality constraints, when appropriately enforced, not only preserve but can enhance predictive performance relative to dense kernel combinations.

Second, regarding the **effect of warm starts (Q2)**, initializing Algorithm 1 with solutions from SDP relaxations yields systematic benefits when computationally feasible. Warm starts based on the  $3 \times 3$  principal-minor or full SDP relaxations consistently reduce the final optimality gap—halving it in several cases—and can yield accuracy improvements exceeding 5 percentage points. However, solving the full SDP or even the restricted  $3 \times 3$  SDP relaxation becomes computationally infeasible for large datasets. In such cases, the random start version of Algorithm 1 remains highly competitive, achieving strong generalization performance without incurring the substantial overhead of solving costly semidefinite programs.

Third, regarding **optimality (Q3)**, our semidefinite relaxations certify that Algorithm 1 attains solutions that are near-optimal. On six datasets, the objective value achieved by Algorithm 1 matches the full SDP lower bound within solver precision, certifying global optimality. Even in the most challenging instances, the worst-case optimality gap remains below 52%, offering strong empirical evidence that SMKML consistently identifies high-quality solutions within the feasible set.

Taken together, these results establish that Algorithm 1’s exact-sparsity alternating best response strategy produces classifiers that are simultaneously interpretable, computationally efficient, and near-optimal. The method scales favorably to large datasets without sacrificing predictive accuracy, offering a practical and theoretically grounded solution to the sparse multiple kernel learning problem.

## 6 Conclusion

In this paper, we formulate the *sparse multiple kernel learning* problem with an exact cardinality constraint on the kernel weights and present an alternating best response algorithm (Algorithm 1) that produces high quality solutions efficiently. Moreover, we reformulate the sparse multiple kernel learning problem as a mixed integer semidefinite program and derive a hierarchy of convex relaxations that deliver strong lower bounds and effective warm starts for Algorithm 1. Our numerical results on ten UCI benchmarks demonstrate

that the randomly initialized variant of Algorithm 1 outperforms the open-source baselines AVERAGEMKL, EASYMKL, and CKA in out of sample prediction accuracy on average by 3.34 percentage points (relative to the best performing benchmark method) while selecting a small number of kernels in comparable runtime. Using the solution to our semidefinite relaxation as a warm start further shrinks the optimality gap of solutions returned by Algorithm 1 and yields additional out of sample accuracy gains of up to five percentage points.

These results illustrate that by imposing exact sparsity, our approach improves prediction accuracy while simultaneously providing interpretable kernel combinations and certificates of near-optimality. Future work includes developing a custom branch and bound algorithm that leverages the SDP relaxation given by Problem 14 to obtain globally optimal solutions to the MISDP given by Problem 12, and extending our framework to the regression setting.

## References

- Fabio Aioli and Michele Donini. Easymkl: a scalable multiple kernel learning algorithm. *Neurocomputing*, 169:215–224, 2015.
- Francis R. Bach. Consistency of the group lasso and multiple kernel learning. *Journal of Machine Learning Research*, 9:1179–1225, 2008.
- Francis R. Bach, Gert R. G. Lanckriet, and Michael I. Jordan. Multiple kernel learning, conic duality, and the smo algorithm. In *Proc. 21st International Conference on Machine Learning (ICML)*, pp. 41–48, 2004.
- Dimitris Bertsimas and Martin S Copenhaver. Characterization of the equivalence of robustification and regularization in linear and matrix regression. *European Journal of Operational Research*, 270(3):931–942, 2018.
- Grigoriy Blekherman, Santanu S Dey, Marco Molinaro, and Shengding Sun. Sparse psd approximation of the psd cone. *Mathematical Programming*, 191(2):981–1004, 2022.
- Bernhard E. Boser, Isabelle M. Guyon, and Vladimir N. Vapnik. A training algorithm for optimal margin classifiers. In *Proceedings of the 5th Annual ACM Workshop on Computational Learning Theory (COLT)*, pp. 144–152, 1992.
- Olivier Bousquet and André Elisseeff. Stability and generalization. *Journal of machine learning research*, 2 (Mar):499–526, 2002.
- Stephen Boyd, Laurent El Ghaoui, Eric Feron, and Venkataramanan Balakrishnan. *Linear matrix inequalities in system and control theory*. SIAM, 1994.
- Christopher J. C. Burges. A tutorial on support vector machines for pattern recognition. *Data Mining and Knowledge Discovery*, 2(2):121–167, 1998.
- Chih-Chung Chang and Chih-Jen Lin. LIBSVM: a library for support vector machines. *ACM Transactions on Intelligent Systems and Technology*, 2(3):27:1–27:27, 2011a.
- Chih-Chung Chang and Chih-Jen Lin. Libsvm: A library for support vector machines. *ACM Transactions on Intelligent Systems and Technology*, 2(3):27:1–27:27, 2011b.
- Corinna Cortes and Vladimir Vapnik. Support-vector networks. *Machine Learning*, 20(3):273–297, 1995.
- Corinna Cortes, Mehryar Mohri, and Afshin Rostamizadeh. Algorithms for learning kernels based on centered alignment. *Journal of Machine Learning Research*, 13:795–828, 2012.
- Dheeru Dua and Casey Graff. Uci machine learning repository. <http://archive.ics.uci.edu/ml>, 2017.
- Mehmet Gönen and Ethem Alpaydm. Multiple kernel learning algorithms. *Journal of Machine Learning Research*, 12:2211–2268, 2011.

- Oktaý Günlük and Jeff Linderoth. Perspective reformulation and applications. In Jon Lee and Sven Leyffer (eds.), *Mixed Integer Nonlinear Programming*, pp. 61–89, New York, NY, 2012. Springer New York. ISBN 978-1-4614-1927-3.
- Thorsten Joachims. Text categorization with support vector machines: Learning with many relevant features. In Claire Nédellec and Céline Rouveirol (eds.), *Machine Learning: ECML-98*, pp. 137–142, Berlin, Heidelberg, 1998. Springer Berlin Heidelberg. ISBN 978-3-540-69781-7.
- Marius Kloft, Ulf Brefeld, Soeren Sonnenburg, and Alexander Zien.  $\ell_p$ -norm multiple kernel learning. *Journal of Machine Learning Research*, 12:953–997, 2011.
- Vladimir Koltchinskii and Ming Yuan. Sparsity in multiple kernel learning. *Annals of Statistics*, 38(6): 3660–3695, 2010.
- Anastasios Kyrillidis, Stephen Becker, Volkan Cevher, and Christoph Koch. Sparse projections onto the simplex. In Sanjoy Dasgupta and David McAllester (eds.), *Proceedings of the 30th International Conference on Machine Learning*, volume 28 of *Proceedings of Machine Learning Research*, pp. 235–243, Atlanta, Georgia, USA, 17–19 Jun 2013. PMLR.
- Gert R. G. Lanckriet, Nello Cristianini, Peter Bartlett, Laurent El Ghaoui, and Michael I. Jordan. Learning the kernel matrix with semidefinite programming. *Journal of Machine Learning Research*, 5:27–72, 2004.
- Ailsa H. Land and Alison G. Doig. *An Automatic Method for Solving Discrete Programming Problems*, pp. 105–132. Springer Berlin Heidelberg, Berlin, Heidelberg, 2010.
- Ivano Lauriola and Fabio Aiolli. MKLpy: A python-based framework for multiple kernel learning. arXiv:2007.09982, 2020.
- Darrin P. Lewis, Tony Jebara, and William S. Noble. Support vector machine learning from heterogeneous data: An empirical analysis using protein sequence and structure. *Bioinformatics*, 22(22):2753–2760, 2006.
- John DC Little. *Branch and bound methods for combinatorial problems*. PhD thesis, MIT, 1966.
- Art B. Owen and Patrick O. Perry. Bi-cross-validation of the SVD and the nonnegative matrix factorization. *The Annals of Applied Statistics*, 3(2):564 – 594, 2009.
- Alain Rakotomamonjy, Francis R. Bach, Stéphane Canu, and Yves Grandvalet. Simplemkl. *Journal of Machine Learning Research*, 9:2491–2521, 2008.
- Albert Reuther, Jeremy Kepner, Chansup Byun, Siddharth Samsi, William Arcand, David Bestor, Bill Bergeron, Vijay Gadepally, Michael Houle, Matthew Hubbell, Michael Jones, Anna Klein, Lauren Milechin, Julia Mullen, Andrew Prout, Antonio Rosa, Charles Yee, and Peter Michaleas. Interactive supercomputing on 40,000 cores for machine learning and data analysis. In *2018 IEEE High Performance extreme Computing Conference (HPEC)*, pp. 1–6. IEEE, 2018.
- Bernhard Schölkopf and Alexander J. Smola. *Learning with Kernels: Support Vector Machines, Regularization, Optimization, and Beyond*. MIT Press, Cambridge, MA, 2002.
- Dogyoon Song and Pablo A Parrilo. On approximations of the psd cone by a polynomial number of smaller-sized psd cones. *Mathematical Programming*, 198(1):733–785, 2023.
- Soeren Sonnenburg, Gunnar Raetsch, Christin Schafer, and Bernhard Schölkopf. A general and efficient multiple kernel learning algorithm. In *Advances in Neural Information Processing Systems 18 (NIPS)*, pp. 1273–1280, 2006.
- Ingo Steinwart. On the influence of the kernel on the consistency of support vector machines. *Journal of machine learning research*, 2(Nov):67–93, 2001.
- Niranjan Subrahmanya and Yung C. Shin. Sparse multiple kernel learning for signal processing applications. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 32(5):788–798, 2010. doi: 10.1109/TPAMI.2009.98.

Robert Tibshirani. Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society: Series B (Methodological)*, 58(1):267–288, 1996.

Hui Xue and Yu Song. Non-convex approximation based  $\ell_0$ -norm multiple indefinite kernel feature selection. *Applied Intelligence*, 49(11):4012–4025, 2019.